Hungry Clouds: Labour Market and Environmental Externalities from Data Centers

Eddy (Weijian) Zou*

First Draft: June 2025
This Draft: August 5, 2026

Abstract

Data centers are critical digital infrastructure whose electricity demand imposes environmental costs that vary sharply across space, and whose local welfare impacts are highly contested. This paper asks whether existing place-based tax incentives induce spatial misallocation of data centers and if alternative policies could induce welfare-improving location choices. Using an event-study approach, I find that entry raises county employment by 7.8% of its pre-entry mean over ten years, driven by professional services, real estate, administrative services, and construction sectors. Data center entry also induces net in-migration, after controlling for their impacts on local real wages. To evaluate alternative policies, I build and estimate a quantitative spatial equilibrium model featuring partially tradable compute, capacity-constrained electricity markets, pollution transport, and endogenous data center siting. Using the model, I find that the existing announced capacity expansion raises measured welfare by $4.98 billion per year (0.214% of national income) but concentrates gains in host counties (+0.40% against −0.98% elsewhere) and increases national emissions by 685 million tons of CO_2 per year. Removing state-level tax incentives raises national welfare by $9.99 billion per year (0.492%), cuts local pollution damage by $2.18 billion, and lowers CO_2 emissions by 56 million tons, as data center capacity gets reallocated towards alternative locations.

Keywords— Spatial Environmental Policies; Computing; Data Centers; Distributional Impacts
JEL codes— Q41; R11; R12; H25; L86

*London School of Economics and Political Science & Institute for Fiscal Studies, UK. Email: w.zou2@lse.ac.uk.
For their invaluable encouragement and guidance, I thank Clare Balboni, Oriana Bandiera, Robin Burgess, and Daniel Sturm. For helpful comments, I thank participants at LSE Environment and Energy (EEE), Industrial Organisation Seminar, LSE Energy Markets Miniconference, LSE Environment Camp, 2026 PEDL Funding Workshop, Queen Mary PhD Conference. I gratefully acknowledge the Private Enterprise Development in Low-Income Countries (PEDL) Initiative and STICERD for funding this work.

1 Introduction

The production of goods, services, and new ideas has become data-intensive (Heath 2019; Goldfarb and Tucker 2019; Beraja et al. 2023; Demirer et al. 2024). Since 2010 the number of internet users worldwide has more than doubled and global internet traffic has grown twenty-five-fold (IEA, 2024). The facilities that store and process that data have in turn become one of the fastest-growing loads on the electricity system: US data center consumption rose from 76 TWh in 2018 to 176 TWh in 2023, or 4.4% of national electricity use, and is projected to reach 325–580 TWh by 2028 (Shehabi et al. 2024).[1] The electricity demand is increasingly spatially concentrated. In 2023 data centers took roughly 26% of the electricity sold in Virginia, 15% in North Dakota, and 11–12% in Nebraska, Iowa, and Oregon (Pew Research Center 2025).

The environmental footprint of these facilities is large and varies substantially across space. Within the United States the carbon intensity of a megawatt-hour delivered to a data center ranges from 0.02 to over one ton of CO_2, and water use from 1.8 to 106 m^3, depending on grid composition, climate, and cooling technology (Siddik et al. 2021; Goiri et al. 2011; Holland et al. 2024). Operators choosing sites to minimise private cost may fail to internalise such unpriced externalities, so the market allocation of capacity across grids need not be the socially efficient one. Communities have been increasingly vocal towards the resulting siting decisions: local opposition has blocked or delayed $64 billion of announced US data center projects (Data Center Watch, 2024), and the Netherlands, Ireland, Singapore, and Spain have each imposed temporary moratoria on new developments.[2]

Place-based policies for data centers must account for both their local economic effects and their environmental externalities. Their net impacts on local economies are theoretically ambiguous: on the one hand, they could increase electricity market congestion and cause higher electricity costs for local households and firms: utilities in the PJM interconnection requested over $29 billion in rate increases in the first half of 2025, an increase attributed in part to data center demand.[3] On the other hand, large, concentrated capital investments by data center entry could generate agglomeration economies. In addition to local labour demand shocks for related industries (Greenstone et al. 2010), data centers could generate new agglomeration spillovers through the computation networks. By lowering the effective cost of computation, a nearby data center can benefit service firms whose output degrades with distance to the processing facility, since network latency rises with every additional kilometre of fibre (Moallemi and Sağlam 2013; Hasbrouck and Saar 2013). The magnitude of agglomeration spillovers and congestion effects, along with the spatially differentiated environmental impacts, will determine the local and aggregate welfare effects of data center siting.

This paper measures those effects and evaluates the policies that shape where data centers locate in the United States, the world's largest data center market. I combine causal evidence on the sector-specific local effects of entry with a quantitative spatial equilibrium model in which electricity markets clear locally, computation is traded across space subject to latency, air pollution is transported from generating plants to the counties that breathe it, and data center locations are chosen endogenously. I use the model to quantify the welfare consequences of the announced expansion of US data center capacity, and to evaluate the most widely used instrument governing its geography: state-level tax incentives.

The analysis rests on a county panel that matches proprietary facility-level records of data center entry from 1999 to 2023 (S&P Global 2025) to administrative data on employment, establishments, wages, retail electricity markets, telecommunications and power infrastructure, migration, and pollution damages. I assemble four novel datasets to further discipline my quantitative framework. First, I measure data-

[1] See Shehabi et al. (2024) and https://www.bloomberg.com/graphics/2024-ai-data-centers-power-grids/.

[2] Sources: https://www.datacenterwatch.org/report; https://www.techmonitor.ai/hardware/cloud/inside-the-data-centre-moratorium-movement.

[3] Source: https://www.eesi.org/articles/view/data-center-power-demands-are-contributing-to-higher-energy-bills.

transmission quality directly, using daily server-pair round-trip times from a global monitoring platform, together with the network of long-haul terrestrial fibre routes from Infrapedia and the fibre node locations of the Internet Topology Zoo (Knight et al. 2011). Second, I observe where firms actually put their data: starting from a registry of 640,000 US service-sector firms, I run a multi-vantage-point measurement campaign that geolocates the web servers behind each firm's domain, retaining direct-origin hosting, which yields novel measures of service firms' choice of data centers and trade flows for data center services. Third, I use a metro-by-provider-by-year panel of colocation prices for 2014–2025, which gives the price of computation on the supply side. Fourth, I compile a panel of state data center tax incentives covering all fifty states from 1980 to 2024, in which twenty-three states adopt a programme between 2005 and 2022.

I begin by asking what predicts the location choices of data centers. Telecommunications infrastructure dominates: broadband connection density, the number of available providers, and proximity to fibre nodes are the strongest predictors of first entry, and counties in the top half of the transmission-latency distribution are significantly less likely to host a facility. Proximity to customers matters as well. A one standard deviation increase in the share of local establishments in ICT raises the entry probability by about 0.08 percentage points, against a baseline entry rate of 0.04%. Higher commercial and industrial electricity tariffs deter entry, and counties in states offering tax incentives are significantly more likely to receive facilities. I consistently find no significant correlations between emissions damage of local electricity generation and data center siting, suggesting that operators' private decisions do not internalise the environmental cost of the electricity the facilities consume.

I then estimate the causal effect of entry on local economies using two strategies that exploit different variation. The primary estimates come from the staggered event-study design of Callaway and Sant'Anna (2021), which recovers dynamic effects of first county entry using not-yet-treated and never-treated counties as controls and is robust to treatment-effect heterogeneity across cohorts. I show robustness to an alternative approach that uses the synthetic difference-in-differences estimator of Arkhangelsky et al. (2021), which reweights control counties to match pre-entry trends.

I find that data center entry generates economically significant, sector-specific local effects that extend well beyond the construction phase, in contrast to the view of data centers as low-employment "server farms" with little local footprint (Joint Legislative Audit and Review Commission 2024). First, treated counties do not exhibit significant pre-entry trends relative to not-yet-treated ones, supporting the validity of the parallel trends assumption. Data center entry causes a progressive increase in county employment over a ten-year post-entry horizon, at about 8% of the pre-entry mean. This increase is driven by significant gains among a range of service sectors, including professional services, real-estate, administrative services, information, and construction services. In contrast, data center entry causes a significant decrease in the number of manufacturing establishments, consistent with an electricity-congestion mechanism.

Next, I estimate how data center entry affects county-to-county migration flows, using the IRS Statistics of Income for 1999–2021. I find that net migration rate rises by 0.0029 after data center entry, at 42% of the baseline rate. The response is driven by colocation facilities, which raise the in-migration rate by 20.9% of the control mean, while hyperscale entry has a small and imprecise effect. Controlling for their impacts on real wages, I find that data center entry induces net increase in migration, suggesting that it has positive net amenity impacts on local residents.

Reduced-form estimates cannot fully capture the general-equilibrium adjustments that determine national welfare: workers and firms re-sort across counties, electricity and computation prices move, and emissions are transported to populations far from the plants that produce them. I therefore develop a quantitative spatial equilibrium model, building on Becker et al. (2021); Arkolakis and Walsh (2023); Quentel (2023); Hollingsworth et al. (2026). Monopolistically competitive firms in tradable, high-computation service, and low-computation service sectors hire workers who choose location and sector subject to Fréchet-

distributed idiosyncratic tastes and switching costs. Three features are added for this setting. First, computation is a produced input that is partially tradable: service firms can buy from data centers anywhere, but its delivered quality degrades with network latency, so local entry lowers the local price of computation. Second, electricity markets clear at the balancing-authority level along an upward-sloping supply curve, so data center entry could generate pecuniary externalities on electricity tariffs across space. Third, data center locations are chosen rather than assumed: entrants trade off market access through the latency network, operating costs through electricity prices and climate-driven cooling requirements, and the tax treatment they face, with idiosyncratic location payoffs generating a logit allocation weighted by county suitability. The equilibrium allocation of households, firms, and facilities then determines electricity demand, the set of generating plants dispatched to serve it, and the resulting distribution of air pollution and carbon emissions.

I calibrate the model to the 2024 cross-section of data center locations, sectoral employment, and wages, recovering location fundamentals by inversion. The model is able to reproduce the spatial allocation of announced entrant capacity and the spatial variation in electricity demand.

Using the model, I first quantify the aggregate welfare effects of the current proposed expansion of data center capacity across the US, taking location and size of the planned pipeline as given. The expansion raises aggregate welfare by \$4.98 billion per year (0.214% of national income). This conceals a sharp split: workers in counties that host new capacity gain 0.395% of their baseline income, while workers in non-host counties lose 0.976% of their income. Decomposing the aggregate effects, I find that the amenity gain of \$17.215 billion accrues only where facilities are built, whereas the pollution cost of \$6.266 billion and the real-wage cost of \$6.283 billion fall on host and non-host counties alike, because electricity prices rise across the grid by 8.5% on average (\$5.34/MWh) and emitted pollutants travel. Turning to global pollutants, the expansion adds 685.1 million metric tons of CO_2, 14% of the 4,904 million tons of energy-related US emissions in 2025.[4]

I then hold the total announced capacity fixed and remove state-level data center tax incentives, rebating the recovered revenue to workers as a state lump sum. Removing the incentives raises national welfare by \$9.985 billion per year, 0.492% of income and roughly double the gain from the expansion itself: \$6.961 billion through amenities, \$2.177 billion through reduced pollution damage, and \$0.208 billion through real wages. Capacity moves off the dirtier Southeastern grids where the abatements had concentrated it and onto cleaner Western ones, which is why the pollution channel turns positive here having been negative under the expansion, and national CO_2 emissions fall by 56.2 million tons, about 1.1% of energy-related emissions, worth a further \$8.1 billion per year. Counties that gain capacity gain 0.575% of income, and counties that lose it lose only 0.071%, as the loss in amenity values outweighs the impacts of lower electricity tariffs.

Finally, I highlight the importance of accounting for multiple general equilibrium channels, by conducting the same policy experiment, each time removing one channel of adjustment. I find that ignoring any single channel of (1) electricity market congestion; (2) computation spillovers; (3) local environmental damages, leads to overvaluing the welfare effects of tax incentive removals.

The results have several implications for place-based policies towards data centers. First, local opposition to data center entry reflects a real distribution of costs and not only environmental preferences: the gains are concentrated in service firms and their workers, while higher electricity prices and additional pollution are borne by households and firms across the whole grid, particularly in non-host counties. As the aggregate effects of data center siting are determined by both the electricity and compute markets, place-based incentives can benefit from accounting for the specific spatial incidence of such spillovers.

[4]The 2025 US total is from the EIA, https://www.eia.gov/environment/emissions/carbon/. Emissions are monetised at the EPA social cost of carbon of \$145 per metric ton throughout.

Related literature. This paper speaks to three literatures. The first develops quantitative spatial models of trade, energy, and environmental externalities (Caliendo et al. 2019; Arkolakis and Walsh 2023; Hollingsworth et al. 2025, 2026; Colas and Reynier 2026). I develop a new framework that embeds computation as a partially tradable input, with the degree of tradability estimated from where firms actually host their data, and the model's equilibrium pollution determined by plant-level dispatch model and a dispersal model, so that a change in the geography of electricity consumption maps into a spatial distribution of damages. The framework therefore jointly accounts for new general equilibrium channels through which data center siting affects wrlfare. The second literature studies place-based fiscal policy and the incidence of taxes and subsidies aimed at firms and households (Slattery and Zidar 2020; Gargano and Giacoletti 2025; Greenstone et al. 2010). My contribution is to evaluate such policies when the externalities they create travel through electricity and computation networks rather than through on-site emissions, and the channel-shutdown exercise above shows that accounting for both networks changes the estimated benefit of reform by a factor of roughly three. The third literature measures the economic and environmental effects of data centers directly (Muller 2025; Alvarez et al. 2026; Kay et al. 2026; Knittel et al. 2025; Siddik et al. 2021). Relative to that work I provide causal estimates of the sector-specific local labour-market effects of entry, evidence that market access and network latency shape location choice alongside operating costs, and a general-equilibrium evaluation of the policies that govern siting rather than a description of its consequences.

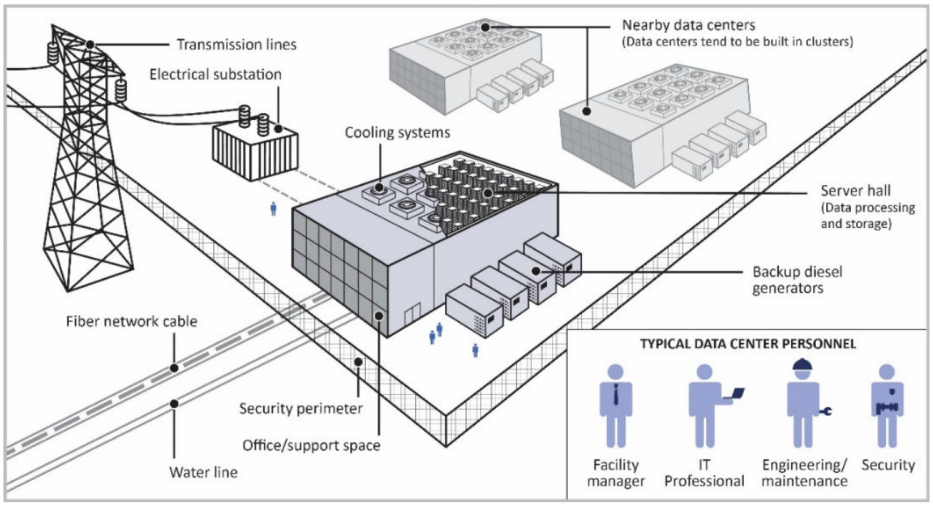
The remainder of the paper is organised as follows. Section 2 provides background on the US data center industry. Section 3 describes the data. Section 4 presents the evidence on location choice and sets out the two identification strategies. Section 5 reports the estimated local effects of entry. Section 6 develops the quantitative spatial model, describes the estimation of its key parameters and the pollution dispersal block, and presents the counterfactuals. Section 7 concludes.

2 Background

Data centers are facilities that store, manage, process, and transmit large amounts of digital data. They support virtually all core digital services for households and firms, including website hosting, video and audio streaming, file storage, cloud computing, and the training and deployment of Artificial Intelligence (AI) models. Facilities vary widely in scale and function. *Colocation* (or "retail") data centers are operated by third-party providers who rent space and power to multiple tenants; *enterprise* data centers are owned and operated by a single organisation for its own use; *cloud* (or "wholesale") data centers are large-scale facilities operated by providers such as Amazon Web Services, Microsoft Azure, and Google Cloud to deliver on-demand computing services; and *hyperscale* data centers are very large, purpose-built facilities (typically exceeding 100 MW of power capacity) designed to support the infrastructure needs of major cloud and AI providers.[5] My baseline treatment definition pools these categories, taking the earliest first entry of a colocation, enterprise, or hyperscale facility in a county; telecom exchanges are excluded throughout. I report treatment-effect heterogeneity by facility type, separately for colocation and hyperscale cohorts, in Section 5.

[5] A fifth category, *edge* data centers, consists of smaller distributed facilities located close to end-users to minimise latency for applications such as the Internet of Things and real-time video streaming.

Figure (1) Schematics of a Data Center. Source: Joint Legislative Audit and Review Commission (2024)



As shown in Figure 1, a typical data center is an industrial building housing computing equipment (servers, storage drives, network hardware), supported by dedicated electrical infrastructure—often including direct connections to nearby high-voltage transmission lines and on-site power substations. On-site backup generators ensure continuity of operations during grid outages. Data centers also require substantial cooling systems, which account for approximately 38–40% of a facility's total electricity consumption (Congressional Research Service 2025). Data centers typically employ a skilled but small workforce relative to their physical footprint: technicians, electricians, network engineers, and security personnel. In Virginia, the largest US cluster of data centers, an average facility occupies approximately 23,000 square meters and employs around 50 workers (Joint Legislative Audit and Review Commission 2024).[6] Nationally, total data center employment reached approximately 500,000 workers by 2023, a 60% increase from 306,000 in 2016 (U.S. Census Bureau 2024). While direct employment per facility is modest, data centers generate substantial indirect employment during construction and through local demand for ancillary services.[7]

2.1 The United States Data Center Industry

The United States hosts the largest concentration of data centers globally, accounting for roughly 40% of the world's data center infrastructure.[8] The industry has grown rapidly in both physical scale and economic significance: total data center electricity consumption increased from approximately 60 TWh in 2016 to 176 TWh by 2023, representing 4.4% of total US electricity consumption (Shehabi et al. 2024). By 2024, the International Energy Agency estimates this figure rose to 183 TWh (International Energy Agency 2025), with projections suggesting it could reach 325–580 TWh by 2028—between 7% and 12% of total US electricity demand (Shehabi et al. 2024).

Three waves of construction. The expansion of the US data center industry can be broadly divided into three periods. The first wave coincided with the internet and dot-com booms of the late 1990s and early

[6] This figure masks considerable heterogeneity: smaller retail colocation facilities may employ 15–25 workers, while large hyperscale campuses can employ 100–200 permanent staff.

[7] The Stargate data center project in Texas, for example, projects peak construction employment of 5,000–10,000 workers and approximately 600 permanent operational jobs once fully built.

[8] Estimates vary by source and definition. Cloudscene reports 5,427 facilities; other sources report slightly different totals depending on whether edge facilities and smaller installations are included.

2000s, which led to the initial large-scale development of retail colocation facilities—often in major connectivity hubs such as lower Manhattan (e.g., 60 Hudson Street), Northern Virginia, and the San Francisco Bay Area.

The second wave, spanning roughly 2010–2019, was driven by the emergence and rapid scaling of cloud computing. Following the launch and expansion of Amazon Web Services (beginning 2006), Microsoft Azure (2010), and Google Cloud Platform (2012), demand shifted from smaller, distributed retail facilities toward larger wholesale and cloud-dedicated data centers. This period saw the rise of Northern Virginia as the dominant US market, facilitated by Virginia's sales tax exemption for data center equipment (enacted in 2009), abundant and relatively affordable electricity from the PJM interconnection, and proximity to the federal government and major internet exchange points.

The third and current wave, beginning around 2020 and accelerating sharply since 2023, is characterised by the construction of hyperscale facilities to support the training and real-time deployment of large language models (LLMs) and other AI applications, alongside continued growth in cloud computing and edge infrastructure. The scale of current investment is unprecedented: according to CBRE, the amount of data center capacity under construction in primary North American markets reached a record 6,350 MW at year-end 2024—a twelve-fold increase from 457 MW in 2020 (CBRE 2024). Quarterly new lease and pre-lease activity in the US has grown from below 100 MW in 2016 to below 200 MW in 2020 to over 1,000 MW by 2024 (CBRE 2024).[9]

Macroeconomic significance. The scale of data center and related AI infrastructure investment has become macroeconomically significant. Using Bureau of Economic Analysis data, Furman (2025) estimates that investment in information-processing equipment and software—though only 4% of GDP—accounted for 92% of US GDP growth in the first half of 2025.[10] Major hyperscale operators have committed unprecedented capital: Meta allocated $65 billion, Microsoft $80 billion, and Amazon over $30 billion for 2025 data center investments alone, with aggregate hyperscaler capital expenditure approaching $400 billion annually (CBRE 2024).

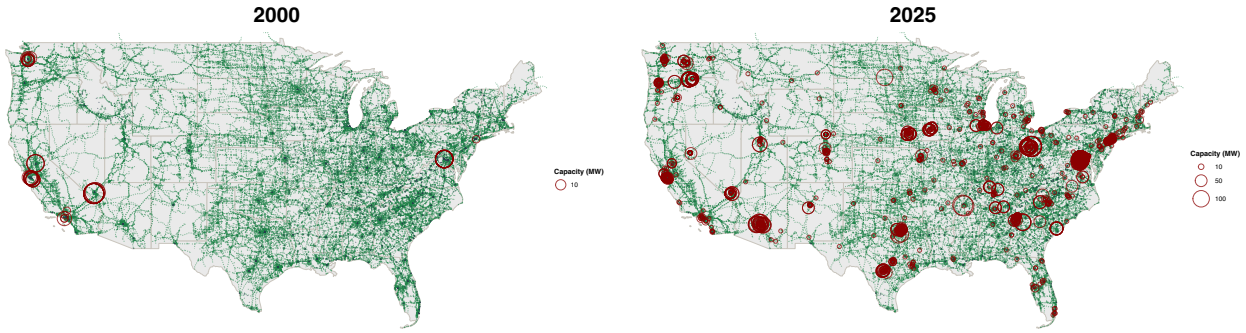
Spatial concentration. Data center activity is highly spatially concentrated, and its footprint has tracked the infrastructure that supports it. Figure 2 plots operational facilities against the terrestrial fibre and high-voltage transmission network in 2000 and 2025. In 2000 a small number of facilities sat in a few coastal connectivity hubs and around Northern Virginia; by 2025 the industry had expanded enormously, filling in along the same fibre and transmission corridors rather than dispersing across space. Northern Virginia remains the dominant market, with approximately 2,930 MW of inventory at year-end 2024, followed by Dallas–Fort Worth, Atlanta, Phoenix, and the Chicago metropolitan area (CBRE 2024). This concentration produces large local loads: in 2023 data centers consumed approximately 26% of total electricity supply in Virginia, 15% in North Dakota, and 11–12% in Nebraska, Iowa, and Oregon (Pew Research Center 2025). The clustering of such loads in specific counties and utility service territories is what generates the variation my empirical analysis exploits.

[9] Primary vacancy rates fell to a record low of 1.9% at year-end 2024, declining across all eight primary markets tracked by CBRE for the first time since tracking began in 2013.

[10] A related estimate by S&P Global, using final private domestic demand, attributes 80% of the increase in private domestic demand in H1 2025 to data centers and related high-tech spending (S&P Global 2025).

Figure (2) Data Center Entry and Network Infrastructure, 2000 and 2025

(a) 2000 (b) 2025



Notes: Operational US data center facilities (red circles, sized by installed power capacity in MW) plotted over the terrestrial fibre and high-voltage transmission network (green), in 2000 (Panel a) and 2025 (Panel b). The comparison shows the growth of the industry over the sample period and the co-location of facilities with pre-existing fibre and transmission corridors.

3 Data

I assemble six datasets for the analysis: (i) a facility-level dataset of US data center locations and characteristics; (ii) a county-year panel of retail electricity market outcomes; (iii) a county-year panel of establishment counts and employment by sector; (iv) county-level measures of telecommunications infrastructure, demographics, and state-level data center tax incentives; (v) a firm-level dataset of where US service-sector firms host their data; and (vi) plant-level generation, emissions, and pollution-transport data used to compute local environmental damages. This section describes the construction of each dataset and discusses key measurement issues. Table A.10 presents summary statistics for the main county-year panel.

3.1 Data Center Facilities

I use a proprietary facility-level dataset from S&P Global covering all operational, planned, and under-construction data center facilities globally. For each US facility operational between 1990 and 2025, I observe facility name, geographic coordinates, and the affiliated firm. I supplement this dataset with a facility-level dataset from ScholarEdge, which uses a combination of OpenStreetMap records, FCC regulatory filings, Baxtel and other commercial platforms' open records, to measure the facility size, installed capacity, estimated Power Usage Effectiveness (PUE), construction year, and specific facility type. I harmonise the two datasets and use variables on data center type, installed capacity, PUE, year of operation in my empirical analysis. Appendix Figure A.14 plots the growth of installed capacity by facility type. Total installed capacity (excluding telecom exchanges) rose from about 1.4 GW in 2000 to 36.5 GW in 2025; colocation and enterprise facilities account for roughly 13 GW of the 2025 total, with hyperscale facilities making up the remainder.

The dataset covers 1,399 unique data center entries and 592 county-level first-entry events between 2000 and 2023. These 592 events are the treatment cohorts for the staggered difference-in-differences design, and I use this window throughout; where an analysis requires lagged covariates it uses the subset of events for which those covariates exist, and I state the effective sample in each case.

Figure 2 presents the spatial distribution of operational data centers in 2000 and 2025. By 2010, 601 facilities were operational. The pace of new entry was approximately 357 facilities per five-year period

through 2020, before accelerating to 411 new facilities between 2021 and 2025, reflecting the recent surge in demand for cloud computing and AI workloads.

For the counterfactual analysis, I supplement the historical S&P Global records with announced post-2024 project data from CleanView, a commercial data center intelligence platform. CleanView tracks planned and under-construction projects reported in planning applications, press releases, and regulatory filings. I use the expected-build scenario, which covers approximately 150 GW of announced US data center capacity at reported county-level locations.

3.2 Retail Electricity Markets

I construct a county-year panel of retail electricity market outcomes from 2000 to 2024 using data from EIA Form 861. The raw data report utility-level information on retail sales (MWh), revenues, number of customers, and average retail price, separately for industrial, commercial, and residential customers. I match each utility to the US counties it serves using the annual EIA utility-county concordance, and aggregate to the county-year level by computing unweighted averages across utilities operating in each county.[11]

3.3 Firm and Worker Outcomes

I measure local economic activity using two complementary sources. My primary dataset is the Business Dynamics Statistics (BDS), from which I construct a county-year panel of the number of operational establishments per 10,000 population and total employment, by two-digit NAICS sector, from 1999 to 2023. The BDS has near-complete coverage of private-sector establishments and allows me to measure sectoral composition at the county level over the full sample period.

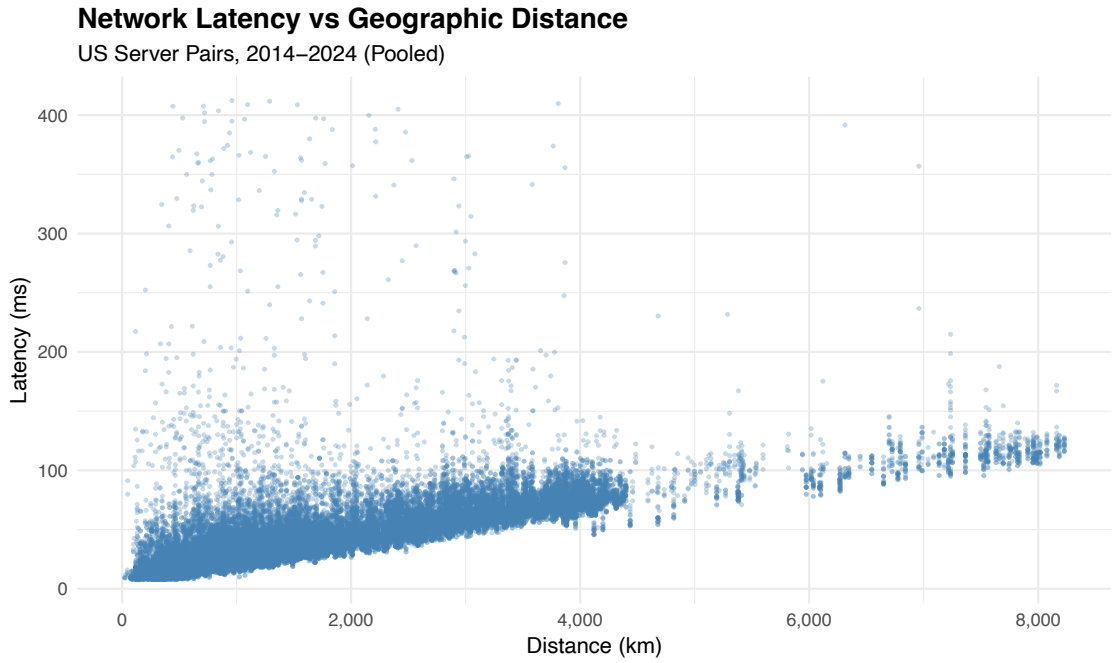
I supplement the BDS data with the Quarterly Census of Employment and Wages (QCEW) for 2000–2024, which provides employment, establishment counts, and total wages at the county-by-three-digit-NAICS level. The finer sectoral detail in the QCEW allows me to distinguish, for example, between electricity-intensive and non-intensive manufacturing subsectors when testing mechanisms.

3.4 Infrastructure, Demographics, and Policy

Telecommunications infrastructure. I measure broadband infrastructure using two FCC Form 477 sources. For 2009–2023, I use county-level data on the number of broadband connections by speed tier. For the earlier period (1999–2007), when county-level connection counts are unavailable, I construct a proxy from zip-code-level FCC data measuring the number of companies providing high-speed internet service, aggregated to the county level. The early-period measure serves as a pre-determined predictor of data center location and as a component of the suitability index used in the shift-share identification strategy (Section **??**).

[11] In my sample, 2,249 counties are served by a single utility, while 1,054 are served by more than one. I verify that results are robust to weighting electricity tariffs by the number of customers served by each utility.

Figure (3) Scatterplot, Server-Pair Latency and Distance



Data Transmission Latency. I construct latency measures from a server-monitoring platform that records daily round-trip ping times between a global network of servers. The first measure uses the full panel of 276 servers in 103 countries, covering 2010 to 2024. I aggregate the server-pair level data using distance-based interpolation to obtain average latency from each US county centroid to all other servers, then take an unweighted county-to-county average to produce a county-year latency measure. I use the long-run average (2010–2024) as a predetermined predictor of data center suitability and as a component of the shift-share instrument (Section **??**). Figure 3 plots server-pair latency against distance in kilometres and shows a clear positive relationship, consistent with the physical propagation constraints of fibre-optic transmission. The second measure restricts to 83 US servers for 2014–2024 and estimates a bilateral gravity equation on server-pair latency to construct the county-pair network quality index $\hat{\Phi}_{\ell,c}$ used in the quantitative spatial model (Section 6.9.1).

Fibre optic network. I use two sources on the physical fibre network. The location of fibre nodes comes from the Internet Topology Zoo (Knight et al. 2011), a curated collection of data-transmission network topologies. From these I compute the distance from each county centroid to the nearest node, which proxies for the cost of connecting to backbone infrastructure and enters the site-selection analysis and the suitability index. For the bilateral computation-quality index in the quantitative model (Section 6.9.1), I use the network of long-haul terrestrial cable routes from Infrapedia and compute the shortest-path distance between locations through this network by Dijkstra's algorithm.[12] Appendix Figure A.15 plots the terrestrial cable network. Because both sources are largely time-invariant over the sample, I treat fibre connectivity as a pre-determined county characteristic.

Demographics. County-level demographic controls—share of college graduates, total population, and median income—come from the 1990 and 2000 decennial censuses. I use 1990 values as pre-determined

[12]Infrapedia (https://www.infrapedia.com) maps global internet infrastructure, including terrestrial and submarine fibre-optic cable routes.

characteristics in the suitability index and as controls in the shift-share design, and 2000 values for the site selection analysis.

County-to-County Migration. I use the Internal Revenue Service (IRS) Statistics of Income county-to-county migration files covering 1990 to 2022. These files record, for each origin-destination county pair, the number of personal income tax returns filed and the number of exemptions by households that changed their county of residence between consecutive tax years. I aggregate to annual bilateral migration flows at the county level and use the 2017–2022 average to calibrate the observed migration matrix π^{obs} in the Fréchet migration framework. The longer 1990–2022 panel is used to estimate the amenity spillover β_B from the bilateral migration equation of Section 6.9.4, which is estimated by OLS with origin-by-destination pair and year fixed effects. The event-study estimates of the migration response in Section 5 use the 1999–2021 window, over which the destination-county aggregates are complete.

State-level tax incentives. I compile a novel dataset of state-level tax incentive programmes targeting data center construction and operations, covering all 50 states from 1980 to 2024. For each state, I record the year of first adoption, the type of incentive (sales tax exemptions, property tax abatements, income tax credits), and whether the programme remains active. Twenty-three states adopted data center-specific incentives during the sample period, with adoption dates ranging from 2005 to 2022. This dataset forms the basis of the shift-share instrumental variable strategy described in Section **??**.

Data Center Choice of Service-Sector Firms. To measure the hosting choices of US service-sector firms, I assemble a dataset linking firm headquarters to data center locations via IP address geolocation. The starting point is the Revelio Labs firm registry, which provides employment records for approximately 640,000 US service-sector firms (NAICS 51–62), including employer name, NAICS code, headquarters city, and corporate website. For each firm, I run a multi-vantage point measurement campaign to identify the city of IP addresses associated with web servers of each firm. I exclude firms using CDN and Anycast traffic (Akamai, Cloudflare), which represents approximately 25% of the raw sample and does not correspond to a fixed hosting city. The final dataset contains 625,163 firm-headquarters-to-hosting-city pairs, which I use to estimate the elasticity of substitution η across data center source locations in the gravity specification of Section 6.9. Further details on the IP geolocation pipeline are in Appendix B.5.

3.5 Environmental Damage Data

Plant-level generation and emissions. I use the Environmental Protection Agency's Continuous Emissions Monitoring System (CEMS) to obtain plant-level annual generation (MWh) and emission rates for $PM_{2.5}$, nitrogen oxides (NO_x), and sulfur dioxide (SO_2). CEMS records hourly generation and emissions for all regulated power plants, covering 2000 to 2024. I supplement CEMS with EIA Forms 860 and 923 for plant-level capacity, fuel type, and location, and use the plant-county crosswalk to aggregate emissions to the county level.

AP3 source-receptor matrices. Local health damages are computed using the AP3 Integrated Assessment Model (Muller, 2025), which provides county-level source-receptor (SR) matrices for $PM_{2.5}$, NO_x, and SO_2. The matrices are 3,109 $\times$ 3,109 (all US counties), and each entry $SR_{p,c,c'}$ gives the marginal health damage in dollars per metric ton of pollutant p emitted in source county c' and borne by residents of receptor county c, accounting for atmospheric dispersion, population exposure, dose-response relationships, and the value of a statistical life.

4 Empirical Strategy

4.1 Determinants of Data Center Location Choice

To understand which local characteristics predict data center entry, I estimate linear probability models examining pre-entry levels and changes in county characteristics. This descriptive analysis tests whether data center operators consider private operational costs, proximity to clients, and local infrastructure quality before making location decisions.

4.1.1 Empirical Approach

For each county receiving a data center in year k, I construct pre-entry averages ($\bar{X}_d$) and changes (ΔX_d) in local characteristics, comparing these against counties that never experienced data center entry (never treated counties). I then form a stacked cross-section of all treatment cohorts, estimating a linear probability model with treatment cohort fixed effects:

$$1[\text{Treated in Year } k]_d = \Gamma_1 \bar{X}_d + \Gamma_2 \Delta X_d + \mu_k + u_d \tag{1}$$

This approach increases statistical power by enabling individual tests of each entry event, and the stacked structure flexibly accommodates staggered treatment timing. The estimation sample is the subset of the 592 first-entry events from 2010 onward, for which the lagged FCC broadband covariates are available. I separately estimate models with lagged levels (2 years pre-entry), lagged changes (4 to 2 years pre-entry), and both specifications pooled together. I measure determinants across five conceptual categories:

Private Operational Costs. I use county-level industrial and commercial electricity tariffs to measure variable cost of data center operations.

Workforce Characteristics. College graduate share captures the availability of skilled workers.

Client Proximity. I measure the presence of potential client industries through county employment shares in finance, insurance, real estate, and information technology sectors, as proximity to major service sector employers may reduce latency and improve service delivery.

Infrastructure Quality. I measure digital infrastructure using county-level broadband connections and the distance of county centroid to the nearest fiber optic node. For electricity infrastructure, I include transmission line density and the density of electricity substations as measures of transmission capacity and reliability.

Environmental Externalities and Policies. I include county-level estimates of environmental damage from electricity generation, which account for local and global pollutants using power plant emissions factors and pollution dispersal models. To test the role of state-level regulatory policies, I construct an indicator for whether a state-level tax incentive for data centers was in place within five years of entry.

Network Quality. Following industry practice, I measure data transmission latency using quartile indicators for the average latency between the county centroid and all other county centroids.

Given the large number of predictors included, I complement the linear probability models with LASSO regressions that use cross-validation to select relevant predictors of data center entry, appropriately penalising inclusion of irrelevant predictors.

4.1.2 Results: Determinants of Data Center Site Selection

Figure 4(a) presents coefficients from the specification that only includes pre-entry average levels of characteristics. I present both estimates from linear probability models and LASSO models, which exhibit highly similar coefficients, alleviating concerns about overfitting due to the inclusion of many predictors.

Overall, telecommunications infrastructures (broadband connections density, availability of broadband providers, and a lesser extent, distance to nearest fiber nodes) remain the strongest predictors of entry. This finding aligns with operator statements that fiber connectivity, access to low-latency internet, are essential for service delivery. Similarly, counties with higher transmission latency (in quartiles 3 to 4 of transmission latency distribution) are significantly less likely to host data centers.

I next present evidence that proximity to data center clients affects their location choices. Counties with larger establishment shares in ICT and finance sectors show significantly higher entry probabilities. The magnitudes are economically meaningful: a one standard deviation increase in ICT establishment share is associated with a 0.08 percentage points increase in entry probability, relative to the baseline entry rate of 0.04%. This suggests that data centers do respond to the geographic distribution of potential customers, consistent with a model where firms value proximity to major service sector employers.

Data centers are also significantly less likely to locate in counties with higher commercial and industrial electricity tariffs, interestingly, transmission latency and telecommunication infrastructures appear more quantitatively important than electricity costs in predicting location choice. Finally, counties in states offering data center tax incentives are significantly more likely to receive facilities.

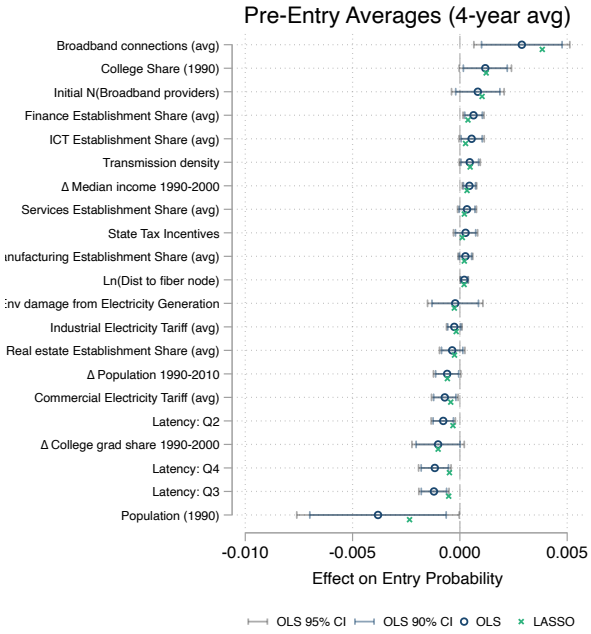
Finally, increased availability of skilled workforce, proxied by lagged county-level share of college graduates is significantly and positively associated with data center entry, while data centers are significantly less likely to locate in areas that had higher past levels of population (negative coefficient on 1990 population), and less likely to locate in counties which have seen faster population growth.

Figure 4(b) presents coefficients from the linear probability models and LASSO models, including pre-entry changes in characteristics (measured 4 to 2 years before entry). I continue to find that the level of initial telecommunications infrastructures, low internet latency, the level of skilled workforce, the presence of state-level tax incentives, and availability of electricity transmission infrastructures, as the most important predictors of data center location.

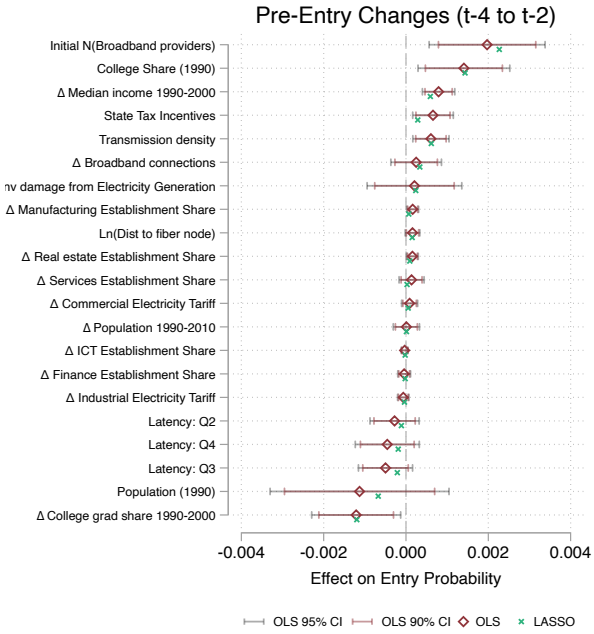
Lagged changes in broadband connections, proxying recent changes in telecommunications infrastructure availability, changes in electricity tariffs, and changes in share of establishments in potential client industries (ICT, finance, real estate) exhibit small and statistically insignificant associations with the probability of data center entry.

Figure (4) Correlates of County-level Data Center Entry

(a) Pre-Entry Levels (2 years before)



(b) Pre-Entry Changes (4 to 2 years before)



Notes: This figure presents linear probability model estimates from Equation 1. Panel (a) shows coefficients on pre-entry averages of variables measured 2 years before data center entry. Panel (b) shows coefficients on changes in variables measured 4 to 2 years before entry. The sample is a stacked cross-section of counties, with each stack corresponding to a cohort of counties receiving their first data center in a given year. Standard errors are clustered at the county level. The joint F-tests indicate whether all coefficients in each specification differ significantly from zero.

4.2 Local Impacts of Data Center Openings

I outline two complementary empirical strategies used to identify the effects of data center entry on local labour markets and establishment outcomes.

4.2.1 Callaway-Sant'Anna Difference-in-Differences (CSDID)

I use the Callaway and Sant'Anna (2021) difference-in-differences estimator (CSDID) to estimate dynamic treatment effects of data center entry on local outcomes. The estimator defines group-time average treatment effects $\text{ATT}(g,t)$ for each cohort g (defined by the year of first data center entry) at each calendar period t. The aggregate ATT and the dynamic event-study path are then formed as weighted averages across cohorts, using the not-yet-treated comparison group: counties that have not experienced a first entry by period t. This avoids the negative-weighting problems that arise under two-way fixed effects when treatment effects are heterogeneous across cohorts.

The estimation sample covers 592 county-level first-entry events between 2001 and 2023. I examine employment levels (workers per 10,000 population) and establishment counts by sector: total, manufacturing, all services, professional services, finance and real estate, information, administrative services, and construction. My baseline specifications control for county, year fixed effects, the interaction of baseline county-characteristics (1990 median income and college education shares) interacted with year fixed effects.

4.2.2 Synthetic Difference-in-Difference Event Studies

As a second strategy, I exploit spatial and temporal variation induced by data center openings in a Difference-in-Differences Design. Specifically, I use counties that had planned or not-yet-constructed data centers (not-yet-treated) and those that never had a data center opening over the sample period (never-treated) as the control group, comparing the evolution of their outcomes before and after a given data center opening. To identify the causal effect of data center entry on local economic outcomes, I employ a Synthetic Difference-in-Differences (SDID) approach, developed by Arkhangelsky et al. (2021). This approach combines elements of the synthetic control method with difference-in-differences estimation to address several concerns that arise in staggered adoption settings with heterogeneous treatment effects.

I exploit location-time variation in data center establishment (W_{it}) to estimate dynamic treatment effects. The SDID estimator is formally defined as:

$$\hat{\tau}_e^{\text{sdid}} = \arg\min_{\tau,\mu,\alpha,\beta} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} \left(Y_{it} - \mu - \alpha_i - \beta_t - \sum_{e=-K}^{L} \tau_e \cdot W_{it}^e \right)^2 \omega_i^{\text{sdid}} \lambda_t^{\text{sdid}} \right\} \tag{2}$$

where W_{it}^e is an indicator equal to 1 if location i is e periods away from data center entry, τ_e represents the treatment effect e periods from entry (our parameters of interest), ω_i are unit weights chosen to balance pre-treatment outcomes, and λ_t are time weights that prioritize periods predictive of post-treatment outcomes.

The implementation follows a two-step process. First, I construct a synthetic control group by creating a weighted combination of never-treated and not-yet-treated locations that approximates the counterfactual trajectory of treated locations. The weights ω_i are specifically selected to balance pre-treatment trends between locations with and without data centers. Second, I estimate dynamic treatment effects using the specification:

$$Y_{it} = \alpha_i + \lambda_t + \sum_{e=-K}^{L} \tau_e \cdot D_{it}^e + \varepsilon_{it} \tag{3}$$

where D_{it}^e is a dummy equal to 1 if location i is e periods from data center entry, K represents the maximum lead periods (pre-treatment), and L denotes the maximum lag periods (post-treatment).

This SDID approach offers several advantages over traditional two-way fixed effects (TWFE) models. As demonstrated by Arkhangelsky et al. (2021), SDID relaxes the parallel trends assumption through unit weighting, remains robust to heterogeneous treatment effects across cohorts and time, and avoids the negative weighting problems that can arise in staggered adoption settings (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020).

To account for potential differential outcome trends associated with local determinants of data center entry, and local policies or other contemporaneous shocks that coincide with data center entry, I control for a vector of state-by-year fixed effects. These account for state-specific policies that may have influenced data center locations or macroeconomic shocks that differentially affect locations.[13]

[13] My estimates are unchanged by excluding state-year fixed effects, suggesting that the synthetic DD estimator selects reasonable control counties.

5 Results

This subsection presents event-study estimates of the dynamic effects of first data center entry across sectors, focusing on employment and establishment counts.

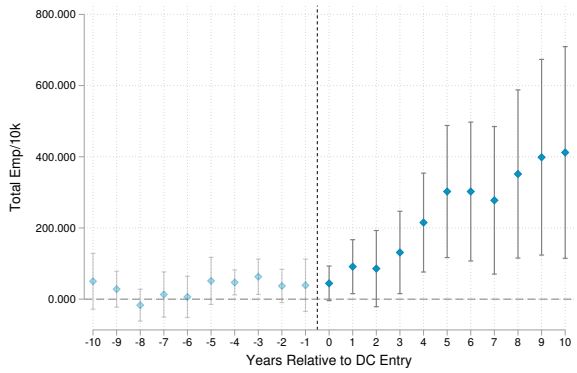
Employment (Levels, per 10,000 Population)

I first present event-study estimates of the effects and 95% Confidence Intervals of data center entry on local employment in Figure 5. We observe small and statistically insignificant pre-entry differences between treated counties and control counties, suggesting that, conditional on control variables, the not-yet-treated and never-treated counties are good comparison groups for the eventually treated counties.

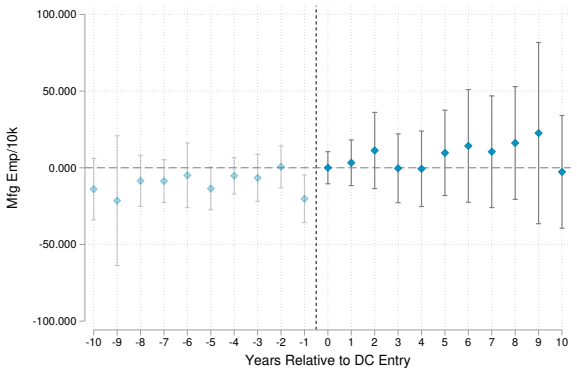
From Panel (a), we observe a significant increase in total employment from the first year post entry. The effect builds gradually rather than arriving at once: total employment rises by roughly 90 workers per 10,000 people within two years of entry, reaches about 300 per 10,000 five years out, and continues to grow to roughly 400 per 10,000 by year ten. This increase is not driven by the manufacturing sector (Panel b), which shows small and insignificant changes throughout the post period. It is mostly driven by the progressive increase in total service sector employment (Panel c), whose path is nearly identical to the total. Panels (d) to (h) disaggregate the increase in service employment into different sectors. The biggest contributors to positive employment effects are the professional service sector, administrative service, and real-estate service sector. Information-sector employment shows a positive but delayed response: coefficients are close to zero for the first five years and only become large and statistically significant effects from 7 years after the first entry event.

Figure (5) Dynamic Effects of First Data Center Entry on Employment (Levels, per 10,000 Population), Baseline

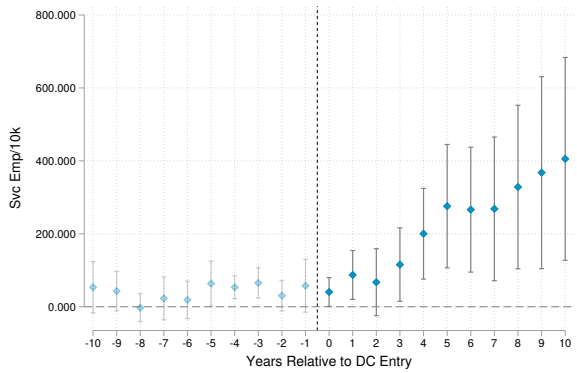
(a) Total



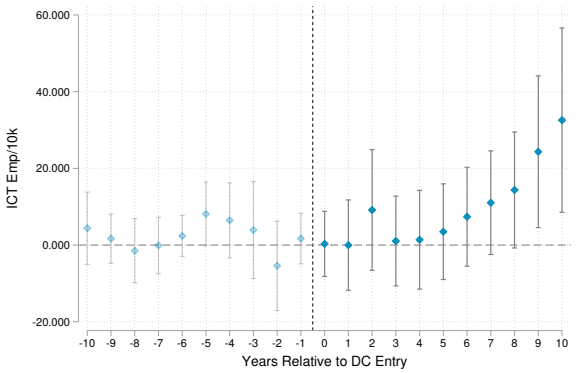
(b) Manufacturing



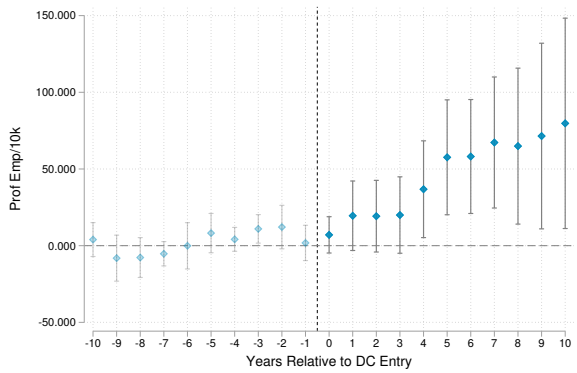
(c) Services



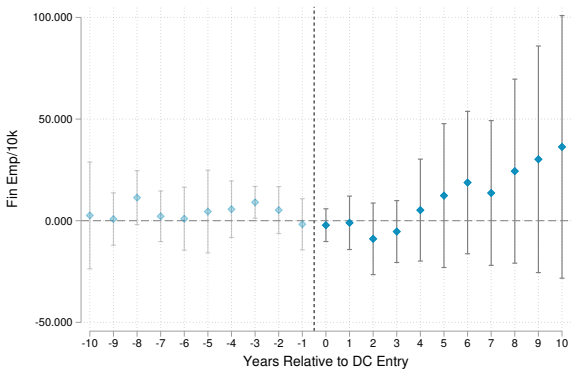
(d) Information



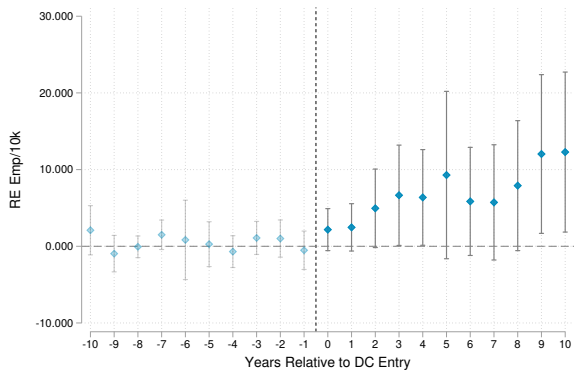
(e) Professional Services



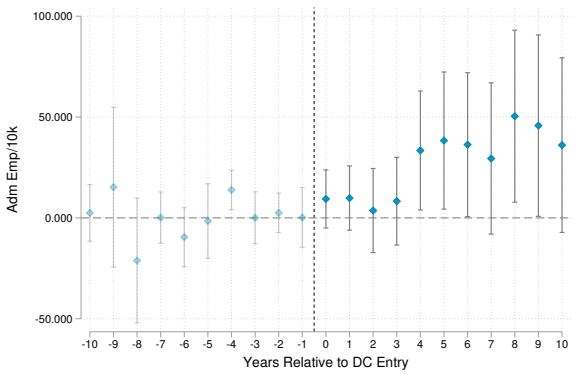
(f) Finance



(g) Real Estate



(h) Administrative Services



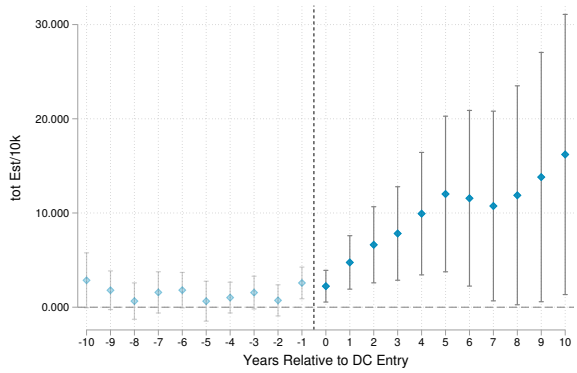
Notes: This figure presents Callaway–Sant'Anna (CSDID) event-study estimates of the effect of first data center entry on employment / 10k, across sectors, under the *Baseline (Min. 3 Pre-Periods)* sample restriction. Each panel plots dynamic coefficients relative to first entry (year −1 omitted). Shaded/whisker regions show 95% confidence intervals.

Effects of Data Center Entry on Establishments

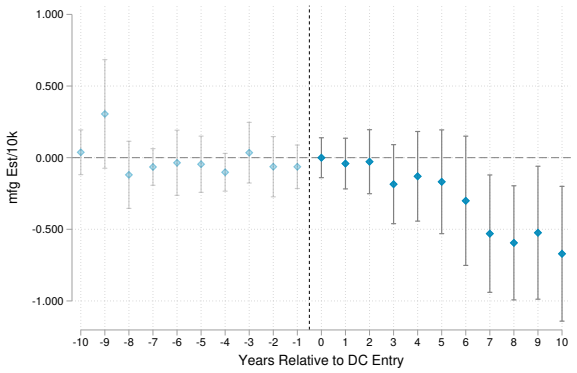
Next, I examine the efects of data center entry on establishment counts in 6. The absence of statistically or economically significant pre-entry differences in outcome trends supports the use of not-yet-treated and never-treated counties as a comparison group. We observe a significant increase in total establishment counts in Panel a, an increase of 16 establishments per 10,000 people by year 10, with the increase being driven by sector-sector firm entry. The timing differs across service sectors in an intuitive way. Professional service (Panel e) and real estate (Panel g) establishments grow steadily over the entire post period, consistent with data centers affecting the location choice of high-skilled service establishments. Finance establishments (Panel f) show positive but statistically insignificant changes. Manufacturing establishments (Panel b) are the one margin that moves in the opposite direction: three years after the first data center entry, we observe a reduction of -0.5 to -0.7 establishments per 10,000. This suggests that data center entry crowded out the operations of some manufacturing firms, although the implied magnitude is small relative to the service-sector gains.

Figure (6) Dynamic Effects of First Data Center Entry on Establishments (Levels, per 10,000 Population), Baseline (Min. 3 Pre-Periods)

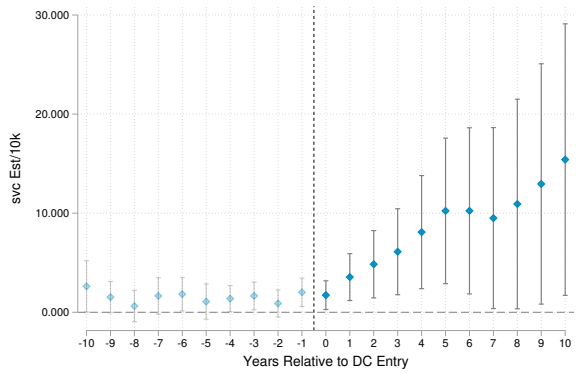
(a) Total



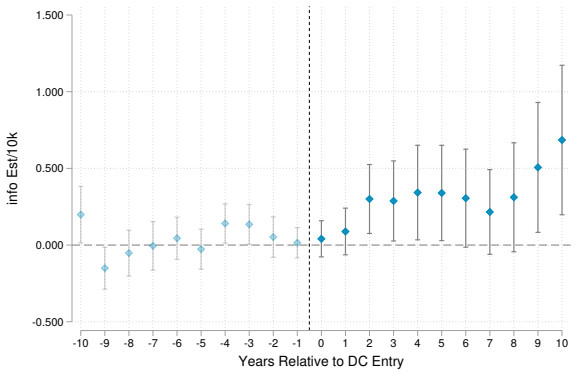
(b) Manufacturing



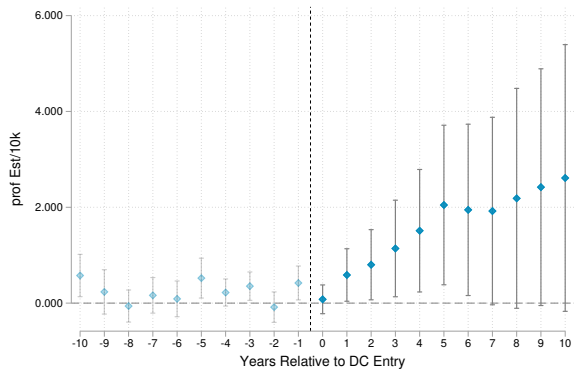
(c) Services



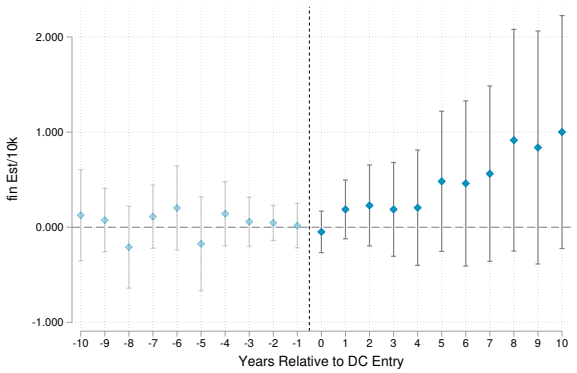
(d) Information



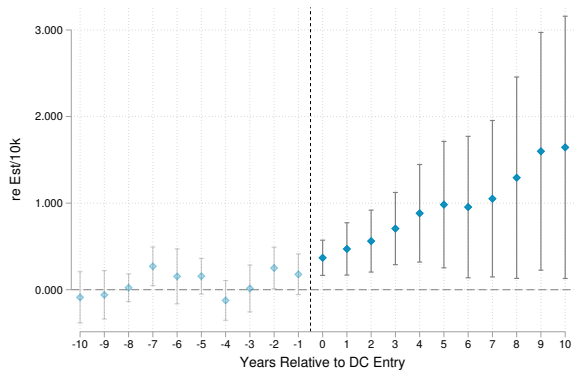
(e) Professional Services



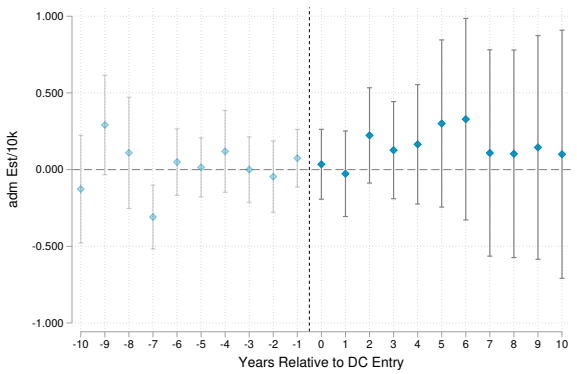
(f) Finance



(g) Real Estate



(h) Administrative Services



Notes: This figure presents Callaway–Sant'Anna (CSDID) event-study estimates of the effect of first data center entry on establishments / 10k, across sectors, under the *Baseline (Min. 3 Pre-Periods)* sample restriction. Each panel plots dynamic coefficients relative to first entry (year −1 omitted). Shaded/whisker regions show 95% confidence intervals.

Tables 1 and 2 report the treatment-cohort-weighted ATTs, constructed as the average of the event-study estimates over horizons 0 to 10. First data center entry raises total employment by 238 workers per 10,000 population, 7.8 percent of the pre-entry control mean. Almost all of this comes from services (220 per 10,000, or 9.7 percent). The largest proportional responses are in professional services, where the ATT of 46 workers per 10,000 amounts to 40.6 percent of the pre-entry mean, followed by information (17.5 percent, though imprecisely estimated) and real estate (16.2 percent). The manufacturing employment ATT is 7.7 per 10,000, 1.4 percent of the mean, and statistically indistinguishable from zero. On the establishment margin, first entry raises total establishment counts by 10.6 per 10,000, or 4.4 percent of the pre-entry mean, again concentrated in services (8.8 per 10,000, or 4.5 percent), with precisely estimated increases in construction (8.5 percent), professional services (9.2 percent), and real estate (10.2 percent). The manufacturing establishment ATT is -0.32 per 10,000 (-2.7 percent), significant at the ten percent level.

Comparing the two margins, the employment response is proportionally larger than the establishment response: total employment rises by 7.8 percent of the pre-entry mean against 4.4 percent for establishment counts. Data center entry therefore raises local employment both by attracting new establishments and by expanding employment within existing ones, with the within-establishment margin accounting for the larger share of the total effect.

Table (1) ATT of First Data Center Entry on Sector Employment per 10,000 Population

	Total	Mfg	Svc	ICT	Prof	Finance	RE	Constr
ATT	237.515***	7.655	220.285***	9.553	45.611**	11.218	6.882*	17.516*
	(87.306)	(15.344)	(80.074)	(7.348)	(19.089)	(16.264)	(3.639)	(9.622)
Effect (% of pre-entry mean)	7.8	1.4	9.7	17.5	40.6	9.6	16.2	9.2
Control mean	3044.127	539.359	2276.691	54.556	112.214	116.823	42.542	190.223
N	59,041	55,306	59,041	49,124	56,310	56,866	51,957	57,674
Coll. quartile×Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inc. quartile×Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports Callaway–Sant'Anna (CSDID) ATT estimates of the effect of first data center entry on sector employment per 10,000 population. The treatment cohort is defined by the first entry of any facility type, taking the earliest of the colocation, hyperscale, and enterprise entry dates, and the sample requires at least three pre-treatment periods (*min3*). Outcomes are residualized on quartile(baseline college share)×year and quartile(baseline income)×year fixed effects before estimation, and not-yet-treated counties serve as the control group. The ATT is the average of the event-study estimates over horizons 0 to 10. "Effect (% of pre-entry mean)" expresses the ATT relative to the control-group mean reported in the row below it. Standard errors are in parentheses. Sector abbreviations: Mfg = manufacturing; Svc = services; ICT = information; Prof = professional services; RE = real estate; Constr = construction. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Table (2) ATT of First Data Center Entry on Sector Establishments per 10,000 Population

	Total	Mfg	Svc	Prof	Finance	RE	Constr
ATT	10.558***	-0.322*	8.827**	1.660**	0.359	0.972***	2.352***
	(4.003)	(0.187)	(3.581)	(0.776)	(0.376)	(0.376)	(0.657)
Effect (% of pre-entry mean)	4.4	-2.7	4.5	9.2	2.3	10.2	8.5
Control mean	239.178	11.839	195.945	18.092	15.581	9.547	27.713
N	59,041	55,335	59,041	56,335	56,876	52,016	57,690
Coll. quartile×Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inc. quartile×Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports Callaway–Sant'Anna (CSDID) ATT estimates of the effect of first data center entry on sector establishment counts per 10,000 population. The treatment cohort is defined by the first entry of any facility type, taking the earliest of the colocation, hyperscale, and enterprise entry dates, and the sample requires at least three pre-treatment periods (*min3*). Outcomes are residualized on quartile(baseline college share)×year and quartile(baseline income)×year fixed effects before estimation, and not-yet-treated counties serve as the control group. The ATT is the average of the event-study estimates over horizons 0 to 10. "Effect (% of pre-entry mean)" expresses the ATT relative to the control-group mean reported in the row below it. Standard errors are in parentheses. Sector abbreviations: Mfg = manufacturing; Svc = services; Prof = professional services; RE = real estate; Constr = construction. Establishment counts for the information and management sectors are not available in the underlying data. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Taking stock, I find that data center entry generates a sizable and persistent expansion of local service activity: total employment rises by about 8 percent and establishment counts by about 4 percent of their pre-entry means, and both effects accumulate over the ten years post entry. The gains are concentrated in high-skilled service sectors, professional services and real estate in particular, alongside a fast construction response that levels off once facilities are built. Manufacturing employment is unaffected, but the number of manufacturing establishments declines at long horizons, suggesting the crowding out of electricity-intensive operations.

One concern with these estimates is anticipation: counties may respond before the recorded entry year, through construction activity, land acquisition, or hiring that begins once a project is announced. As a robustness check, I shift the treatment year forward by one year and re-estimate the event studies. The results, reported in Appendix Figures A.1 and A.2, are close to the baseline: pre-entry coefficients remain small and insignificant, and the post-entry paths of total employment and establishments build at the same pace and to the same levels, roughly 420 workers and 16 establishments per 10,000 population ten years out. The baseline estimates are therefore not driven by responses that predate the recorded entry year.

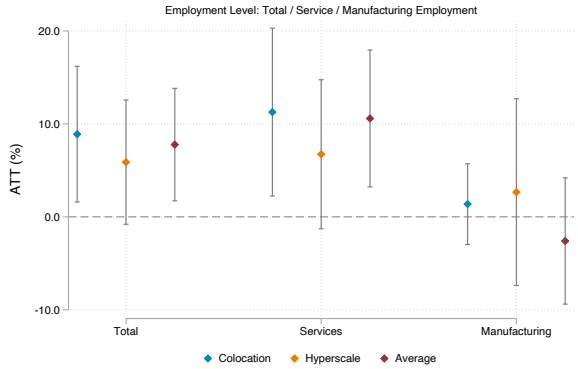
Facility-Type Heterogeneity. A policy-relevant dimension of treatment effect heterogeneity relates to heterogeneous effects by facility type. This could be driven by factors such as the size of the facility, the type of clients served by the facility.

I test for differential local employment impacts of data center entry by estimating event studies and ATTs for the first entry event of a given type of facility.

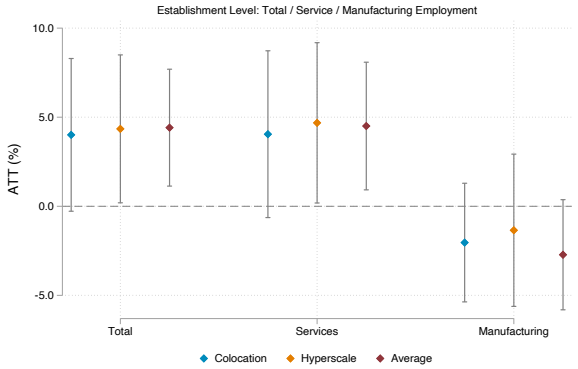
Figure 7 shows the results, comparing static ATT estimates (expressed as a percent of the control-group mean) across three facility-type cohorts: colocation, hyperscale, and an average cohort pooling the earliest first-entry event across colocation, enterprise, and hyperscale facilities (*min3*). The top row reports effects on total, service, and manufacturing employment (left) and establishments (right); the bottom row reports effects on professional, finance, real estate, and management employment (left) and establishments (right).

Figure (7) ATT Estimates by Facility Type: Employment and Establishment Effects

(a) Employment: Total / Service / Manufacturing



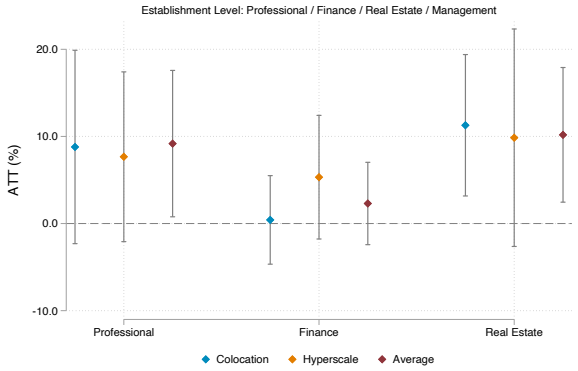
(b) Establishments: Total / Service / Manufacturing



(c) Employment: Professional / Finance / Real Estate / Management



(d) Establishments: Professional / Finance / Real Estate / Management



Notes: Each panel plots Callaway–Sant'Anna ATT estimates (as a % of the control-group mean) for the first entry of a given facility type, by sector. Colocation and Hyperscale are facility-type-specific cohorts; Average pools the earliest first-entry date across colocation, enterprise, and hyperscale facilities (*min3*). Whiskers show 95% confidence intervals. No enterprise-only cohort estimates were available; the Average series is used as the pooled reference point instead.

Panel (a) shows that total and service employment respond positively to first data center entry across all three facility-type cohorts, with colocation entries associated with the largest total-employment gain (approximately 9% of the control mean), hyperscale the smallest (around 6%), and the average cohort in between (about 8%). The service-sector estimates are of similar magnitude and precision to the total-employment estimates, confirming that the aggregate employment response is driven almost entirely by services rather than manufacturing. Manufacturing employment effects are close to zero and imprecisely estimated for all three cohorts, with the average cohort's point estimate turning slightly negative (around -3%).

Panel (b) shows a qualitatively similar pattern for establishment counts: total and service establishments increase by 4–4.5% of the control mean under all three cohorts, with colocation, hyperscale, and average estimates at the same magnitude. Manufacturing establishment counts decline modestly (by 1.5–3%) across cohorts, though again these effects are not statistically significant. The similarity of the colocation and hyperscale point estimates in Panel (b), relative to the more pronounced colocation–hyperscale gap in Panel (a), suggests that facility-type heterogeneity is more pronounced on the employment margin (intensity of hiring per establishment) than on the extensive margin of establishment entry.

Panels (c) and (d) turn to professional, finance, real estate, and management sectors, where the facility-

type comparison becomes noisier and more heterogeneous. Professional-service employment (Panel c) rises substantially under colocation entry (point estimate near 40% of the control mean) versus a smaller and less precise hyperscale response (around 20%), with wide confidence intervals for both; the corresponding establishment-count effects (Panel d) are more tightly estimated and of comparable magnitude across cohorts (8–9%), again indicating that the extensive margin is more stable across facility types than the employment-intensity margin. Finance employment effects are close to zero and imprecise for all three cohorts, while finance establishment counts show a positive but statistically insignificant hyperscale effect (around 5%) against a near-zero colocation effect. Real estate employment and establishment effects are positive and of similar magnitude (roughly 10–15%) across colocation, hyperscale, and average cohorts, the most internally consistent sector in this group. Management employment is the most volatile cell in the table: the hyperscale point estimate is large and negative (around -170%) with an enormous confidence interval, while colocation is modestly negative and the average cohort's estimate is positive, reflecting the small underlying cell sizes for management-sector employment around data center entry events and warranting caution in interpreting this coefficient.

Taken together, these results point to two distinct forms of facility-type heterogeneity. First, colocation-facility entry is associated with systematically larger and more precisely estimated positive employment effects than hyperscale entry in total, service, and professional-service employment, while the two cohorts converge on establishment-count effects; this is consistent with colocation facilities drawing a denser mix of on-site service and professional activity relative to the more capital-intensive, land-extensive hyperscale campuses. Second, treatment effect heterogeneity is sector-dependent rather than uniform: real estate and core services effects are stable and directionally consistent across facility types, whereas niche sectors with small underlying cell counts (management, finance) generate noisy, sign-inconsistent estimates that should be interpreted cautiously. The average (*min3*) cohort generally tracks a weighted blend of the colocation and hyperscale estimates, as expected given it pools the earliest entry event across facility types, and provides the most stable summary series where the two facility-type-specific estimates diverge.

5.1 Local Migration Responses to Data Center Entry

The results so far suggest that data centers create local "structural transformation" by shifting the structure of employment away from manufacturing sectors towards service sectors. I next ask whether the local employment gains from data center entry primarily affect existing residents in the counties, or whether it could be partly driven by increase in-migration. Data center entry could also drive increased out-migration if it produces visual disamenities, noise and air pollution, that are quantitatively meaningful. Finally, estimating the impacts of local data center entry on net migration, controlling for their impacts on local real wages, recovers key parameters on the amenity spillovers of data center entry, which are central to evaluating place-based policies that influence their location choice.

Using county-to-county annual migration flows from the IRS Statistics of Income for 1999–2021, I aggregate to the destination-county level and estimate the same Callaway–Sant'Anna event studies on three outcomes: the *in-migration rate* (annual in-migrants divided by base-year county population), the *out-migration rate* (annual out-migrants over base-year population), and the *net-migration rate* (in-migrants minus out-migrants, over base-year population).

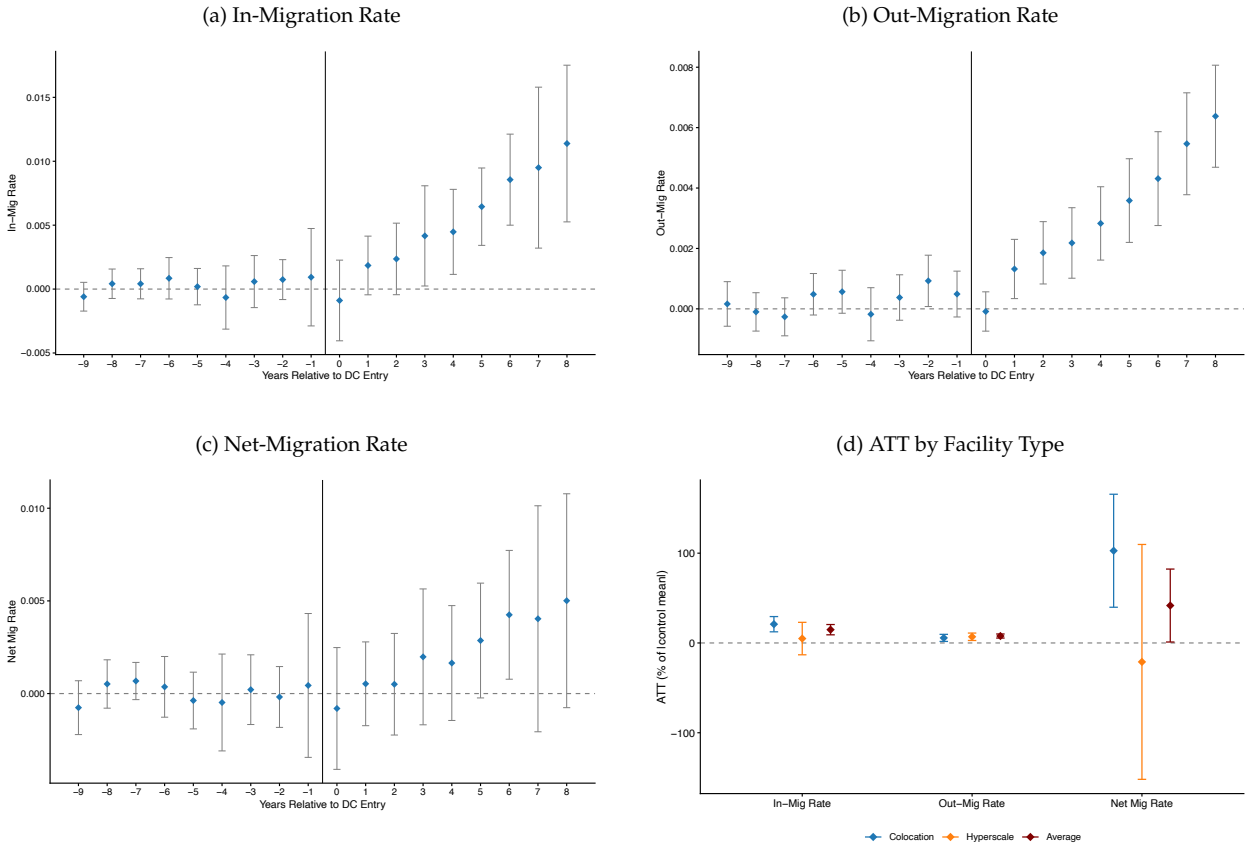
Figure 8 and Table 3 report the event-study and ATT estimates. Panels (a)–(c) of Figure 8 plot the dynamic effects on the in-, out-, and net-migration rates, pooling different types of data center facilities; panel (d) collapses the estimates to static ATTs and compares them across facility types (colocation, hyperscale, and the pooled average).

Pre-entry coefficients in panels (a)–(c) are small and statistically insignificant, so treated and comparison

counties are on parallel migration trends before entry. After entry, *both* gross flows rise: the in-migration rate increases by 0.0073 (a 14.8% increase relative to the control-group mean of 0.049), and the out-migration rate increases by 0.0044 (7.7% of its mean of 0.056); both are precisely estimated (Table 3). Data center entry therefore raises local population churn in both directions. As the in-migration response is nearly twice the out-migration response, the net effect is positive: the net-migration rate rises by 0.0029. The control-group net-migration rate is slightly negative (-0.007), so this ATT represents a shift of about 42% of the absolute baseline toward net in-migration. The extensive-margin indicators point the same way: the probability of any recorded in-migration rises by 2.1 percentage points and of any out-migration by 1.4 points. On net, the labor-demand pull created by the local service-sector expansion documented above dominates any disamenity-driven push.

Panel (d) shows that this pattern is mainly driven by the entry of colocation data centers, with limited evidence of significant net-migration responses from hyperscale data center entry. Colocation facilities raise the in-migration rate by 20.9% of the control mean and the net-migration rate by almost 100% of control counties' mean, while hyperscale entry has a small and imprecise effect on in-migration rate (4.9%), and a negative but imprecisely estimated effect on net-migration rate. This colocation–hyperscale gap mirrors the employment results: colocation campuses draw a denser mix of on-site service activity, and hence more in-migration, than the more capital-intensive, land-extensive hyperscale facilities.

Figure (8) Dynamic Effects of First Data Center Entry on Migration Rates



Notes: Callaway–Sant'Anna (CSDID) event-study estimates of the effect of first data center entry on county migration rates, using IRS county-to-county migration flows, 1999–2021. Migration rates are defined relative to base-year (1990) county population: the *in-migration rate* is annual in-migrants over base-year population; the *out-migration rate* is annual out-migrants over base-year population; the *net-migration rate* is (in-migrants − out-migrants) over base-year population. Panels (a)–(c) plot dynamic coefficients relative to first entry (year -1 omitted) for the pooled first-entry (*min3*) cohort, with 95% confidence intervals. Panel (d) plots static ATTs by facility type (Colocation, Hyperscale, and the pooled Average), each expressed as a percentage of the *absolute value* of the control-group mean of the outcome; whiskers are 95% confidence intervals. The absolute mean is used because the control-group mean of the net-migration rate is negative (≈ -0.007), so dividing by the signed mean would flip the sign of a positive ATT; the control-group means are ≈ 0.049 (in), ≈ 0.056 (out), and ≈ -0.007 (net). Not-yet-treated counties serve as the comparison group.

Table (3) CSDID ATT: DC Entry on Migration — Min(Colo, Ent, Hyper) Entry

	In Rate	Out Rate	I(In >0)	I(Out >0)	Net Rate
ATT	0.007***	0.004***	0.021***	0.014***	0.003**
	(0.001)	(0.001)	(0.006)	(0.003)	(0.001)
Mean (control group)	0.049	0.056	0.949	0.970	-0.007
N (obs)	70,248	70,248	70,248	70,248	70,248
N (counties)	3,061	3,061	3,061	3,061	3,061
Years	1999–2021	1999–2021	1999–2021	1999–2021	1999–2021

Notes: CSDID (Callaway & Sant'Anna 2021), not-yet-treated comparison group. Controls: county FE, quartile(1990 college share)×year FE, quartile(1990 median HH income)×year FE. County-clustered SEs. In (Out) Rate is annual in- (out-) migrants divided by base-year county population; Net Rate is (in − out) migrants over base-year population. I(In >0) and I(Out >0) are indicators for any recorded in- and out-migration. $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Finally, in Section 6.9, I use the same empirical research design to estimate how county-level data center entry affects net migration, controlling for their impacts on real wages, in order to estimate the amenity spillovers of data center entry.

5.2 Robustness Checks

I conduct some robustness checks to examine the sensitivity of the event-study results and consider alternative explanations driving the result.

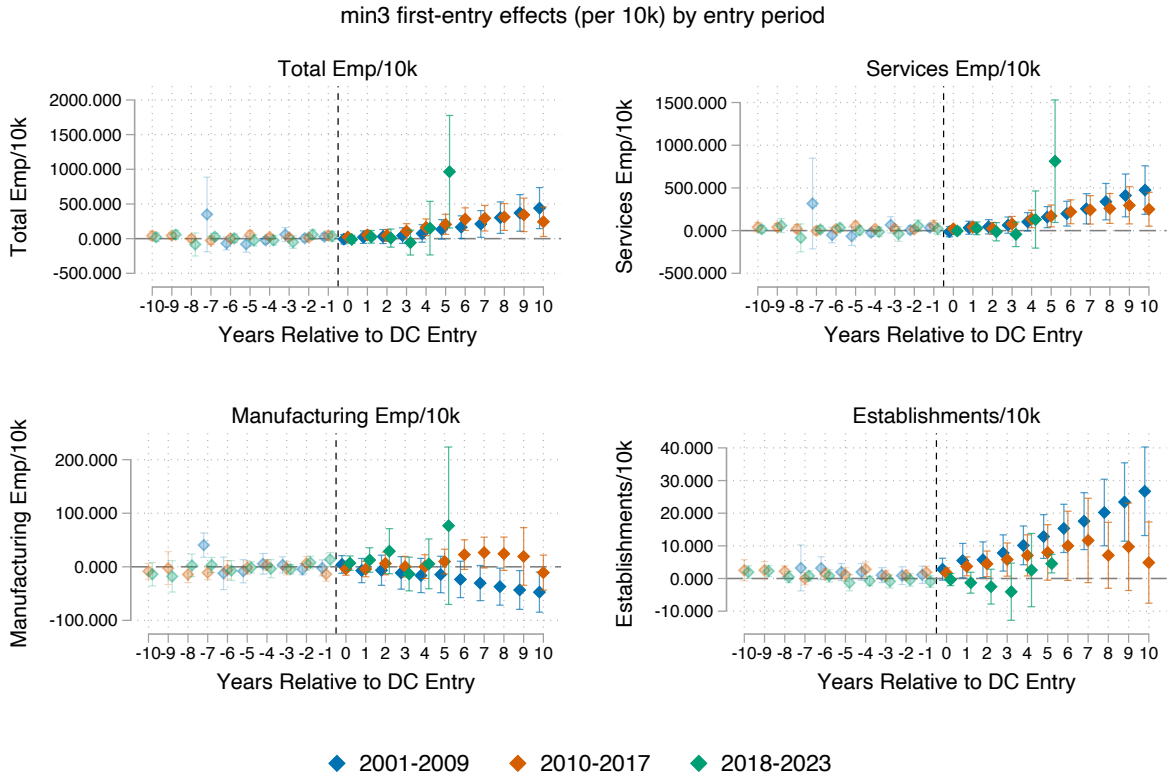
Synthetic Difference-in-Differences. To address potential concerns about the endogeneity in the placement of data centers, I estimate Synthetic Difference-in-Differences event studies following Arkhangelsky et al. (2021), which optimally chooses weights on untreated counties to minimise pre-entry differences in outcome trends between treated and control units. Appendix section B.1 presents the results, which remain qualitatively and quantitatively similar to the main event study results. This suggests that the event study estimates are not sensitive to the alternative choice of control counties.

Controlling for Contemporaneous Shocks. One might be concerned that data center entry coincides with contemporaneous shocks operating at a broader level, e.g., a state-wide policy change or industry boom common to every county in a state, or divergent trajectories among places with different *ex ante* digital-infrastructure endowments, since data centers locate close to fiber and connectivity may itself shape local employment. To address these concerns, I re-estimate the pooled first-entry event studies, adding, in turn, (i) state×year fixed effects, which absorb any shock common to a state in a given year, and (ii) fixed effects for the quartile of a county's initial distance to the nearest long-haul fiber node interacted with year, which allow counties with different pre-existing connectivity to follow separate time paths. Appendix section B.2 presents the results. The dynamics and magnitudes are essentially unchanged: for total employment the average post-entry effect moves only from 168.5 additional jobs per 10,000 residents in the baseline to 153.9 and 177.0 respectively, service employment behaves almost identically, and pre-entry coefficients stay flat under both sets of controls. The estimated effects are thus not an artefact of coincident state-level shocks, nor of differential trends among counties that were better connected to begin with.

Heterogeneity by Entry Cohort. The data-center build-out spanned two decades, evolving from an early enterprise- and colocation-driven phase to the recent hyperscale boom. To verify that the pooled estimate is not driven by any single era, and to ask whether the effect has shifted over time, I split treated counties by the period of first entry into three cohorts (2001–2009, 2010–2017, and 2018–2023) and re-estimate the event study for each cohort against never-treated counties, dropping counties whose first entry falls outside the window in question.

Figure 9 reports the results. The employment response is a consistent feature across all three cohorts: the average post-entry effect on total employment is 162, 189, and 180 jobs per 10,000 residents for the 2001–2009, 2010–2017, and 2018–2023 cohorts respectively, and service employment tracks closely (188, 159, and 151). The most recent cohort is less precisely estimated: counties entering after 2018 are observed for only a handful of post-entry years within the 2001–2023 sample. However, its point estimates are comparable in magnitude to the earlier cohorts. Manufacturing employment is null in every cohort. Overall, the positive employment effect of data center entry is a stable feature across the entire two decades of entry, not the product of any one period.

Figure (9) Heterogeneity by Entry Cohort: Effect of First Data Center Entry by Period of First Entry



Notes: Each panel plots Callaway–Sant'Anna event studies of first data center entry on a per-10,000-resident outcome, estimated separately for treated counties whose first entry falls in 2001–2009 (blue), 2010–2017 (vermillion), or 2018–2023 (green), each compared with never-treated counties; counties entering outside a given window are excluded from that cohort's estimation. Outcomes are residualised on 1990 college-share quartile×year and 1990 median household income quartile×year fixed effects. Markers are dynamic coefficients relative to first entry (year −1 omitted) with 95% confidence intervals; pre-entry markers are faded. The vertical scale differs across panels.

Leave-One-State-Out. A remaining concern is that the estimated effects are driven by a single state whose data center build-out coincided with other local shocks. Virginia is the natural candidate: Northern Virginia hosts a large share of US data center capacity, and its concurrent expansion of professional and information services could in principle generate the observed employment path without any causal role for data center entry. To assess this, I re-estimate the pooled first-entry (*min3*) specification 41 times for each outcome, each time dropping all counties in one of the 41 states that contain at least one treated county. Appendix section B.3 reports the results. No single exclusion overturns the findings: the full-sample effect on total employment of 168.5 jobs per 10,000 residents ranges from 156.7 to 191.0 across the 41 re-estimations, service employment behaves almost identically, and every leave-one-out estimate remains inside the full-sample 95% confidence interval and statistically significant at the 5% level. Establishment counts are tighter still. Dropping Virginia leaves both sets of estimates essentially unchanged, so the results do not rest on any one state's experience.

6 Quantitative Model of Data Centers and the Economy

To interpret the mechanisms underlying the reduced-form estimates and evaluate the welfare consequences of alternative spatial policies, I develop a quantitative spatial equilibrium model. This section outlines the model, the identification and estimation of key parameters, and the counterfactual policy experiments.

The model incorporates five key features motivated by the empirical findings: (i) an electricity market with data-driven supply curves across Balancing Authorities; (ii) partially tradable computation services, subject to network quality degradation, generating input-cost spillovers for data-intensive firms; (iii) local pollution damages from electricity generation, computed via source-receptor matrices for fine particulate matter ($PM_{2.5}$), nitrogen oxides (NO_x), and sulfur dioxide (SO_2); (iv) data center location choice responding to local costs, policies, location fundamentals, and market access; and (v) endogenous migration across locations and sectors, governed by a Fréchet migration framework.

6.1 Set-Up

The economy consists of C locations (US counties) indexed by c. Each location has four production sectors: tradable goods (T), high-compute services ($S1$), low-compute services ($S2$), and electricity generation (E). Workers are heterogeneous across sectors and mobile across both locations and sectors. Total population is $\bar{L} = \sum_c L_c$.

By splitting services into a high-compute tier ($S1$: information, finance, professional services) and a low-compute tier ($S2$: retail, hospitality, other services, utilities), the model can capture the differential exposure of service sub-sectors to computation input costs, a key mechanism through which data centers can generate unequal welfare effects across workers and locations.

Second, by allowing each worker to choose unemployment as an outside option, the model implicitly allows data center entry to affect total stock of employed workers. Unemployed workers are assumed to receive a population-weighted average real wage as their outside option value.

Households. I assume that each representative household in location c has Cobb–Douglas preferences over tradable goods, services, housing, and electricity:

$$U_c = B_c \cdot C_{T,c}^{\mu} \cdot C_{S,c}^{\theta} \cdot H_c^{\nu} \cdot E_{H,c}^{\varepsilon_E} \cdot g(D_c), \qquad \mu + \theta + \nu + \varepsilon_E = 1 \tag{4}$$

where B_c denotes location amenities, $C_{T,c}$ and $C_{S,c}$ are CES consumption indices, H_c is housing, $E_{H,c}$ is residential electricity, and $g(D_c)$ captures local environmental damages, $g'(D_c) < 0$. Services. Services consumption $C_{S,c}$ is itself a Cobb-Douglas aggregate of $S1$ and $S2$ varieties.[14] Accounting for residential electricity consumption provides a direct link between pecuniary price externalities of data center location choice and households' cost of living and welfare.

Budget Constraint and Residential Electricity Demand. Households pay two ad valorem taxes (defined in Section 6.6): a general consumption (sales) tax τ_c on the tradable-and-service composite, from which housing is exempt; and a residential electricity tax $t_c^{E,H}$ that wedges the retail electricity price above the balancing-authority wholesale price $p_{E,b}$,

$$p_{E,c}^{H} = (1 + t_c^{E,H})\, p_{E,b}. \tag{5}$$

[14] I obtain all expenditure shares from the US National Income & Product Accounts (BEA NIPA Table 2.4.5U).

The budget constraint applies the consumption wedge to the taxed composite and the electricity wedge through $p^H_{E,c}$,

$$(1+\tau_c)\left(P_{T,c}C_{T,c}+P_{S,c}C_{S,c}\right)+P_{H,c}H_c+p^H_{E,c}\,E_{H,c}=w_c+T_c, \tag{6}$$

where $T_c = \text{Rev}_c / L_c$ is the per-capita rebate of local tax revenue. Cobb–Douglas preferences imply constant expenditure shares, so residential electricity demand is

$$E_{H,c}=\frac{\varepsilon_E\,(w_c+T_c)}{p^H_{E,c}}. \tag{7}$$

The electricity wedge $t^{E,H}_c$ rations household consumption one-for-one, so data center entry that raises $p_{E,b}$ crowds out residential consumption through the same congestion channel that operates on firms.

Indirect Utility. Substituting optimal demands, the indirect utility of household in location c is:

$$V_c=\frac{B_c\,(w_c+T_c)}{\left((1+\tau_c)P_{T,c}\right)^{\mu}\left((1+\tau_c)P_{S,c}\right)^{\theta}P^{\nu}_{H,c}\left(p^H_{E,c}\right)^{\varepsilon_E}}\cdot g(D_c), \tag{8}$$

with consumption price index $P_c=\left((1+\tau_c)P_{T,c}\right)^{\mu}\left((1+\tau_c)P_{S,c}\right)^{\theta}P^{\nu}_{H,c}\left(p^H_{E,c}\right)^{\varepsilon_E}$. Both household taxes enter directly as multiplicative wedges in the cost of living; the sales tax scales P_c by $(1+\tau_c)^{\mu+\theta}$.

Local Environmental Damage Function. I assume that local pollution concentrations cause a proportional reduction to worker utility, similar to Hollingsworth et al. (2026). $g(D_c)$ is the fraction of reduction in worker utility due to total monetised local pollution damage in c.

$$g(D_c)=1-\frac{D_c}{V^0_c} \tag{9}$$

where D_c is monetised local pollution damage per resident in location c (\$ per capita per year) and V^0_c is baseline indirect utility in dollar terms. This functional form ensures damages are bounded $g(D_c)\in(0,1]$, linearises the relationship between dollar damages and welfare loss, and allows for the dose-response heterogeneity emphasised in Holland et al. (2016).

Current Treatment of Global Carbon Damage The model is able to predict the aggregate value and set of power plants responsible for carbon emissions from electricity generation in the US. For now, I do not include welfare costs from CO2 emissions into household utility. In all counterfactual results, I report changes in CO2 emissions and changes in monetised aggregate emissions damages.

Housing Equilibrium and Local Price Determination. Local housing supply is an increasing, concave function of the quantity of housing consumed in equilibrium, governed by the calibrated fundamentals $(\eta_c, \bar{H}_c)$, where η_c is the county-specific housing supply elasticity, imputed from Baum-Snow and Han (2024).[15]

$$HS_c=\bar{H}_c\cdot Q^{\eta_c}_{H,c}. \tag{10}$$

[15] I provide more details on the imputation procedure in the Data Appendix. Model results are unaffected by using a single national housing supply elasticity, e.g, from Saiz (2010).

On the demand side, I assume each worker spends a fixed Cobb-Douglas share of income on housing. Equating housing demand to supply gives the equilibrium price:

$$P_{H,c} = \alpha_H \left\{ \frac{w_T L_T + w_{S1} L_{S1} + w_{S2} L_{S2}}{H S_c} \right\}^{\frac{1}{1+\eta_c}}. \tag{11}$$

In any counterfactual equilibrium, $(\eta_c, \bar{H}_c)$ are held fixed and equation (A.4) is re-solved using counterfactual sector wage bills. County-specific η_c governs how much the potential wage shocks from data center entry pass through to local housing costs versus housing quantities.

Migration. I allow for workers to choose the location and sector of work, subject to switching costs. Workers draw location-sector-specific taste shocks from a Fréchet distribution with shape parameter ι (the inverse migration elasticity). The bilateral migration share from origin o, sector s to destination c, sector s' is:

$$\pi_{oc,ss'} = \frac{(B_c \cdot w_{s',c}/P_c)^{\iota} \cdot \kappa_{oc}}{\sum_{c'} \sum_{s''} (B_{c'} \cdot w_{s'',c'}/P_{c'})^{\iota} \cdot \kappa_{oc'}} \tag{12}$$

where κ_{oc} captures bilateral migration frictions (distance, state borders, historical ties).

6.2 Production Sectors

Tradable Sector. Tradable firms produce differentiated varieties under monopolistic competition using labour and electricity:

$$y_T = A_{T,c} \cdot \ell^{\alpha_L} \cdot e^{\alpha_E}, \quad \alpha_L + \alpha_E = 1 \tag{13}$$

Firms set prices as a constant markup over marginal cost. The only fiscal instrument they face is an electricity input tax t_c^E, levied on the delivered retail tariff, so that they buy power at $p_{E,c}^{\text{firm}} = (1 + t_c^E)(p_{E,c} + \varrho_c)$, where $p_{E,c}$ is the wholesale price and ϱ_c the retail delivery margin. All five tax instruments in the model are defined together in Section 6.6. Marginal cost is then

$$MC_{T,c} = \frac{w_{T,c}^{\alpha_L} \cdot \left(p_{E,c}^{\text{firm}}\right)^{\alpha_E}}{A_{T,c} \cdot \alpha_L^{\alpha_L} \cdot \alpha_E^{\alpha_E}} \tag{14}$$

Tradables are shipped across locations with iceberg trade costs $\tau_{cd} \geq 1$ calibrated from the 2017 Commodity Flow Survey, following Allen and Arkolakis (2014). Free entry with fixed cost F_T pins down the number of firms $n_{T,c}$. The tradable price index in c aggregates varieties from all origins:

$$P_{T,c} = \left[\sum_{d=1}^{C} n_{T,d} \cdot (p_{T,d} \cdot \tau_{dc})^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \tag{15}$$

Service Sectors ($S1$, $S2$). Service firms in tier $s \in \{S1, S2\}$ produce non-tradable varieties using labour, computation, and electricity:

$$y_s = A_{s,c} \cdot \ell^{\beta_{L,s}} \cdot m^{\beta_{M,s}} \cdot e^{\beta_{E,s}}, \quad \beta_{L,s} + \beta_{M,s} + \beta_{E,s} = 1 \tag{16}$$

The key parameter $\beta_{M,s}$ captures the computation intensity of services in tier s, the demand linkage between service firms and data centers. I allow $\beta_{M,S1} > \beta_{M,S2}$, reflecting the higher computation intensity of information, finance, and professional services relative to retail, hospitality, and other services. The

marginal cost of each tier depends on the local computation price index $P_{M,c}$:

$$MC_{s,c} = \frac{w_{s,c}^{\beta_{L,s}} \cdot P_{M,c}^{\beta_{M,s}} \cdot \left(p_{E,c}^{\text{firm}}\right)^{\beta_{E,s}}}{A_{s,c} \cdot \beta_{L,s}^{\beta_{L,s}} \cdot \beta_{M,s}^{\beta_{M,s}} \cdot \beta_{E,s}^{\beta_{E,s}}} \tag{17}$$

Services are non-tradable, so $P_{s,c}$ depends only on local supply. The overall services price index is $P_{S,c} = P_{S1,c}^{\omega_{S1}} \cdot P_{S2,c}^{1-\omega_{S1}}$.

Data Center Sector. Data centers produce computation using labour and electricity in fixed proportions (Leontief).[16]

$$q = \min\left(\frac{\ell}{\lambda_L}, \frac{e}{\text{PUE}_c \cdot CF_{DC} \cdot 8760}\right), \tag{18}$$

Total employment for a facility of capacity $K_{DC,c}$ Megawatts (MW) is $\lambda_L K_{DC,c}$ workers, the total electricity consumption is $\text{PUE}_c \cdot CF_{DC} \cdot 8760 \cdot K_{DC,c}$ Mwh per year.

CF_{DC} is the utilisation rate of installed capacity, which converts nameplate megawatts into annual energy, this is set fixed for each type of facilities.

PUE_c is the county-specific Power Usage Effectiveness predicted from local climate, which affects the marginal electricity consumption for each unit of IT capacity.[17] λ_L is the per-MW employment parameter, which is taken from the median value among industry estimates.

A data center operator in location c sells computation at the local compute price $p_{M,c}$, hires at $w_{S2,c}$; and it buys electricity at the delivered, tax-inclusive rate

$$p_{E,c}^{DC} = \left(1 + t_c^{E,DC}\right)\left(p_{E,c} + \varrho_c\right), \tag{19}$$

where $p_{E,c}$ is the wholesale price that the dispatch model of Section 6.2 delivers, ϱ_c is the retail delivery margin, and $t_c^{E,DC}$ is the data center electricity tax. We describe the main tax instruments for data centers (the electricity tax $t_c^{E,DC}$ and the property-and-equipment charge ϕ_c^K used below) in Section 6.6.

Compute Prices Across Space. Computation is tradable with quality degradation due to increased latency. Output from d delivered to c arrives at quality $\Phi_{dc} \leq 1$, with $\Phi_{cc} = 1$, where Φ_{dc} is the latency-dependent quality of service on the $d \to c$ route. The compute price index faced by service firms in c is therefore

$$P_{M,c} = \left[\sum_{d=1}^{C} n_{DC,d} \cdot \left(\frac{p_{M,d}}{\Phi_{dc}}\right)^{1-\eta}\right]^{\frac{1}{1-\eta}}, \tag{20}$$

where $\eta > 1$ is the elasticity of substitution across source locations.

[16] I provide evidence in support of the assumption that electricity and labour are the main inputs of production in the data center sector in the Appendix, by estimating hedonic price regressions relating variation in data center pricing to variation in electricity tariffs, facilities' Power Usage Effectiveness, and local wages. I find that electricity tariffs and PUE values are the main significant correlates of data center facility prices, supporting this functional form assumption.

[17] I provide more details of the prediction of county-specific PUE_c in the Appendix.

Flow Profit. Annual profit of a facility of capacity $K_{DC,c}$ is revenue net of the two variable inputs, the equipment tax, and a common fixed cost:

$$\pi_c^{DC} = \underbrace{p_{M,c}\, q_c}_{\text{(i) Revenue}} - \underbrace{w_{S2,c}\, \lambda_L\, K_{DC,c}}_{\text{(ii) labour}} - \underbrace{p_{E,c}^{DC} \cdot \text{PUE}_c \cdot CF_{DC} \cdot 8760 \cdot K_{DC,c}}_{\text{(iii) electricity}} - \underbrace{\phi_c^K \cdot 10^6 \cdot K_{DC,c}}_{\text{(iv) property and equipment tax}} - \underbrace{F_{DC}}_{\text{(v) fixed cost}}. \tag{21}$$

Location Choice. We assume that entry is a static, atomistic, non-strategic problem. Each operator simultaneously chooses where to enter, taking the location-specific compute prices and marginal costs of operation as given. I hold the national total fixed: aggregate US data center capacity, $\sum_c K_{DC,c}$, is exogenous, and entry reallocates it across space.[18]

A potential entrant j evaluating location c receives payoff $\pi_{jc} = \bar{\pi}_c + \delta_c + \varepsilon_{jc}$. The ex-ante profit, including local tax burdens and potential tax rebates, is:

$$\bar{\pi}_c = \pi_c^{\text{var}} - t_c^{\text{prop}}, \qquad t_c^{\text{prop}} = \frac{\phi_c^K \cdot 10^6}{R_{\text{perMW}}^{DC}}, \tag{22}$$

t_c^{prop} is the ad-valorem equipment/capital tax, expressed as a burden rate on annual revenue per MW. [19] The term δ_c is a location fixed effect capturing unobserved location-specific attributes not included, such as water rights, permitting speed, proximity to an existing campus, which is recovered by inversion so that the model reproduces the observed distribution of capacity. Before making location choices, we assume that each developer draws an unobserved profit shock, which follows an Extreme Value Type-I Distribution, $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$, standard deviation is denoted σ_ϵ. Finally, we assume that the equilibrium probability of location is constrained by ζ_c, the fraction of county c's land area suitable for data center construction, built from IM3's rasterised data (at 100 meter resolution) on terrain, flood risk, and proximity to transmission and water bodies. The inclusion of suitability weights directly prevents the model from making erroneous predictions that data centers could locate in areas with protected/restricted land from development, or areas that are already built by 2024.

Given the EV Type-I assumption about the unobserved profit shocks, the equilibrium probability that a given data center is located in c being:

$$\Pr(c) = \frac{\zeta_c \cdot \exp\big((\bar{\pi}_c + \delta_c)/\sigma_\varepsilon\big)}{\sum_{c'} \zeta_{c'} \cdot \exp\big((\bar{\pi}_{c'} + \delta_{c'})/\sigma_\varepsilon\big)}. \tag{23}$$

Electricity Sector: Hourly Dispatch Model. A key innovation of the model is to capture spatially differentiated effects of data center locations on welfare through its impacts on the electricity grid. The response to new electricity load depends on which units are next in the merit order, which varies enormously across regions and across hours within a region. Following Arkolakis and Walsh (2023), I embed an hourly economic dispatch of the US thermal plant fleet, which produces location-specific merit-order supply curves as a function of local electricity demand. I assume that each Balancing-Authority solves constrained cost-minimisation problems every hour and the marginal cost of the most expensive generating unit in the merit-order framework delivers the predicted wholesale prices. Appendix C provides more details on the construction, validation and robustness of the model.

[18] This closes the aggregate margin that the fixed cost F_{DC} would otherwise determine.

[19] Normalising the capital charge is necessary because π_c^{var} is compared across locations in proportional terms while ϕ_c^K is a dollar rate.

Marginal cost. Each unit i in hour t is assigned a short-run marginal cost built from fuel costs, heat rates and pollution permit prices,

$$mc_{i,t} = vom_i + hr_i \times fp_{f(i),s(i),t} + ec_{i,t}, \tag{24}$$

with hr_i the heat rate measured from EPA continuous-emissions-monitoring (CEMS) data, fp the state-month fueld price, and ec the potential emissions prices for pollutants under power plant cap-and-trade schemes (set to zero at baseline).

The model uses all dispatchable, fossil-fuel power plants reporting hourly to CEMS: 3,222 units at 1,129 plants and 734.8 GW, across 45 balancing authorities, covering 60% of metered US demand and 95% of operating fossil generation. We define the cap capacity as the maximum of nameplate and the 99th percentile of observed generation, after excluding planned outages. To account for the presence of non-dispatchable generation sources (nuclear, hydro, and renewablws), we use for estimation the total hourly electricity load after excluding non-dispatchable generation, so the merit order responds only to the residual thermal target.

The hourly problem. In each hour, the system planner minimises variable cost subject to meeting every region's observed thermal generation $G_{r,t}$ from units inside it,

$$\min_{\{g_{i,t}\}} \sum_i mc_{i,t}\, g_{i,t} \quad \text{s.t.} \quad \sum_{i \in \mathcal{I}_r} g_{i,t} = G_{r,t} \;\; \forall r, \qquad 0 \le g_{i,t} \le \bar{C}_{i,t}. \tag{25}$$

Because the constraints are separable across region-hours and units enter linearly, the solution is the exact merit order. The hourly price is the marginal cost of the most expensive dispatched unit, $p_{b,t} = \max_i\{mc_{i,t} : g_{i,t} > 0\}$; I annualise it weighting by observed load and add a balancing-authority wedge μ_b fixed in the base year so $\bar{p}_b + \mu_b$ matches the observed industrial tariff. All counterfactual price *changes* thus come from the dispatch, while the base-year level matches the data by construction.

The supply curve. The model delivers for each Balancing-Authority, the predicted electricity tariff $p_{b,t} + \mu_b$, as a function of observed BA-level electricity load. In any counterfactual equilibrium, we let k denote the ratio of counterfactual demand to the baseline equilibrium demand, and the supply curve can be indexed using the observed ratio ($\bar{p}_b(\kappa)$).

6.3 Local Pollution Externalities

The dispatch model above determines which generating units meet electricity demand in each Balancing Authority. This section combines the model with an integrated pollution dispersal model, the AP3 model, to translate plant-specific emissions from electricity generation, into equilibrium spatial distribution of pollution concentrations, and monetised environmental damages from pollution externalities.

From Generation to Emissions. Given power plant i's generation to meet electricity demand $G_i(D_b(i))$, the total annual emissions of pollutant p are

$$e_{i,p} = r_{i,p} \cdot G_i(D_b(i))\,, \tag{26}$$

where $r_{i,p}$ is a plant-level emission rate in metric tons per MWh. I include five major local pollutants: $p \in \{\mathrm{PM}_{2.5}, \mathrm{NO}_x, \mathrm{SO}_2, \mathrm{NH}_3, \mathrm{VOC}\}$: the first three are primary emissions and direct precursors, and ammo-

nia and volatile organic compounds are included because secondary $PM_{2.5}$ formation depends on them.[20]

Each unit is also assigned a stack-height tier $h(i)$. Stack height matters for a reason distinct from the quantity emitted: the same ton released from a tall stack is carried further before it settles, and disperses over a wider and typically less densely populated area, so tier changes the *mapping* from emissions to concentrations rather than merely scaling it. A short-stack unit in a populated county and a tall-stack unit in the same county emitting identically do not cause the same damage.

Next, I use the AP3 integrated assessment model (Muller, 2025), to construct source-receptor matrices, that determines how a per-unit change in pollution *emissions* translates into changes in pollution *concentrations*, and monetises the value of pollution concentrations experienced in each location. Let $S_{c'c}^{h,p}$ be the damage in dollars borne by residents of receptor county c per metric ton of pollutant p emitted from a tier-h source located in county c'. Each matrix already composes atmospheric dispersion, the resulting change in ambient concentration, the exposed population, the concentration-response function, and the monetary value of the health endpoints. The total local environmental damamge in county c from electricity generation is the sum over pollutants and over all source counties,

$$D_c \;=\; \sum_{p} \sum_{c'} S_{c'c}^{h,p}\; e_{p,c'}, \tag{27}$$

with the tier-specific matrix applied to the share of county c''s emissions released at each stack height.

Carbon Emissions from Electricity Generation. The electricity dispatch model also predicts the total level of CO_2 emissions from electricity consumption. National emissions are the power plant-level sum:

$$Z_{CO_2}(D_{b(i)}) = \sum_i r_i^{CO_2}\, G_i\Big(D_{b(i)}\Big) \tag{28}$$

with $r_i^{CO_2}$ the unit's carbon rate in tons per MWh. Total carbon emissions is monetised at a social cost of carbon for reporting, but is excluded from household utility in the baseline specification: it is a global externality, so including it in local amenities would misattribute a national damage to the residents of particular counties. I report changes in tonnage and in monetised carbon damage in counterfactual exercises, and report the specification that includes carbon damages into welfare as a robustness exercise.

6.4 Productivity and Amenity Externalities

A key channel through which data centers affect local economies, beyond their direct demand for labour and electricity, is via spillover effects on the productivity of neighbouring firms and on the amenity value of the location for residents. I assume that local TFP is affected by labour force through classical agglomeration forces, and that data centers do not generate local productivity spillovers. However, local data center capacity is allowed to flexibly account local amenities, such as due to noise pollution or local public goods via higher tax revenues.

Productivity Externalities. Total factor productivity in sector $s \in \{T, S1, S2\}$ at location c is:

$$A_{s,c} = \bar{A}_{s,c} \cdot L_c^{\mu_L} \tag{29}$$

[20] Rates come from a harmonised plant-level emissions frame rather than from the monitoring data alone, because continuous emissions monitoring covers NO_x, SO_2 and CO_2 but never measures particulate matter, ammonia or volatile organic compounds directly; those rates are constructed from the National Emissions Inventory at the process level. Appendix C documents the construction and the pollutant coverage.

where $\bar{A}_{s,c}$ is the baseline productivity fundamental and μ_L is the agglomeration spillover parameters with respect to the size of the total labour force in c, following Redding and Rossi-Hansberg (2017).

Amenity Spillovers. Local amenity depends on local data center capacity in the following form:

$$B_c = \bar{B}_c \cdot (K_{DC,c} + 1)^{\beta_B} \tag{30}$$

where β_B is the amenity spillover of data center capacity. The sign of β_B is theoretically ambiguous: positive forces include local employment growth, infrastructure investment, and tax revenue; negative forces include congestion, noise, and visual disamenity. The amenity spillover is conceptually distinct from the pollution-damage term $g(D_c)$: the pollution channel operates through electricity-induced emissions and the source-receptor matrix, while the amenity channel captures direct effects of DC presence on residential quality of life unrelated to grid emissions.

In counterfactual hat algebra, the effective amenity hat combines the amenity spillover channel and the environmental damage channel:

$$\hat{B}_c = \hat{g}_c \cdot \left(\frac{K^{\text{cf}}_{DC,c} + 1}{K^{\text{obs}}_{DC,c} + 1} \right)^{\beta_B} \tag{31}$$

where $\hat{g}_c = g(D^{\text{cf}}_c)/g(D^{\text{obs}}_c)$ is the change in environmental quality.

6.5 Wages and Market Access

Sector-specific wages are determined by a Firm Market Access (FMA) equation derived from the free-entry and gravity conditions. For sector $s \in \{T, S1, S2\}$:

$$w_{s,c} = \left(\frac{\text{FMA}_{s,c} \cdot A_{s,c}}{L_{s,c} \cdot p_{E,c}^{\xi_{E,s} \cdot \theta_{\text{tr}}} \cdot P_{M,c}^{\xi_{M,s} \cdot \theta_{\text{tr}}}} \right)^{\frac{1}{1+\beta_{L,s} \cdot \theta_{\text{tr}}}} \tag{32}$$

where $\theta_{\text{tr}} = \sigma - 1$ is the trade elasticity, $\xi_{E,s}$ and $\xi_{M,s}$ are the electricity and computation input shares, and $\text{FMA}_{s,c}$ is the Firm Market Access:

$$\text{FMA}_{s,c} = \sum_d \tau_{cd}^{-\theta_{\text{tr}}} \cdot \frac{\text{Ex}_{s,d}}{\text{CMA}_{s,d}} \tag{33}$$

Here $\text{Ex}_{s,d}$ is expenditure on sector s goods in location d and $\text{CMA}_{s,d}$ is the Consumer Market Access.

6.6 Local Taxation

Local governments levy ad valorem tax rates, each a proportional wedge on a price or spending base. There are five instruments, all introduced above where they enter each agent's problem: a household consumption (sales) tax τ_c on the tradable-and-service composite (housing exempt); a residential electricity tax $t_c^{E,H}$; a firm electricity tax t_c^E; a data center electricity tax $t_c^{E,DC} \leq t_c^E$; and a data center property-and-equipment wedge ϕ_c^K. There is no revenue tax on tradable or service firms and no corporate profit tax: these producers are competitive and earn zero economic profit, so the only fiscal wedge they face is on their electricity input.

Data center incentives as rate discounts. State incentive programs lower the effective tax rates faced by operators. The net property/M&E rate is the gross statutory rate less the abatement granted, $\phi_c^K = \phi_c^{K,\text{full}} - (\text{abatement})_c \leq \phi_c^{K,\text{full}}$, and the electricity tax exemption sets $t_c^{E,DC} \leq t_c^E$. Removing incentives in

a counterfactual restores $\phi_c^K \to \phi_c^{K,\text{full}}$ and $t_c^{E,DC} \to t_c^E$. The rates are calibrated at county-year grain from the Net Present Value of state-level incentive packages (electricity, equipment, and infrastructure sales-tax exemptions), annualised over a 15-year horizon at a 7% discount rate (NPV factor = 9.11), combined with combined state-plus-local sales-tax rates.

Government budget. Local tax revenue sums the five instruments over their bases and is rebated per worker:

$$\text{Rev}_c = \underbrace{\tau_c C_c}_{\text{sales}} + \underbrace{t_c^E R_c^{E,\text{firm}} + t_c^{E,DC} R_c^{E,DC} + t_c^{E,H} R_c^{E,H}}_{\text{electricity}} + \underbrace{\phi_c^K \cdot 10^6 \cdot K_{DC,c}}_{\text{DC property/M\&E}}, \qquad T_c = \frac{\text{Rev}_c}{L_c}. \tag{34}$$

where C_c is the taxable consumption composite, $R_c^{E,\text{firm}}, R_c^{E,DC}, R_c^{E,H}$ are firm, data center, and residential electricity spending, and $K_{DC,c}$ is installed capacity. Incentives enter through the reduced ϕ_c^K and $t_c^{E,DC}$ that multiply the bases, not as a separate $-\text{subsidy}_c$ term, so a high-incentive county mechanically collects less and rebates a smaller T_c, with no double counting.

6.7 Spatial Equilibrium

An equilibrium consists of sector-location allocations:
$\{L_{s,c}, n_{s,c}, K_{DC,c}, D_{E,c}, D_c\}$ and prices $\{w_{s,c}, p_{E,c}, P_{T,c}, P_{S1,c}, P_{S2,c}, P_{M,c}, P_{H,c}, T_c\}$ such that:

(i) migration shares satisfy equation (12);

(ii) labour markets clear in each sector-location;

(iii) free entry holds in tradable and service sectors;

(iv) data center allocation satisfies the logit conditions;

(v) electricity markets clear at the BA level, with the wholesale price given by the dispatch of Section C.3.1;

(vi) the population constraint $\sum_c L_c = \bar{L}$ holds;

(vii) compute markets clear in each sector-location;

(viii) local pollution damages are consistent with equilibrium electricity demand via the source-receptor matrix;

and (viv) amenity and productivity spillovers are consistent with equilibrium data center allocations. Appendix Section C.8 outlines the solution algorithm to solve for the equilibrium.

6.8 Parameter Estimation & Calibration

I use a combination of calibrated parameter values from the existing literature and conduct estimation of key parameters relating to data center location choice using my data. Table 4 summarises parameter values.

Table (4) Model Parameters

Parameter	Description	Value	Source
Externally calibrated			
σ	Elasticity of substitution (T, S)	5	Head and Mayer (2014)
μ_L	Labour agglomeration externality	0.05	Redding and Rossi-Hansberg (2017)
μ	Tradable expenditure share	0.30	BEA NIPA
θ	Service expenditure share	0.40	BEA NIPA
γ	Housing supply elasticity	location-specific	Baum-Snow and Han (2024)
ι	Inverse migration elasticity	1.0	Jaworski et al. (2023)
α_L, α_E	Factor shares (Tradable Sector)	Normalised from BEA I-O tables	
$\beta_{L,S1}, \beta_{E,S1}, \beta_{M,S1}$	Factor shares (S1 Sector)	Normalised from BEA I-O tables	
$\beta_{L,S2}, \beta_{E,S2}, \beta_{M,S2}$	Factor shares (S2 Sector)	Normalised from BEA I-O tables	
λ_E	Labour elasticity (E)	0.70	EIA data
λ_L	Labour per MW (DC)	1.05	Industry estimates
CF_{DC}	DC capacity factor	0.65	Industry estimates
PUE_c	Power usage effectiveness	location-specific	CERF-DC raster data
$SR_{p,c,c'}$	Source-receptor matrices	—	AP3 model
Estimated from data (Section 6.9)			
$\Phi_{o,d}$	Bilateral compute quality	—	Server-Pair Latency Data
β_B	Amenity spillover of DC capacity	0.016	Bilateral migration regression
σ_ε	dispersion (DC location profit)	1.840 \$M/MW	Conditional logit on entry
η	DC elasticity of substitution	1.295	Gravity on firm hosting flows
Recovered via model inversion			
$A_{T,c}, A_{S1,c}, A_{S2,c}$	Sector productivities	—	Match sectoral employment
$A_{E,c}\bar{K}_{E,c}$	Electricity fundamentals	—	Match BA-level generation

Factor shares (α_L, α_E) for tradables and $(\beta_{L,s}, \beta_{E,s}, \beta_{M,s})$ for each service tier are normalised from BEA Input-Output tables to sum to unity within each sector (Cobb-Douglas constant returns to scale). This normalisation is stored in a shared parameter file and loaded by the solver.

Given observed data on $\{L_{s,c}, w_{s,c}, p_{E,c}, K_{DC,c}\}$, I recover location fundamentals $\{A_{T,c}, A_{S1,c}, A_{S2,c}, \bar{B}_c\}$ via model inversion. The inversion uses the wage equation (32) and the spatial equilibrium conditions, inverting sequentially: productivities from wages and FMA, amenities from the migration equilibrium condition.

6.9 Identification of Key Parameters

6.9.1 Network Quality Index ($\Phi_{\ell,c}$)

The bilateral network quality index ($\Phi_{\ell,c}$) shapes the effective quality of a firm in an origin o procuring data compute services from destination d, and can be thought of as implicit trade costs in the data center service sector. The quality could be determined by factors such as geographic distance between two locations, and availability and quality of fiber networks.

I use server-pair latency data from WonderNetwork (83 US servers, 2014–2024) to estimate the bilateral network quality index $\Phi_{\ell,c}$ for each orogin and destination county, proceeding in two steps.

First, using server-pair latency data, I estimate the following gravity equation:

$$\ln(\text{Latency})_{o,d,t} = \alpha_{1,o} + \alpha_{2,d} + \mu \ln(\text{Distance})_{o,d} + \varepsilon_{o,d,t} \tag{35}$$

where $\alpha_{1,o}$ and $\alpha_{2,d}$ are origin and destination server fixed effects.

I include two different measures of distance between servers, the first is the straight-line distance between server coordinates, the second computes the shortest distance it takes to travel from one server to another, via the existing configuration of terrestrial fiber networks in the US.

Table (5) Gravity Estimates: Server-Pair Latency and Distance

	(1) GC only	(2) Fiber only	(3) Both	(4) Gravity
$\ln D^{GC}$	0.602***	—	0.072	—
	(0.019)		(0.086)	
$\ln D^{Fiber}$	—	0.621***	0.548***	—
		(0.020)	(0.091)	
$\ln \text{Pop}_o$			—	-0.019***
				(0.005)
$\ln \text{Pop}_d$			—	-0.023**
				(0.009)
$\ln \text{Area}_o$			—	0.084***
				(0.005)
$\ln \text{Area}_d$			—	0.077***
				(0.005)
Same state			—	-0.929***
				(0.017)
R^2	0.725	0.728	0.728	0.120
N	44153	41478	41478	41394
Origin FE	Yes	Yes	Yes	No
Dest FE	Yes	Yes	Yes	No
Year FE	Yes	Yes	Yes	Yes

Notes: Dependent variable: $\ln(\text{latency}_{ijt})$, where latency is the mean measured round-trip time (ms) between WonderNetwork servers. Geographic distance (D_{ij}^{GC}) is Haversine (great-circle). Fiber distance (D_{ijt}^{Fiber}) is the Dijkstra shortest path over Infrapedia terrestrial cable routes, recalculated yearly as cables enter service (2014–2024). Column (4) replaces distance with pure gravity covariates: log county population, log county land area, and a same-state indicator, without origin or destination fixed effects. Standard errors (parentheses) are two-way clustered by origin and destination server. Columns (2) and (3) exclude server-pair-years that are unreachable in that year's fiber graph. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5 presents the distance elasticity of latency from the two different measures. Column 1 uses geographic distance, column 2 uses fiber network shortest distance, column 3 includes both measures of distance. Entered separately, the two elasticities are almost identical and both precisely estimated (0.602 and 0.621). Entered together in column 3, fiber network distance retains an elasticity of 0.548 and remains significant at the 1% level, while the coefficient on geographic distance falls to 0.072 and is no longer distinguishable from zero. Conditional on the route a packet actually travels through the fibre network, straight-line distance carries no additional information about latency, which is what the physical propagation constraint implies. Column 4 replaces distance with pure gravity covariates—log county population, log county land area, and a same-state indicator—and drops the server fixed effects. The substantially lower R^2 in this specification confirms that spatial connectivity, not local demographic mass, is the primary driver of server-pair latency.

Second, I predict the *level* of latency of sending data from any given origin, by using the estimated origin & destination fixed effects. I generate predicted location fixed effects at the county level using observable infrastructure characteristics (distance to nearest fiber node, broadband density, population density).

Finally, the predicted quality index is:

$$\hat{\Phi}_{o,d,t} = \exp\left(\hat{\alpha}_{1,o} + \hat{\alpha}_{2,d} + \hat{\delta}_t + \hat{\mu} \ln(\text{Distance}_{o,d})\right) \tag{36}$$

Appendix Figure A.16 maps the resulting county-level network latency and the implied computation-quality index $\hat{\Phi}$. The index is highest in the dense coastal corridors and lowest in remote interior counties, mirroring the geography of the fibre backbone in Appendix Figure A.15.

6.9.2 Elasticity of Substitution Across DC Source Locations (η)

The parameter η governs how substitutable computation services from different locations are from the perspective of firms. Under CES aggregation across source locations, the expenditure share directed from a firm in origin county c to data-center location ℓ is:

$$s_{\ell,c} = \frac{K_\ell^{DC} \cdot mc_\ell^{1-\eta} \cdot \Phi_{\ell,c}^{\eta-1}}{\sum_k K_k^{DC} \cdot mc_k^{1-\eta} \cdot \Phi_{k,c}^{\eta-1}} \tag{37}$$

Taking logs and absorbing origin and destination terms into fixed effects yields the gravity equation:

$$\ln s_{\ell,c} = \underbrace{(1-\eta)\ln mc_\ell + \ln K_\ell^{DC}}_{\alpha_\ell \text{ (destination FE)}} + (\eta-1)\ln \Phi_{\ell,c} - \underbrace{\ln \Omega_c}_{\alpha_c \text{ (origin FE)}} \tag{38}$$

The coefficient on $\ln \hat{\Phi}_{\ell,c}$ identifies $\eta - 1$, so that $\hat{\eta} = \hat{\beta} + 1$.

I estimate this gravity equation using county-pair data on firm hosting choices. For each origin county c, I observe the number of firms choosing to host in destination county ℓ, matched to the predicted network quality index $\hat{\Phi}_{\ell,c}$ from the latency gravity model (Section 6.9.1). Origin and destination fixed effects absorb all county-level confounders. Since $\ln N_{\ell,c} = \ln s_{\ell,c} + \ln N_c$ and $\ln N_c$ is absorbed by the origin fixed effect, regressing log flow counts on $\alpha_c + \alpha_\ell + \beta \ln \hat{\Phi}_{\ell,c}$ recovers the same $\hat{\beta}$ as using log shares directly.

Table (6) Network Quality and Firm Hosting Choices

	(1) Distance only	(2) $\hat{\Phi}_A$	(3) $\hat{\Phi}_B$
$\ln d_{\ell c}$ (geographic)	-0.177^{***}		
	(0.007)		
$\ln \hat{\Phi}_A$		0.295^{***}	
		(0.011)	
$\ln \hat{\Phi}_B$			0.295^{***}
			(0.051)
R^2	0.687	0.687	0.687
Observations (pairs)	65,964	65,964	65,964
Origin FE	✓	✓	✓
Destination FE	✓	✓	✓

Notes: OLS estimates of Equation (38). Dependent variable is $\ln N_{\ell c}$, the log number of firms headquartered in origin county c that host data in destination county ℓ. $\hat{\Phi}_A$ and $\hat{\Phi}_B$ are alternative network quality indices from the latency gravity model (Section 6.9.1). Column (3) uses fiber network distance to predict network quality (Φ_B); its coefficient is suppressed for brevity. Standard errors clustered by origin county in parentheses. Sample: Tier-1 cloud and direct-hosted US firms, 2024. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Results. Table 6 reports the main estimates. Column (1) confirms that geographic distance is a strong, precisely estimated negative predictor of hosting location ($\hat{\beta}_d = -0.177$, SE $= 0.007$), consistent with latency-sensitive demand. Column (2) replaces distance with the network quality index $\hat{\Phi}_A$, yielding $\hat{\beta} = 0.295$ (SE $= 0.011$) and an implied elasticity of substitution $\hat{\eta} = 1.295$. The near-identical overall

R^2 across columns (1) and (2) indicates that $\hat{\Phi}_A$ and geographic distance capture a similar amount of residual variation within the fixed effects, which is expected given that network quality is partly a function of physical distance. Column (3) instead $\hat{\Phi}_B$, an alternative network quality index constructed using fiber network distance, not geographic distance. The coefficient on Φ remains the same. The estimate $\hat{\eta} = 1.295$ is the baseline used in the calibration of the quantitative spatial model.

Identification of η using cost pass-through . As a second complementary approach that does not require demand data, I use the markup passthrough relationship. Under CES demand with monopolistic competition, the model implies $p_{M,\ell t} = \frac{\eta}{(\eta-1)} \cdot mc_{\ell t}$.

A regression of data center price on marginal costs identifies $\mu \equiv \eta/(\eta-1)$ and therefore $\eta = \mu/(\mu-1)$.

I estimate this using a metro-by-provider-by-year panel of colocation pricing for 2014–2025 from Telegeography, constructing marginal cost as $\hat{mc}_{\ell t} = \lambda_L \cdot w_{\ell t} + \text{PUE}_{\ell t} \cdot p_{E,\ell t}$ with λ_L being the workers per MW. The preferred specifications in Columns (3)–(4) of Table 7, which include metro $\times$ year fixed effects to absorb demand-side confounders, imply $\hat{\eta} = 1.342$ (contemporaneous, s.e. $= 0.037$) and $\hat{\eta} = 1.326$ (one-year lag). Without fixed effects (Columns 1–2) the estimates are larger ($\hat{\eta} \approx 1.79$–1.95), consistent with demand shocks positively co-varying with marginal costs when metro trends are uncontrolled. The baseline calibration uses $\eta = 1.295$ from the hosting-flows gravity in Table 6; the markup passthrough estimates of 1.33–1.34 sit just above it and provide an over-identifying check on the CES–monopolistic-competition assumption, since they use only supply-side variation and no demand data.

Table (7) DC Price Pass-Through to Marginal Cost: Identifying η

	No fixed effects		Metro $\times$ year FE	
	(1) Contemp.	(2) Lag 1	(3) Contemp.	(4) Lag 1
Total MC (contemp.)	2.271***		3.920***	
	(0.100)		(0.316)	
Total MC (1-yr lag)		2.056***		4.068***
		(0.041)		(< 0.001)
Implied $\hat{\eta}$	1.787	1.947	1.342	1.326
SE($\hat{\eta}$)	0.062	0.037	0.037	—
Observations	65	41	65	41

Panel: operator $\times$ year. MC in USD/kW/month (labour + electricity cost).
$\hat{\eta} = \hat{\mu}/(\hat{\mu}-1)$; SE by delta method. "—" = not estimable (singular design matrix).
Stars test $H_0: \mu = 0$; every estimated markup also rejects $H_0: \mu = 1$ (perfect competition) at $p < 0.01$.
Standard errors clustered at metro level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

6.9.3 DC Location Dispersion (σ_ε)

The dispersion parameter σ_ε governs how sensitive DC location choice is to cost differences. I estimate it from conditional logit/Poisson models of data center entry against time-varying determinants, exploiting variation in state-level tax incentive packages. Specifically, I compute the Net Present Value of each state's tax incentive package for each county-year, and estimate:

$$n_{DC,c,t} = \exp\left(\mathbf{X}'_{c,t}\boldsymbol{\beta}/\sigma_\varepsilon + \xi_c + \delta_t\right) + \varepsilon_{c,t} \tag{39}$$

where $\mathbf{X}_{c,t}$ includes electricity prices, broadband density, and the annualised tax incentive value. The coefficient on the NPV variable, expressed in \$/MW/yr, identifies $1/\sigma_\varepsilon$ when the profit function is specified in the same units.

Data and results. The entry panel covers 2014–2024, with DC counts from SP Global. Tax incentive NPVs are annualised at a 7% discount rate over a 15-year horizon (NPV factor = 9.11). Regressors include the two-year lagged values of the local commercial electricity price, county broadband density, and the annualised tax NPV in \$/MW/yr.

Table (8) Poisson Fixed Effects Regressions of Data Center Entry, County-Year Panel (2014-2024)

	(1)	(2)	(3)
NPV of total tax saving per MW IT capacity	0.5436*	0.8940**	1.1930*
	(0.315)	(0.392)	(0.622)
Industrial Electricity Rate (t-2)		-0.0013	-0.0051
		(0.002)	(0.003)
High-Speed Internet Connections (t-2)			0.0003
			(0.002)
Implied $\hat{\sigma}_\varepsilon$ (\$M/MW)	1.840	1.119	0.838
County FE	✓	✓	✓
Year FE	✓	✓	✓
Observations	6990	5548	2522

Poisson fixed-effects estimates of county data center entry counts, 2014–2024.
Standard errors clustered at the county level in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The baseline estimate is $\hat{\sigma}_\varepsilon = \$1.84\text{M/MW}$ in NPV units, equivalent to approximately \$200K/MW/yr at the calibrated discount factor. To interpret this magnitude: a \$1M/MW NPV tax advantage (broadly the scale of a generous state incentive package) raises a county's share of DC capacity by $\exp(1/1.84) - 1 \approx 72\%$ relative to its baseline share, holding all other determinants fixed. This implies moderate but meaningful location sensitivity, consistent with the cross-state dispersion in observed DC density.

6.9.4 Amenity Spillover (β_B)

The model allows data center capacity to flexibly affect local amenities. On the positive side, data centers generate fiscal revenues (property and sales taxes, utility payments) and can prompt infrastructure improvements, raising the attractiveness of the host county to workers. On the negative side, they introduce noise, visual disamenity, increased heavy-vehicle traffic, and elevated water consumption. Whether people want to move *toward* or *away from* DC-dense counties — conditional on wages — identifies the sign and magnitude of β_B. The parameter is theoretically unrestricted; the data determine its sign.

Estimating equation. The model's bilateral migration equation pins down the structural estimating equation. Given Fréchet moving costs with shape parameter ι, the ratio of the bilateral out-migration share from origin i to destination j relative to the stay-in-i share satisfies:

$$\log\left(\frac{\pi_{ij,t}}{\pi_{ii,t}}\right) = \iota \log\left(\frac{w_{j,t}/P_{j,t}}{w_{i,t}/P_{i,t}}\right) + \beta_B\left[\log(K_{DC,j,t}+1) - \log(K_{DC,i,t}+1)\right] + \xi_{ij} + \phi_t + \varepsilon_{ij,t} \tag{40}$$

where $P_c = P_{T,c}^{\mu} P_{S,c}^{\theta} P_{H,c}^{1-\mu-\theta}$ is the full Cobb-Douglas consumption price index. Origin-by-destination fixed effects ξ_{ij} absorb time-invariant bilateral frictions κ_{ij} and the amenity-ratio fundamentals $\bar{B}_j / \bar{B}_i$; year fixed effects ϕ_t absorb common macro trends. The key identifying variation is within-pair changes in DC capacity differentials over time.

I impose $\iota = 1$ following Jaworski et al. (2023) and treat the real-wage differential as a control, identifying β_B from the residual co-movement of migration flows with the DC capacity differential. Table 9 reports results across six measures of DC capacity, comprising two continuous transformations and four binary facility-type indicators.

Results. The baseline estimate in Column (1), using the log-capacity ratio on county pairs both with positive DC capacity, is $\hat{\beta}_B = 0.016$ (s.e. $= 0.007$), significant at the 5% level. Column (2) extends to the full sample using an asinh transformation and yields a similar estimate of 0.013 (s.e. $= 0.005$). Binary DC-entry indicators in Columns (3)–(6) show that colocation and enterprise DC entry raises migration ratios by roughly 2–3 percentage points; hyperscale entry in Column (6) is small and imprecise (0.007, s.e. 0.017), consistent with hyperscale facilities generating less local economic activity per MW.

Table (9) Bilateral Migration Responses to Data Center Capacity

	(1)	(2)	(3)	(4)	(5)	(6)
	Log ratio	Asinh ratio	Any DC	Colocation	Enterprise	Hyperscale
Panel A: Continuous DC capacity (destination − origin)						
$\log(K_{DC,j}/K_{DC,i})$	0.016**					
	(0.007)					
$\text{asinh}(K_{DC,j}+1) - \text{asinh}(K_{DC,i}+1)$		0.013***				
		(0.005)				
Panel B: Facility-type indicators (destination − origin)						
$D_{\text{any},j} - D_{\text{any},i}$			0.020***			
			(0.006)			
$D_{\text{colo},j} - D_{\text{colo},i}$				0.028**		
				(0.014)		
$D_{\text{ent},j} - D_{\text{ent},i}$					0.023***	
					(0.007)	
$D_{\text{hyp},j} - D_{\text{hyp},i}$						0.007
						(0.017)
Observations	506,347	1,563,326	1,563,326	1,563,326	1,563,326	1,563,326
R^2	0.972	0.978	0.978	0.978	0.978	0.978
Pair (O×D) FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	$K_j, K_i > 0$	Full	Full	Full	Full	Full

Notes: Dependent variable: $\log(\pi_{ij}/\pi_{ii})$, the log ratio of the bilateral migration share from i to j over the stay-in-i share. All specifications include origin×destination pair FE and year FE, with real wage ratios absorbed by the pair FE (imposing $\iota = 1$). Standard errors clustered by origin state × destination state. Column (1) restricted to county pairs where both counties have positive DC capacity; columns (2)–(6) use the full sample. Panel A estimates a continuous log and asinh DC capacity differential; Panel B estimates binary indicators for county pairs where at least one county has a facility of the given type. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.10 Goodness of fit

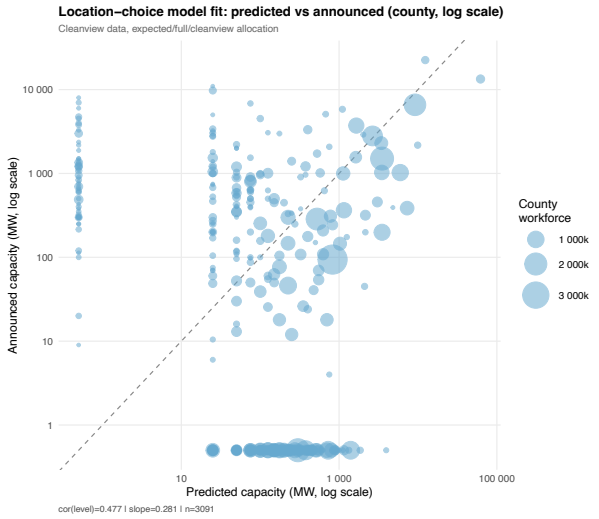
The baseline equilibrium of the model is calibrated to exactly match the cross-sectional distribution of sectoral employment and wages. In this subsection, we examine the goodness of fit of the model for two sets of moments not directly targeted: the county- and state-level distribution of predicted data center entry, the balancing-authority level electricity tariffs, and total electricity load across balancing authorities. The model provides reasonable fit in the across-county and across-state allocation of proposed data center capacity, as well as the across-BA variation in electricity tariffs and total electricity demand.

6.10.1 Location choice: predicted vs. announced capacity

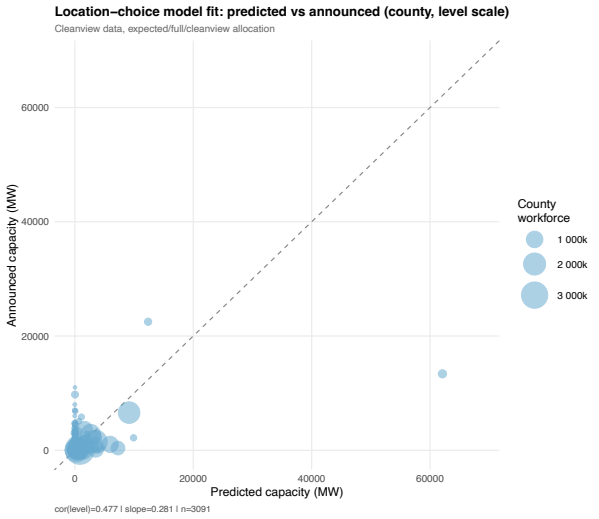
The first key moment is the predicted allocation of a given exogenous increase in data center capacity. We use the location-choice model to allocate the same total planned capacity from CleanView data, holding observed taxes and fundamentals fixed at their 2024 values. The model thus generates the predicted entrant capacity assuming that they were making entry decisions in 2024. We compare the predicted county (and state) allocation of megawatts against the announced locations. At the county level, the OLS regression of announced on predicted capacity gives $R^2 = 0.23$ (level correlation 0.48; rank and log-scale correlations 0.21 and 0.23; Figure 10, Table 10). At the state level, $R^2 = 0.62$ and the level correlation is 0.79. States with a large incumbent presence (Virginia above all) receive a large predicted share and a large share of announcements. Texas, Georgia, and Pennsylvania sit above the 45° line (more announced capacity than their incumbent share predicts), while California sits below it.

Figure (10) Location-choice fit: predicted vs. announced planned capacity

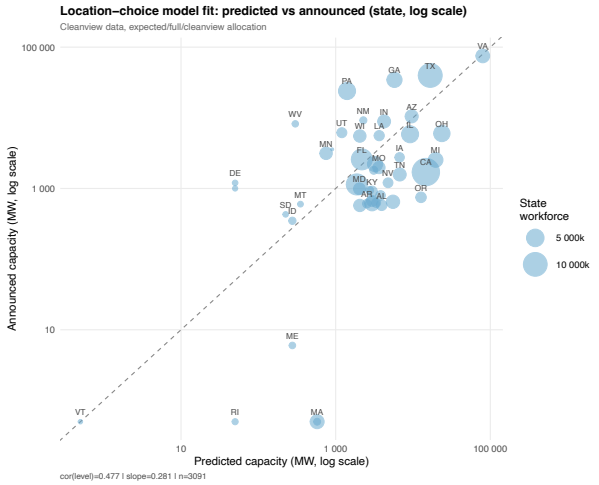
(a) County, log scale



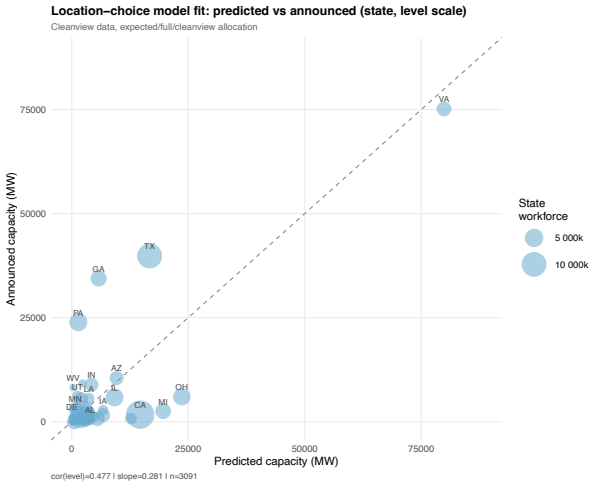
(b) County, level scale



(c) State, log scale



(d) State, level scale



Notes: Predicted county (state) allocation of the announced total planned capacity, holding observed taxes and fundamentals fixed, vs. the county (state) each facility was actually announced in. Bubble size is proportional to county (state) workforce. Dashed line is the 45° reference. Fit statistics are reported in Table 10.

Table (10) Location-choice fit metrics, predicted vs. announced capacity

Level	R^2 (OLS)	Corr. (level)	Corr. (log)	Spearman	N
County	0.227	0.477	0.231	0.209	3,091
State	0.624	0.790	0.625	0.480	47

Notes: Predicted vs. announced planned capacity (MW). R^2 is from a cross-sectional OLS regression of announced on predicted capacity (equal to the squared level correlation). Corr. is the Pearson correlation in levels and in logs; Spearman is the rank correlation. State aggregates sum county capacity within each state.

6.11 Policy Counterfactual: Removal of State-Level Tax Incentives

Given the model, I ask whether existing state-level tax incentives for data centers induce spatial misallocation.

Design. I remove state-level tax abatement, solving for the new spatial equilibrium. I hold constant the total national capacity addition from the proposed expansion pipeline, and ask how the removal of tax rebates induces re-sorting of data centers. The incremental tax revenue collected is rebated as an equal lump-sum transfer to workers in each affected state.

Mechanisms. Removing the abatements works through four channels. The *fiscal channel*: counties that currently grant abatements lose their cost advantage and may lose data centers, while the additional statutory revenue collected is returned to workers. The *reallocation channel*: data centers move toward counties with lower operating costs or better fundamentals (suitability, latency), which can increase allocative efficiency. The *electricity-price channel*: counties that lose data centers see lower electricity demand and prices, benefiting their remaining firms, while counties that gain them face higher prices. Non data-center firms may consequently sort across locations, amplifying the welfare effects. The *environmental channel*: if capacity shifts between grids of different carbon and pollution intensities, both global emissions and local damages change. The net welfare effect depends on whether the abatement-driven allocation sits closer to or farther from the social optimum than the allocation the market would choose under a level tax.

Welfare Definition. For any counterfactual equilibrium E' relative to a baseline E, I measure welfare by the per-location log compensating equivalent variation:

$$CEV_c(E \to E') = \log V_c^{E'} - \log V_c^{E}, \tag{41}$$

where V_c is the indirect utility of a worker in location c.

Welfare Decomposition. I split the change in welfare into three channels:

$$\frac{\Delta W}{W} = \underbrace{\frac{\Delta I}{I} - \left[\mu \frac{\Delta P_T}{P_T} + \theta \frac{\Delta P_S}{P_S} + (1-\mu-\theta)\frac{\Delta P_H}{P_H}\right]}_{\text{Real wages}} + \underbrace{\frac{\Delta g}{g}}_{\text{Pollution}} + \underbrace{\frac{\Delta B}{B}}_{\text{Amenity}}, \tag{42}$$

with each term population-weighted. The real-wage channel is nominal income (inclusive of the lump-sum rebate) net of the local cost of living, which absorbs the pass-through of electricity and compute prices; the pollution channel aggregates changes in local air-pollution damage across counties; and the amenity channel captures the local amenity value of data-center capacity. The equivalent variation $EV = \bar{L} \cdot \bar{I} \cdot (V'/V - 1)$ gives the dollar value of the welfare change.

6.12 Counterfactual Results

Business-As-Usual Counterfactual: Proposed Data Center Growth. I begin with the announced expansion pipeline. Figure 11 shows the location of proposed facilities (orange circles, sized by planned MW) against existing operational facilities (black dots). Proposed capacity is highly clustered and tracks the existing footprint: Virginia and its neighbors, Texas, a broad concentration across the Southeast (Georgia in particular), and the Ohio Valley.

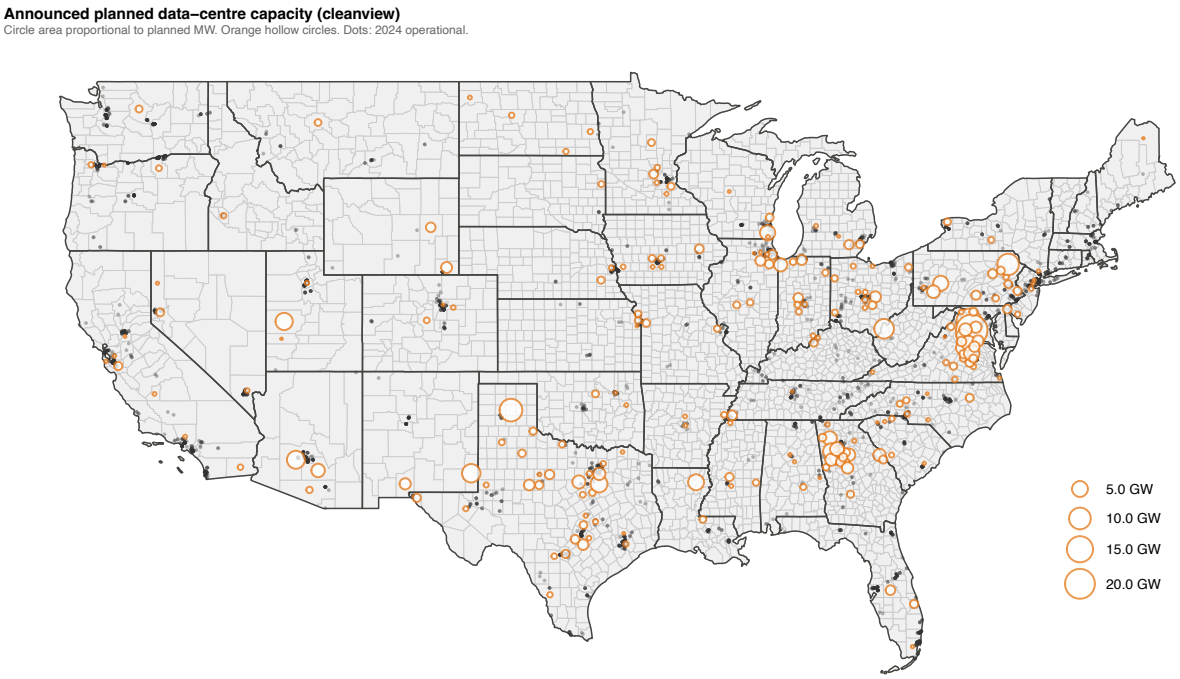


Figure (11) Announced planned data-center capacity. Orange hollow circles are planned facilities, sized by planned MW; black dots are 2024 operational facilities.

Taking the announced locations as given, I assume all proposed facilities become operational and compare the resulting equilibrium to the baseline. This is the Business-As-Usual (BAU) counterfactual $(E_A - E_0)$.

Table 11 presents the estimated welfare impacts of the BAU counterfactual equilibrium, relative to the 2024 baseline equilibrium. Each row is a population subgroup. *Aggregate* is the population-weighted national total; the next four rows split workers by sector (tradables, high-compute services, low-compute services, and the unemployed); the final two rows split counties by whether they host any data-center capacity in the counterfactual allocation. Each column is a component of welfare. The first two columns report the headline change, in dollars per year ($B/yr) and as a percentage of baseline income (ΔCEV%). The next three columns decompose the dollar figure into the channels of equation (42): Amenity, Pollution (local air pollution), and Wages (real wages). To a first-order approximation these three sum to the headline figure. The Electricity Price columns report the change in the average electricity price faced by consumers, in levels ($/MWh) and percent; this is a mechanism that enters the Wages channel rather than a separate welfare term. The CO_2 columns are left blank in this pipeline run (see the closing paragraph).

BAU results. National welfare rises by $4.977 billion per year, or 0.214% of 2024 national income. The gain is concentrated by location: workers in host counties gain $6.657 billion (0.395% of their baseline income), while workers in non-host counties lose $1.680 billion (0.976%). Across sectors, the per-worker gain is largest for high-compute services (+0.325%) and smallest for tradables (+0.039%); in dollar terms the low-compute service sector accounts for the most (+$2.964 billion) because it is the largest by employment.

The decomposition explains the host/non-host gap. The amenity gain (+$17.215 billion) accrues entirely to host counties: amenities respond only to local data-center capacity, so the amenity figure is identical for the Aggregate and DC Host rows and zero for non-host counties. Against this, both the pollution chan-

nel (−$6.266 billion) and the real-wage channel (−$6.283 billion) are negative and fall on host and non-host counties alike, because higher electricity prices and added pollution spill beyond the counties that host the facilities. Host counties keep the amenity gain and so come out ahead; non-host counties bear only the costs.

Carbon. The last two columns report the global externality that the welfare columns exclude. Serving the announced capacity raises national CO_2 emissions from electricity generation by 685.1 million metric tons per year, 14.0% of the 4,904 million tons of energy-related US emissions in 2025.[21] At a social cost of carbon of $145 per metric ton this is $99.3 billion per year of damage, twenty times the $4.977 billion welfare gain measured in the first column. The ordering is not close enough for the choice of carbon price to matter: the expansion would still be welfare-reducing at any social cost of carbon above roughly $7 per ton. The sign of the BAU result therefore depends on whether carbon is priced at all, not on how it is priced.

Table (11) Welfare effects — C1. BAU planned capacity addition ($E_A - E_0$).

	Total Welfare		Welfare Change Decomposition ($B/yr)			Electricity Price		CO_2 (national)	
	$B/yr	ΔCEV%	Amenity	Pollution	Wages	Δ$/MWh	Δ%	ΔMMT	Damage $B
Aggregate	+4.977	+0.214%	+17.215	-6.266	-6.283	+5.3419	+8.505%	+685.062	+99.334
Tradables	+0.219	+0.039%	+1.682	-0.677	-0.821	+5.7616	+9.148%	—	—
High-Computation	+1.771	+0.325%	+4.558	-1.495	-1.362	+5.2082	+8.298%	—	—
Low-Computation	+2.964	+0.201%	+10.906	-4.071	-4.074	+5.3190	+8.471%	—	—
Unemployment	+0.197	+0.043%	+1.004	-0.387	-0.449	+5.7706	+9.014%	—	—
DC Host Counties	+6.657	+0.395%	+17.215	-5.497	-5.359	+5.2510	+8.366%	—	—
Non-DC Host Counties	-1.680	-0.976%	+0.000	-0.770	-0.924	+5.9393	+9.418%	—	—

Notes: Total Welfare $B/yr equals (to first order) the sum of the three decomposition channels. Amenity $= \sum_c \left[\beta_B \log\left(\frac{K_c^{cf}+1}{K_c^{base}+1}\right)\right] \cdot \bar{w}_c L_c$ with $\beta_B = 0.016$. Pollution $= \sum_c \log(g_c^{cf}/g_c^{base}) \cdot \bar{w}_c L_c$ and covers *local* air-pollution damage only. Wages $= \sum_c \log(\bar{w}_c^{cf,\text{real}}/\bar{w}_c^{base,\text{real}}) \cdot \bar{w}_c L_c$. DC Host Counties defined by counterfactual DC allocation ($K_c^{cf} > 0$). Sector rows dollar-weight welfare by sector baseline employment. CO_2 columns: ΔMMT is the change in national emissions from electricity generation in million metric tons; Damage $B $= \Delta\text{MMT} \times 10^6 \times \$145/\text{tCO}_2$ (EPA social cost of carbon), so the two carry the same sign and a positive entry is an *increase* in damage. Carbon is a global externality: the pipeline reports the same national total on every row, so it is shown here for the Aggregate row only, and it is *not* included in the Total Welfare columns (see Section 6).

The maps show how these effects are distributed across space. Figure 12 maps the per-worker welfare change. Most counties are shaded light red, indicating small losses, with concentrated gains (blue) in and around host counties. The map and the positive aggregate are consistent: widespread small losses are outweighed by large gains in a smaller set of host locations.

[21] Source: US Energy Information Administration, https://www.eia.gov/environment/emissions/carbon/.

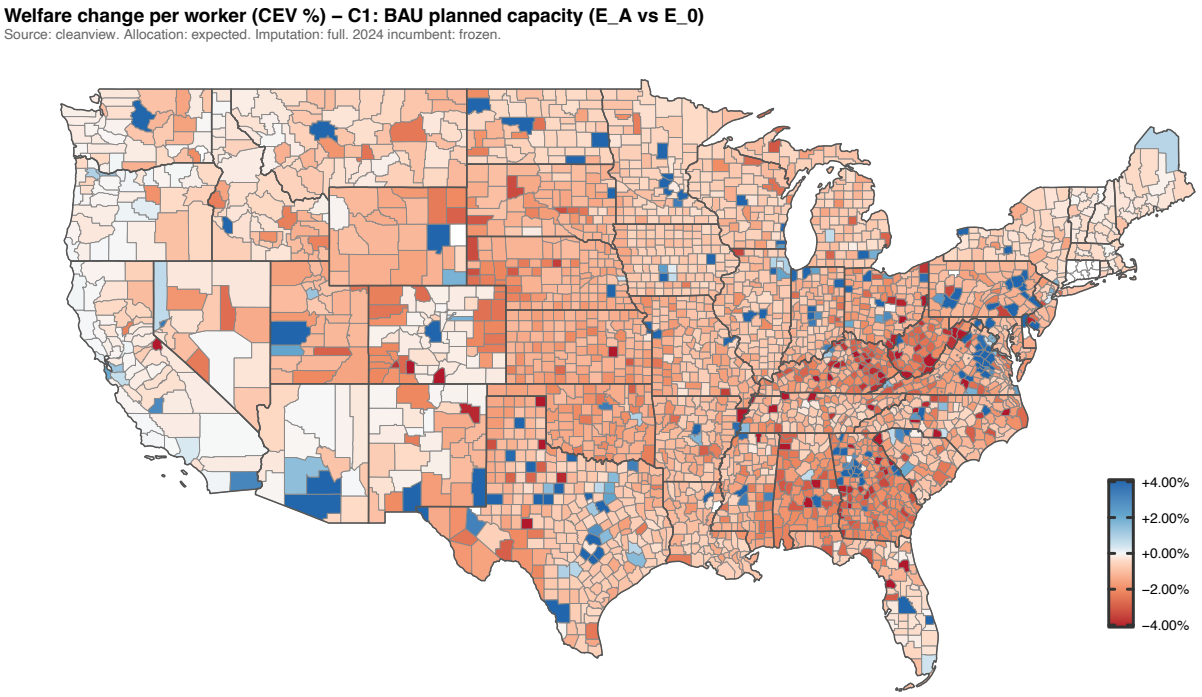


Figure (12) C1 (BAU): welfare change per worker (CEV, %) by county.

Figure 13 maps the change in local environmental damage, which rises (blue) across most of the eastern and midwestern United States as new capacity draws additional fossil generation onto the grid; a handful of Pacific-Northwest counties see damage fall.

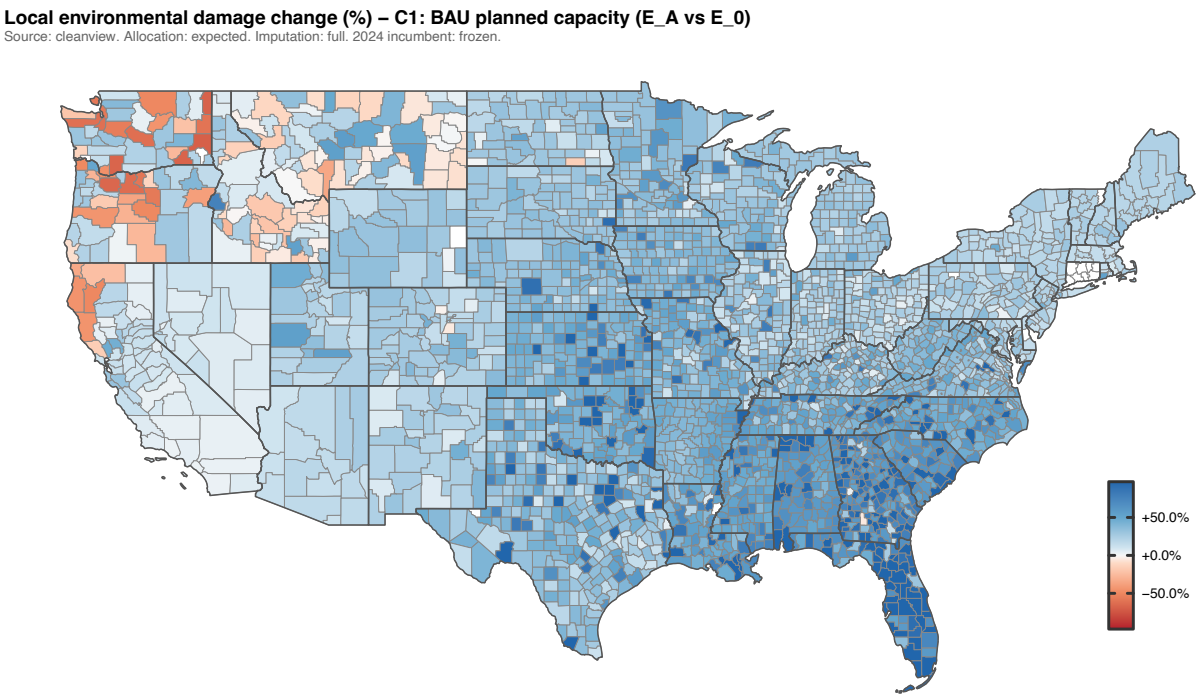


Figure (13) C1 (BAU): change in local environmental damage (%) by county.

Figure 14 maps the change in electricity prices. Prices rise everywhere, most in the Plains, the Ohio Valley, and the Southeast (Georgia and Alabama), where the largest blocks of planned capacity locate; the population-weighted aggregate increase is 8.5% (+$5.34/MWh).

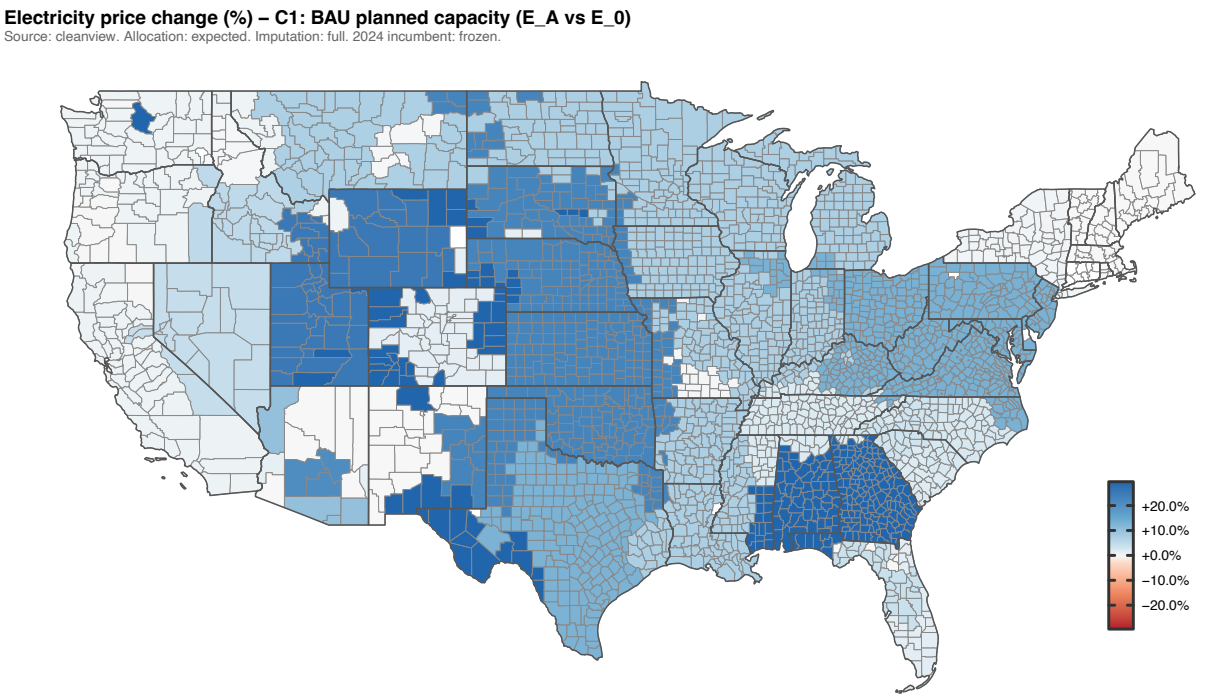


Figure (14) C1 (BAU): change in electricity price (%) by county.

Welfare Effects of Removing State-Level Tax Incentives. I next hold the size of the expansion pipeline fixed but assume developers no longer receive tax rebates when making their location choice decisions. This frees up local revenues that could have been spent on data center tax incentives, which are rebated equally as lump-sum transfers to workers in affected states. Because the total capacity added is the same as in the BAU case, comparing this equilibrium to the BAU equilibrium isolates whether existing state-level abatements induce spatial misallocation of data centers.

Figure 15 maps the resulting change in capacity. Proposed data centers in some counties are responsive to the removal of tax incentives— most visibly those in Northern Virginia hub — moving to other low-cost locations, with California being the main predicted destination that sees an increase in proposed data center capacity.

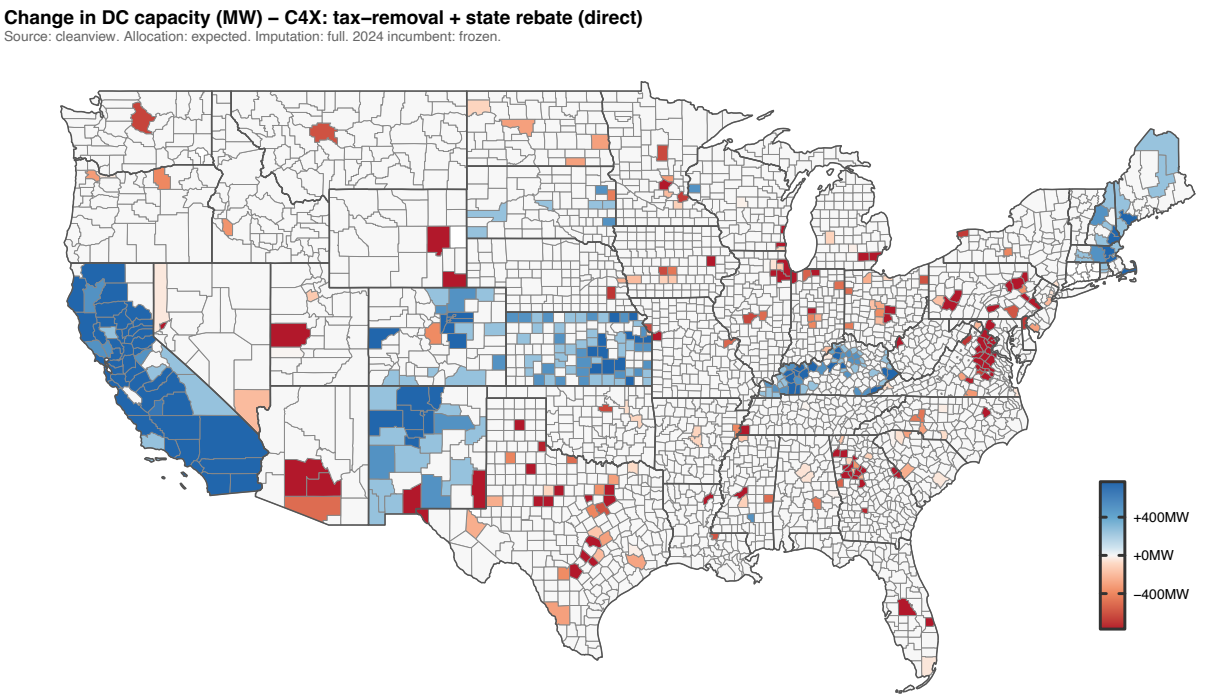


Figure (15) C4X (tax removal + rebate): change in data-center capacity by county.

Table 12 reports the welfare effect. Removing the abatements raises national welfare by $9.985 billion per year (0.492% of baseline income) — roughly double the BAU gain. The gain comes from a $6.961 billion rise in amenity values, a $2.177 billion gain through the pollution channel, and a $0.208 billion rise in real wages. The real-wage channel is the sum of two large and nearly offsetting forces: reallocation lowers real wages by $2.575 billion (moving capacity out of large incumbent hubs costs more in wages than an equivalent gain in a smaller destination is worth), while the lump-sum rebate of the $2.78 billion in incremental statutory revenue returns almost exactly that amount to workers, leaving a small net positive. The pollution channel turns *positive* here, in contrast to the BAU case: reallocation moves capacity out of the dirtier Southeastern grids, where the abatements had concentrated it, into cleaner Western ones, lowering net local damage.

Carbon moves in the same direction as local pollution, and for the same reason. Because the reallocation shifts capacity off the fossil-heavy Southeastern grids and onto cleaner Western ones, national CO_2 emissions from electricity generation fall by 56.2 million metric tons per year, about 1.1% of energy-related US emissions, worth a further $8.1 billion per year at $145 per ton. Adding this to the $9.985 billion measured in the welfare columns gives a total of $18.1 billion per year. Carbon accounting therefore reinforces the case for removing the incentives, in contrast to the BAU expansion, where it reverses the sign: the two counterfactuals differ not only in the size of the welfare change but in whether the environmental externality works with or against it.

The split by host status mirrors the BAU pattern but with sharper price effects. Counties that gain data centers under the full-rate allocation gain $10.128 billion (0.575%), while counties left without capacity lose only $0.143 billion (0.071%). The host-county real-wage channel is negative (−$1.174 billion) because gaining capacity raises local electricity prices (+5.26%), whereas non-host counties see prices *fall* (−6.65%) as capacity departs, turning their real-wage channel positive (+$1.382 billion); together with the rebate this nearly offsets the amenity loss they post (−$2.033 billion) from shedding capacity they had held under

BAU, leaving them close to break-even.

Table (12) Welfare effects of removing tax incentive + direct rebate (C4X, $= (E_B - E_A) +$ rebate).

	Total Welfare		Welfare Change Decomposition ($B/yr)			Electricity Price		CO_2 (national)	
	$B/yr	ΔCEV%	Amenity	Pollution	Wages	Δ$/MWh	Δ%	ΔMMT	Damage $B
Aggregate	+9.985	+0.492%	+6.961	+2.177	+0.208	+6.7383	+3.729%	-56.179	-8.146
Tradables	+0.943	+0.490%	+0.476	+0.255	+0.149	+4.8204	+1.961%	—	—
High-Computation	+2.857	+0.505%	+2.415	+0.519	-0.236	+7.2759	+4.209%	—	—
Low-Computation	+6.148	+0.488%	+4.047	+1.395	+0.292	+6.8510	+3.839%	—	—
Unemployment	+0.866	+0.756%	+0.702	+0.134	-0.005	+7.1108	+4.216%	—	—
DC Host Counties	+10.128	+0.575%	+8.995	+1.762	-1.174	+8.4386	+5.259%	—	—
Non-DC Host Counties	-0.143	-0.071%	-2.033	+0.415	+1.382	-4.7947	-6.649%	—	—

Notes: Total Welfare $B/yr equals (to first order) the sum of the three decomposition channels. Amenity $= \sum_c \left[\beta_B \log\left(\frac{K_c^{cf}+1}{K_c^{base}+1}\right)\right] \cdot \bar{w}_c L_c$ with $\beta_B = 0.016$; the baseline for this contrast is the BAU allocation E_A, so counties that shed capacity register a negative amenity term. Pollution $= \sum_c \log(g_c^{cf}/g_c^{base}) \cdot \bar{w}_c L_c$ and covers *local* air-pollution damage only. Wages $= \sum_c \log(\bar{w}_c^{cf,\text{real}}/\bar{w}_c^{base,\text{real}}) \cdot \bar{w}_c L_c$, inclusive of the state lump-sum rebate. DC Host Counties defined by counterfactual DC allocation ($K_c^{cf} > 0$). Sector rows dollar-weight welfare by sector baseline employment. CO_2 columns: ΔMMT is the change in national emissions from electricity generation in million metric tons; Damage $B $=$ ΔMMT $\times 10^6 \times$ \$145/tCO_2 (EPA social cost of carbon), so the two carry the same sign. Both are negative here: the reallocation lowers emissions, and the $-$\$8.146 billion is a *reduction* in damage and therefore a welfare gain of that size. Carbon is a global externality: the pipeline reports the same national total on every row, so it is shown here for the Aggregate row only, and it is *not* included in the Total Welfare columns.

Two features set this counterfactual apart from the BAU case. First, because total capacity is fixed, removing the abatements *redistributes* electricity demand across space rather than adding to it: prices fall where capacity leaves and rise where it arrives. Figure 18 shows the pattern, with sharp declines in Georgia, Alabama, and Texas — the states whose abatement-driven capacity relocates — and increases in California, the northern Plains, and New England. The population-weighted aggregate still rises (+3.7%) because the destinations are, on net, more populated and higher-priced grids. Second, the reallocation moves capacity toward cleaner grids on net, so local pollution damage *falls* in aggregate even as it rises in the Western counties that gain capacity; Figure 17 shows damage falling (red) across the Southeast and rising (blue) across the Mountain West and Great Plains.

Figure 16 maps the per-worker welfare change, which is more dispersed than under BAU and shows gains across much of the country, with the largest concentrations in California and the vacated Virginia hub.

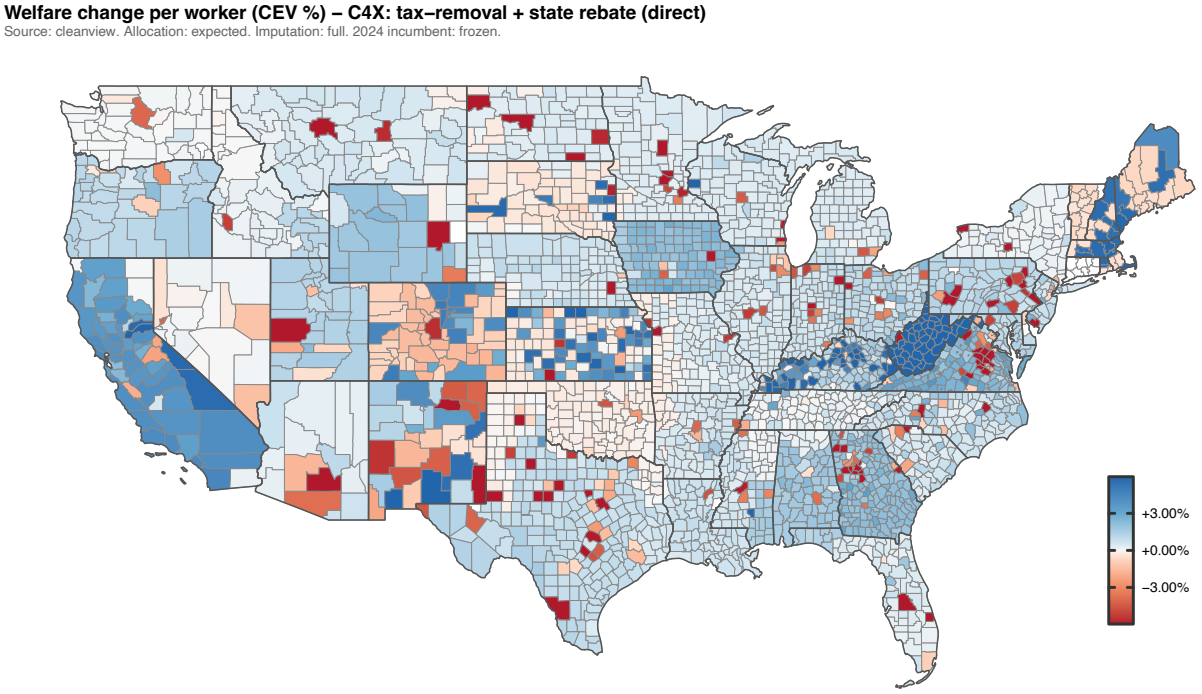


Figure (16) C4X (tax removal + rebate): welfare change per worker (CEV, %) by county.

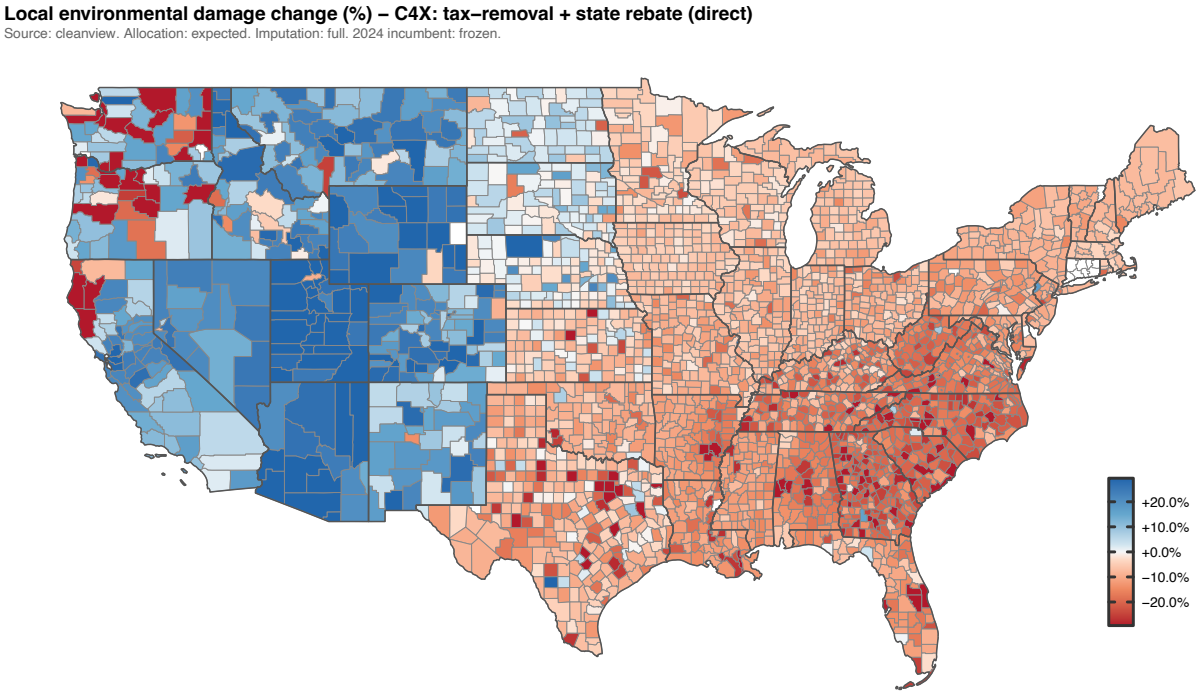


Figure (17) C4X (tax removal + rebate): change in local environmental damage (%) by county. Damage falls (red) across the Southeast, where capacity departs, and rises (blue) in the West, where it arrives; the net is a reduction.

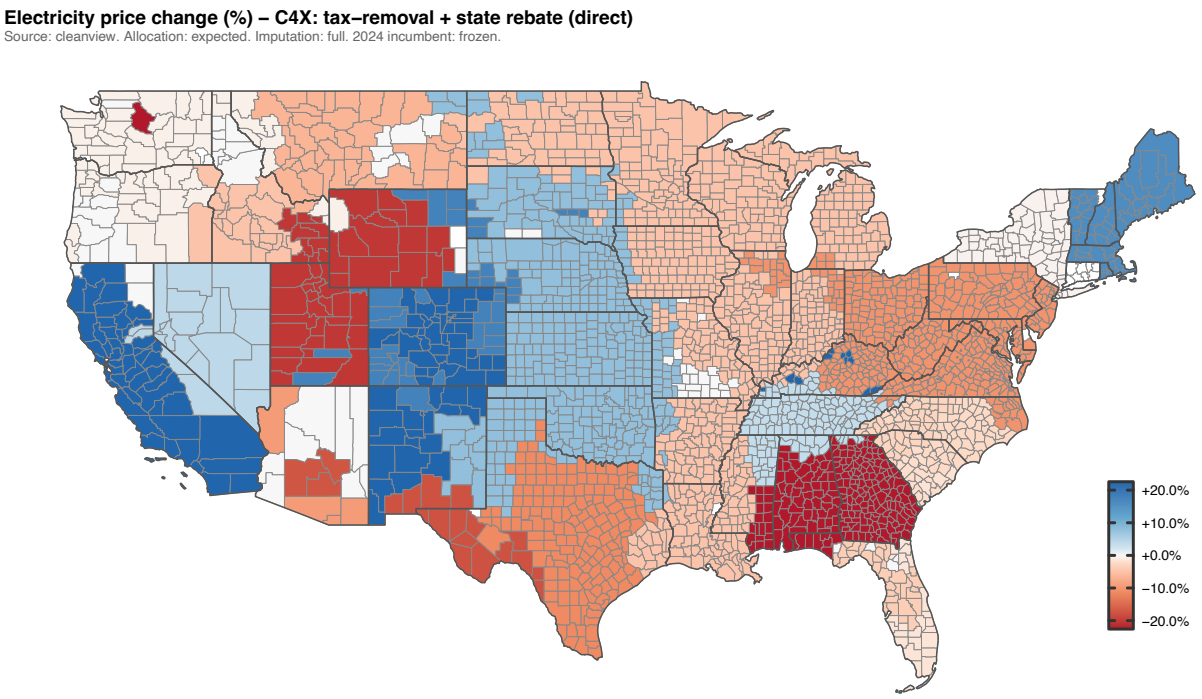


Figure (18) C4X (tax removal + rebate): change in electricity price (%) by county.

Importance of general-equilibrium adjustments. The welfare evaluation above rests on three general-equilibrium channels through which data-center location affects the welfare of local residents: the electricity-market congestion channel, the local air-pollution channel, and the compute-supply (tradability) channel. To gauge how much these adjustments matter for the policy conclusion, I recompute the incentive-removal counterfactual (C4X), holding the policy shock and the total capacity addition fixed but switching off one channel at a time:

1. *No electricity-price response.* Electricity tariffs are held fixed as data centers re-sort under no tax rebates.
2. *No local air-pollution damage.* Local pollution damage is frozen at its observed 2024 level.
3. *Perfectly tradable computation.* Computation is uniformly tradable across space with no quality degradation

Table 13 reports the resulting welfare effects: column 1 reproduces the full-model C4X result and columns 2–4 correspond to the three shutdowns.

First, switching off the electricity market congestion channel nearly triples the aggregate gain from removing incentives, from \$9.99 billion (0.49% of income) to \$28.57 billion (1.59%). In the full model, with total capacity fixed, the reallocation raises electricity prices in the grids that gain data centers by more than it lowers them where capacity departs (a population-weighted +3.7%); this congestion cost is borne by households and non-data-center firms and offsets much of the direct gain. Suppressing the response removes that offsetting cost, making the policy appear far more beneficial than it is. The same channel shapes incidence: in the full model, counties that lose data centers enjoy *falling* electricity prices (−6.6%), which cushions them to a near-break-even loss of only 0.07%; once the price response is switched off, that relief disappears and their loss deepens to 0.52%.

Table (13) Welfare effects of removing state tax incentives (C4X) under the full model and with each general-equilibrium channel switched off.

Subgroup	Measure	Full model	General-equilibrium channel switched off		
			No electricity price response	No local air pollution damage	Perfectly tradable compute
Aggregate	ΔCEV (%)	+0.492	+1.594	+1.560	+1.856
	EV ($B/yr)	+9.985	+28.567	+27.688	+26.280
Tradables	ΔCEV (%)	+0.490	+1.500	+1.463	+2.396
	EV ($B/yr)	+0.943	+2.693	+2.597	+3.572
High-compute svcs.	ΔCEV (%)	+0.505	+1.692	+1.651	+1.438
	EV ($B/yr)	+2.857	+7.959	+7.714	+4.989
Low-compute svcs.	ΔCEV (%)	+0.488	+1.574	+1.542	+1.921
	EV ($B/yr)	+6.148	+17.809	+17.273	+17.612
Unemployed	ΔCEV (%)	+0.756	+1.760	+1.713	+2.129
	EV ($B/yr)	+0.866	+2.007	+1.987	+1.883
DC-host counties	ΔCEV (%)	+0.575	+1.958	+1.911	+1.904
	EV ($B/yr)	+10.128	+29.594	+28.646	+26.427
Non-host counties	ΔCEV (%)	−0.071	−0.518	−0.485	−0.517
	EV ($B/yr)	−0.143	−1.027	−0.958	−0.146

Notes: Each cell reports the welfare effect of removing state-level data-center tax incentives and rebating the recovered revenue as a state lump-sum transfer (contrast C4X $= (E_B - E_A) +$ rebate), relative to the business-as-usual expansion E_A. Column 1 is the full model; columns 2–4 recompute the same counterfactual with one general-equilibrium channel switched off: (2) electricity tariffs held unresponsive to data-center resorting, (3) local air-pollution damage frozen at its 2024 level, and (4) computation made perfectly tradable across space (no latency/quality friction). For each subgroup the first row is the population-weighted per-worker compensating variation (ΔCEV, % of baseline income) and the second row is the aggregate equivalent variation ($B/yr). Headline specification: cleanview / expected / full / frozen.

The other two channels are individually almost as important. Freezing local pollution damage raises the aggregate gain to $27.69 billion (1.56%): in the full model the counties that gain capacity also absorb higher local pollution, which deters in-migration and dampens the reallocation's real-wage gains, so removing this disamenity lets workers and firms re-sort more freely and inflates the measured gain. Making computation perfectly tradable raises the gain to $26.28 billion (1.86%): latency ties data centers and compute-using firms to particular locations in the full model, and removing that friction lets capacity chase low-cost grids more aggressively—the tradable-goods sector, which benefits most from cheap ubiquitous compute, sees its gain rise the most, from 0.49% to 2.40%. As with the electricity channel, both shutdowns also deepen the losses borne by non-host counties, to roughly 0.5%.

Taken together, these exercises show that the three general-equilibrium channels are quantitatively first-order. Switching off any single channel inflates the estimated national welfare gain from removing incentives by a factor of 2.6 to 2.9: from 0.49% of income to between 1.56% and 1.86%, while raising the per-worker loss to communities that lose data centers roughly sevenfold, from 0.07% to about 0.5%. A partial-equilibrium assessment that omitted any one of these channels would overstate the aggregate benefits of reform roughly threefold and understate its costs to the places that lose capacity. Accounting for these general-equilibrium adjustments is therefore essential to a credible welfare evaluation of where data centers locate.

7 Conclusion

Data centers are a rapidly expanding form of digital infrastructure whose spatial allocation has consequences for local economies, for regional air quality, and for national carbon emissions. This paper measures those consequences for the United States, using facility-level records on the near-universe of data centers matched to county outcomes from 2000 to 2023, and uses them to evaluate the place-based policies that govern where facilities are built.

Three empirical findings emerge. First, data center location is determined primarily by telecommunications infrastructure. Broadband density, proximity to fibre nodes, and low transmission latency are the strongest predictors of first entry, ahead of electricity costs, and counties in the upper half of the latency distribution are significantly less likely to host a facility. Proximity to potential clients in ICT and finance also predicts entry, as do state tax incentives. Damages from local electricity generation do not. Absent policy, private siting decisions do not internalise the environmental cost of the power the facilities consume, which is the fact the rest of the paper builds on.

Second, entry produces a large and persistent expansion of local service activity. Via an event-study design, I find that total employment rises to roughly 400 workers per 10,000 residents ten years after entry, averaging 7.8% of the pre-entry mean over the first decade. Services account for almost all of it. The precisely estimated gains are in professional services, real estate, administrative services, and construction; the information sector responds positively but becomes statistically distinguishable from zero only from the seventh year, and the finance response is imprecise throughout. Establishment counts rise by 4.4% of their pre-entry mean, so most of the employment response works through hiring within existing establishments rather than through firm entry. Manufacturing employment is unchanged, while manufacturing establishment counts fall by 2.7% of the pre-entry mean, significant at the 10% level. Both gross migration flows rise and net migration turns positive, so part of the employment gain is met by workers moving in rather than by local residents alone.

Third, these effects are unevenly distributed, and the quantitative model shows that the distribution matters as much as the aggregate. Under the announced capacity pipeline, national welfare rises by $4.98 billion per year, 0.214% of national income, but workers in host counties gain 0.395% of their income while

workers in counties that host nothing lose 0.976% of theirs, because the amenity gain is local while higher electricity prices and additional pollution travel across the grid and downwind. Two adjustments reverse the sign of that aggregate. The expansion adds 685 million metric tons of CO_2 per year, 14% of energy-related US emissions, worth $99 billion at $145 per ton and twenty times the measured gain. Setting the amenity spillover to zero also turns the gain into a loss of $12.4 billion. On the evidence here, the announced expansion is welfare-improving only if carbon goes unpriced and the amenity channel is taken at face value.

The policy counterfactual holds the announced capacity fixed and asks only where it locates. Removing state-level data center tax incentives and rebating the recovered revenue raises national welfare by $9.99 billion per year, 0.492% of income and roughly double the gain from the expansion itself, as capacity moves off the fossil-heavy Southeastern grids where the abatements had concentrated it. Local pollution damage falls by $2.18 billion per year and national CO_2 emissions by 56 million tons, worth a further $8.1 billion. Unlike the expansion result, this one does not rest on the amenity parameter: with the amenity channel switched off entirely the gain is still $3.33 billion per year, carried by lower pollution and higher real wages. Nor is it an artefact of any single general-equilibrium channel, though its magnitude depends on all of them. Switching off the electricity-price response, freezing local pollution damage, or making computation perfectly tradable each raises the estimated gain by a factor of 2.6 to 2.9 and deepens the loss to counties that give up capacity roughly sevenfold. A partial-equilibrium evaluation of the same policy would overstate its national benefit threefold and understate its local cost.

References

ALLEN, T. AND C. ARKOLAKIS (2014): "Trade and the topography of the spatial economy," *The Quarterly Journal of Economics*, 129, 1085–1140.

ALVAREZ, F. E., D. ARGENTE, J. CHOW, AND D. VAN PATTEN (2026): "Data Centers and Local Economies in the Age of AI: A Shift–Share Approach," Working Paper 35194, NBER.

ARKHANGELSKY, D., S. ATHEY, D. A. HIRSHBERG, G. W. IMBENS, AND S. WAGER (2021): "Synthetic difference-in-differences," *American Economic Review*, 111, 4088–4118.

ARKOLAKIS, C. AND C. WALSH (2023): "Clean growth," Tech. rep., National Bureau of Economic Research.

BAUM-SNOW, N. AND L. HAN (2024): "The microgeography of housing supply," *Journal of Political Economy*, 132, 1897–1946.

BECKER, S. O., S. HEBLICH, AND D. M. STURM (2021): "The impact of public employment: Evidence from Bonn," *Journal of Urban Economics*, 122, 103291.

BERAJA, M., D. Y. YANG, AND N. YUCHTMAN (2023): "Data-intensive innovation and the state: Evidence from AI firms in China," *The Review of Economic Studies*, 90, 1701–1723.

BERGÉ, L. (2018): "Efficient Estimation of Maximum Likelihood Models with Multiple Fixed-Effects: The R Package FENmlm," *CREA Discussion Papers*.

CALIENDO, L., M. DVORKIN, AND F. PARRO (2019): "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica*, 87, 741–835.

CALLAWAY, B. AND P. H. C. SANT'ANNA (2021): "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 225, 200–230.

CBRE (2024): "North America Data Center Trends H2 2024," Tech. rep., CBRE Research.

COLAS, M. AND E. REYNIER (2026): "Climate Policy with Electricity Trade," *Journal of Political Economy: Microeconomics*.

CONGRESSIONAL RESEARCH SERVICE (2025): "Data Centers: Federal Policy Issues," Tech. rep., Congressional Research Service.

DE CHAISEMARTIN, C. AND X. D'HAULTFŒUILLE (2020): "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 110, 2964–2996.

DEMIRER, M., D. J. J. HERNÁNDEZ, D. LI, AND S. PENG (2024): "Data, privacy laws and firm production: Evidence from the GDPR," Tech. rep., National Bureau of Economic Research.

FURMAN, J. (2025): "The Economic Case for Data Center Growth," .

GARGANO, A. AND M. GIACOLETTI (2025): "Subsidizing the Cloud: U.S. State Incentives to Data Centers," *SSRN Working Paper 5881105*.

GOIRI, Í., K. LE, J. GUITART, J. TORRES, AND R. BIANCHINI (2011): "Intelligent placement of datacenters for internet services," in *2011 31st International Conference on Distributed Computing Systems*, IEEE, 131–142.

GOLDFARB, A. AND C. TUCKER (2019): "Digital economics," *Journal of economic literature*, 57, 3–43.

GOODMAN-BACON, A. (2021): "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 225, 254–277.

GREENSTONE, M., R. HORNBECK, AND E. MORETTI (2010): "Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings," *Journal of political economy*, 118, 536–598.

HASBROUCK, J. AND G. SAAR (2013): "Low-Latency Trading," *Journal of Financial Markets*, 16, 646–679.

HEAD, K. AND T. MAYER (2014): "Gravity equations: Workhorse, toolkit, and cookbook," in *Handbook of international economics*, Elsevier, vol. 4, 131–195.

HEATH, D. R. (2019): "Prediction machines: the simple economics of artificial intelligence: by Ajay Agrawal, Joshua Gans and Avi Goldfarb, Published in 2018 by Harvard Business Review Press, 272 pp., 30.00(hardcover),KindleEdition: 16.19, ISBN: 978-1-633695672," .

HOLLAND, S. P., E. T. MANSUR, N. Z. MULLER, AND A. J. YATES (2016): "Are there environmental benefits from driving electric vehicles? The importance of local factors," *American Economic Review*, 106, 3700–3729.

HOLLAND, S. P., E. T. MANSUR, V. VERDIER, AND A. J. YATES (2024): "Regularization from economic constraints: A new estimator for marginal emissions," Tech. rep., National Bureau of Economic Research.

HOLLINGSWORTH, A., T. JAWORSKI, C. KITCHENS, AND I. RUDIK (2025): "Air Pollution and the Labor Market: Evidence from Wildfire Smoke," *Working Paper*.

HOLLINGSWORTH, A., C. KITCHENS, T. JAWORSKI, AND I. RUDIK (2026): "Economic geography and air

pollution regulation in the united states," *Journal of Political Economy Microeconomics*, 4, 000–000.
INTERNATIONAL ENERGY AGENCY (2025): "Electricity 2025: Analysis and Forecast to 2027," Tech. rep., IEA.
JAWORSKI, T., C. KITCHENS, AND S. NIGAI (2023): "Highways and globalization," *International Economic Review*, 64, 1615–1648.
JOINT LEGISLATIVE AUDIT AND REVIEW COMMISSION (2024): "Data Centers in Virginia: Commission Draft," Commission report, Virginia General Assembly, Richmond, VA, draft report.
KAY, O., R. REASER, AND R. TAYLOR (2026): "Processing Power: The Effect of Data Centers on Wholesale Electricity Markets," *Working Paper*.
KNIGHT, S., H. X. NGUYEN, N. FALKNER, R. BOWDEN, AND M. ROUGHAN (2011): "The internet topology zoo," *IEEE Journal on Selected Areas in Communications*, 29, 1765–1775.
KNITTEL, C. R., J. R. L. SENGA, AND S. WANG (2025): "Flexible Data Centers and the Grid: Lower Costs, Higher Emissions?" Working Paper 34065, NBER.
MOALLEMI, C. C. AND M. SAĞLAM (2013): "OR Forum—The Cost of Latency in High-Frequency Trading," *Operations Research*, 61, 1070–1086.
MULLER, N. Z. (2025): "Measuring the Impact of Data Centers in the United States Economy: Monetary Damage from Air Pollution and Greenhouse Gas Emissions," Working Paper 35100, NBER.
PEW RESEARCH CENTER (2025): "Americans' Views on Data Centers," .
QUENTEL, M. (2023): "Gone with the wind: renewable energy infrastructure, welfare, and redistribution," in *13th European Meeting of the Urban Economics Association. https://cerge-ei.cz/pdf/events/papers/milan_quentel_job_market_paper. pdf.[Accessed on March 26, 2025]*.
REDDING, S. J. AND E. ROSSI-HANSBERG (2017): "Quantitative spatial economics," *Annual Review of Economics*, 9, 21–58.
SAIZ, A. (2010): "The geographic determinants of housing supply," *The Quarterly Journal of Economics*, 125, 1253–1296.
SANTOS SILVA, J. M. C. AND S. TENREYRO (2006): "The Log of Gravity," *The Review of Economics and Statistics*, 88, 641–658.
SHEHABI, A., A. NEWKIRK, S. J. SMITH, A. HUBBARD, N. LEI, M. A. B. SIDDIK, B. HOLECEK, J. KOOMEY, E. MASANET, AND D. SARTOR (2024): "2024 united states data center energy usage report," .
SIDDIK, M. A. B., A. SHEHABI, AND L. MARSTON (2021): "The environmental footprint of data centers in the United States," *Environmental Research Letters*, 16, 064017.
SLATTERY, C. AND O. ZIDAR (2020): "Evaluating State and Local Business Incentives," *Journal of Economic Perspectives*, 34, 90–118.
S&P GLOBAL (2025): "Global Data Center Market Overview," .
U.S. CENSUS BUREAU (2024): "Annual Business Survey," .

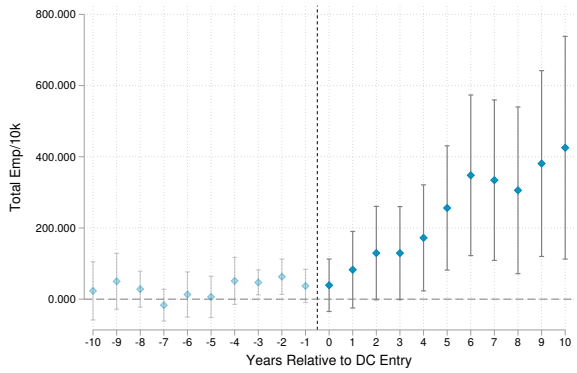
A Appendix

B Event Study Evidence: Local Effects of Data Center Entry

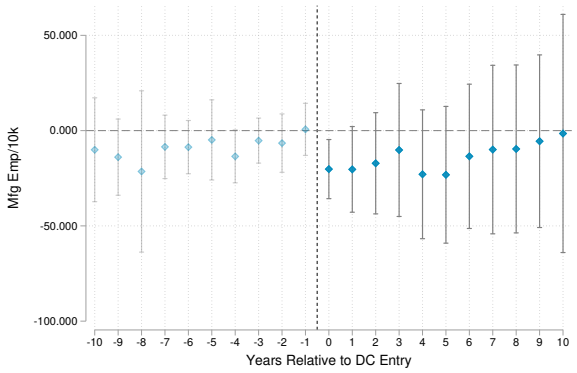
This appendix presents the full set of Callaway-Sant'Anna (CSDID) event study estimates of the local employment effects of first data center entry, testing the robsutness to allowing for a one-year anticipation effect of data center entry, by shifting the treatment year forward by one year to allow for responses that predate the recorded entry year. Next, I present the event-study estimates from Synthetic Difference-in-Differences (SDID).

Figure (A.1) Dynamic Effects of First Data Center Entry on Employment (Levels, per 10,000 Population), One-Year Anticipation Adjustment

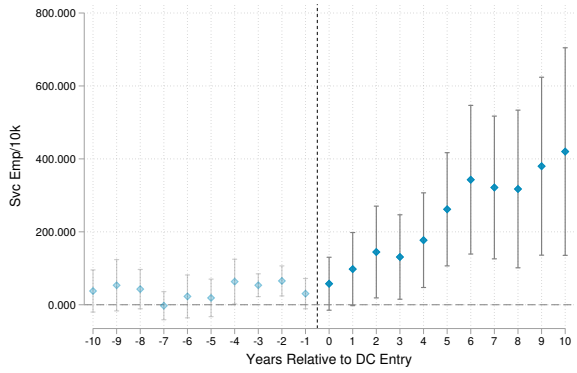
(a) Total



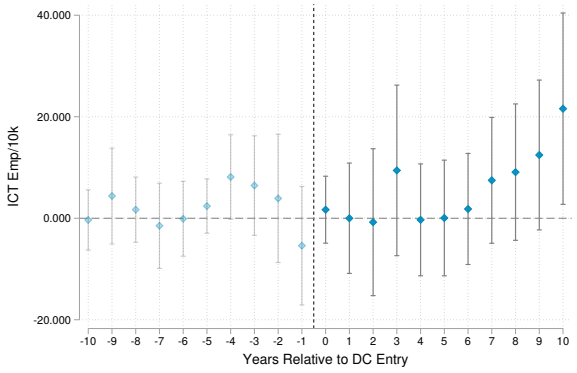
(b) Manufacturing



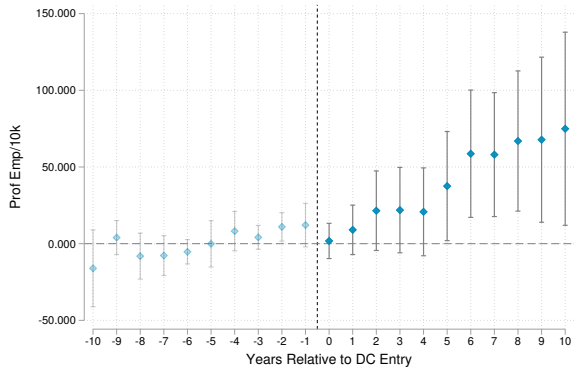
(c) Services



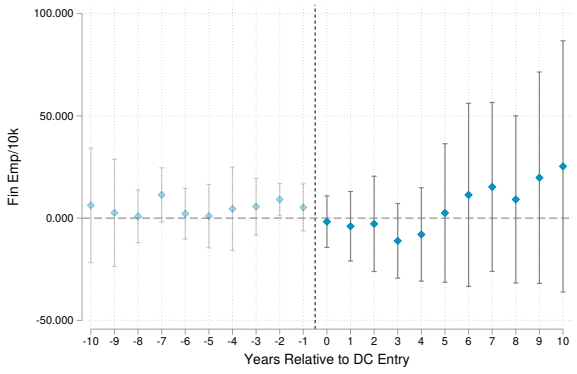
(d) Information



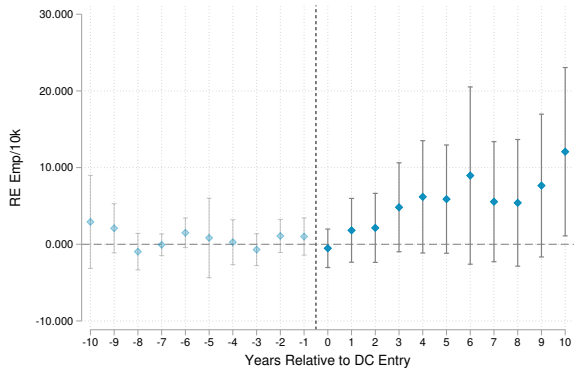
(e) Professional Services



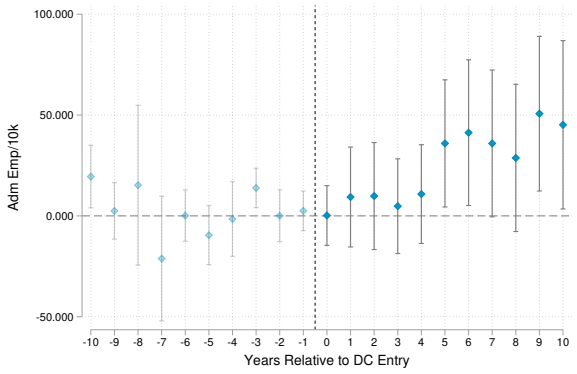
(f) Finance



(g) Real Estate



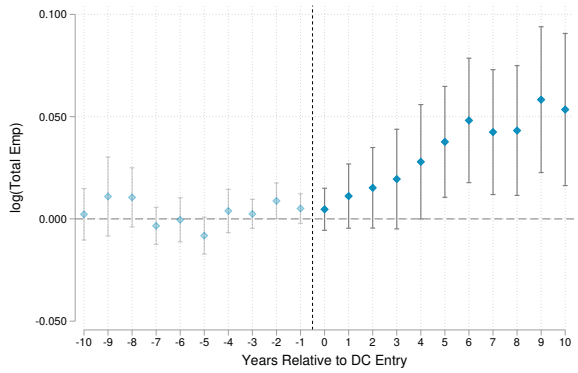
(h) Administrative Services



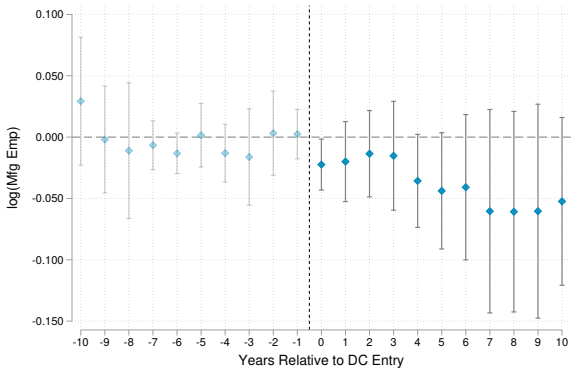
Notes: This figure presents Callaway–Sant'Anna (CSDID) event-study estimates of the effect of first data center entry on employment / 10k, across sectors, under the *One-Year Anticipation Adjustment* sample restriction. Each panel plots dynamic coefficients relative to first entry (year −1 omitted). Shaded/whisker regions show 95% confidence intervals.

Figure (A.2) Dynamic Effects of First Data Center Entry on Employment (Logs), One-Year Anticipation Adjustment

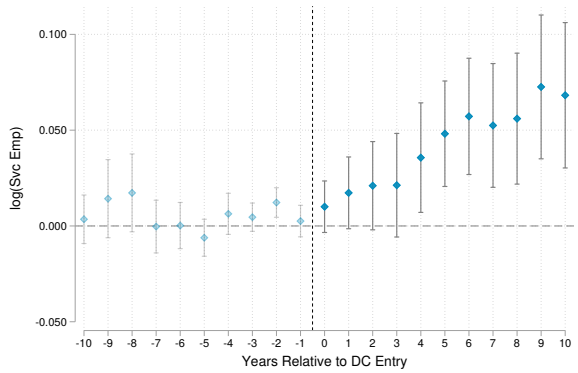
(a) Total



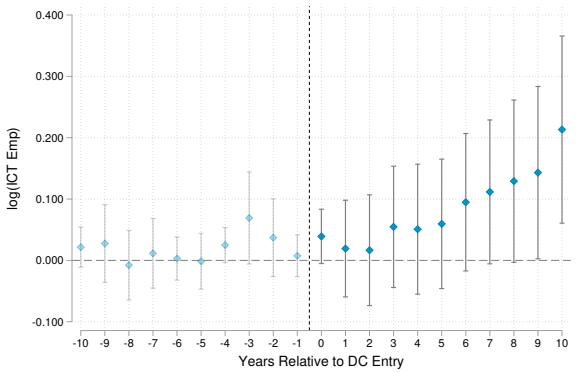
(b) Manufacturing



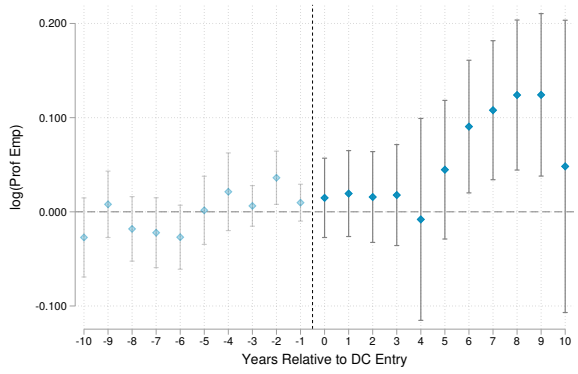
(c) Services



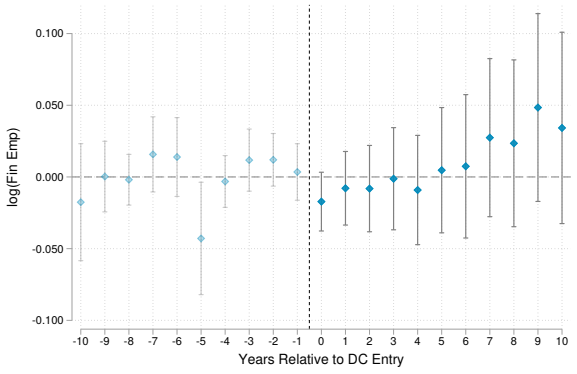
(d) Information



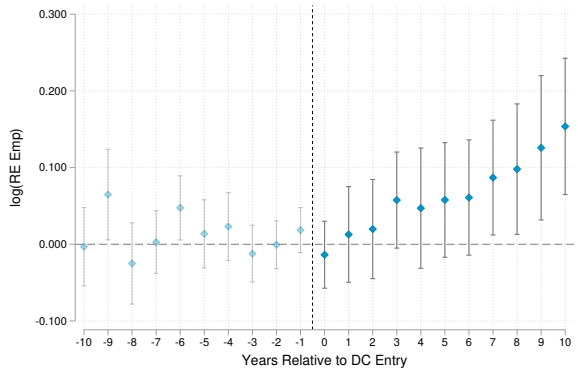
(e) Professional Services



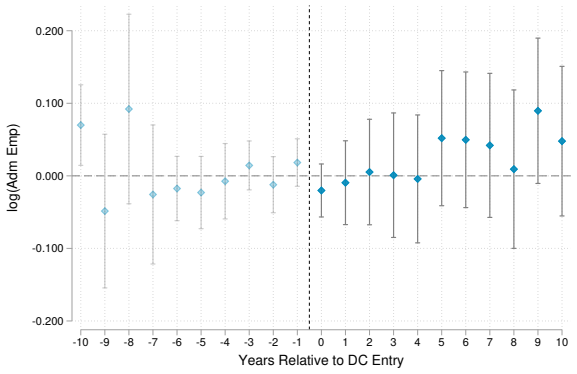
(f) Finance



(g) Real Estate



(h) Administrative Services



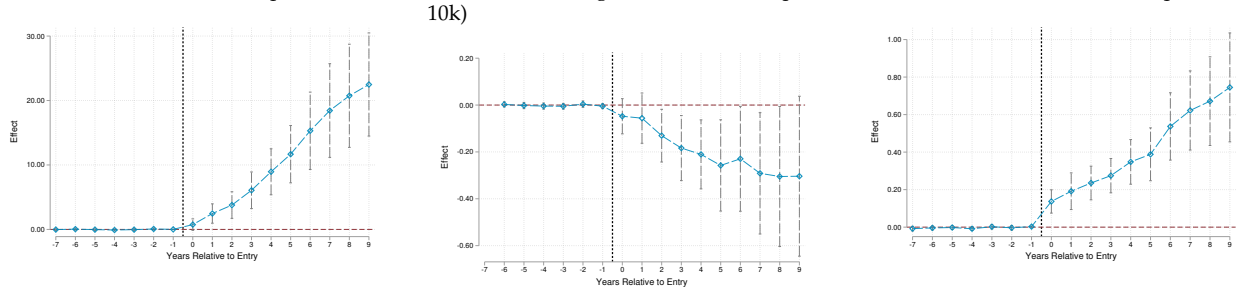
Notes: This figure presents Callaway–Sant'Anna (CSDID) event-study estimates of the effect of first data center entry on log(employment), across sectors, under the *One-Year Anticipation Adjustment* sample restriction. Each panel plots dynamic coefficients relative to first entry (year −1 omitted). Shaded/whisker regions show 95% confidence intervals.

B.1 SDID Event Studies: Establishment Counts and Employment

Figures A.3 and A.4 present Synthetic Difference-in-Differences (SDID) estimates of the effects of first data center entry on establishment counts and employment. The SDID reweights control units via a synthetic control step that balances pre-treatment outcome paths, providing additional robustness against violations of the parallel trends assumption. Qualitative patterns across all outcomes match the CSDID estimates: service-sector establishments and employment rise post-entry while manufacturing shows a gradual decline.

Figure (A.3) SDID: Effect of First DC Entry on Establishment Counts

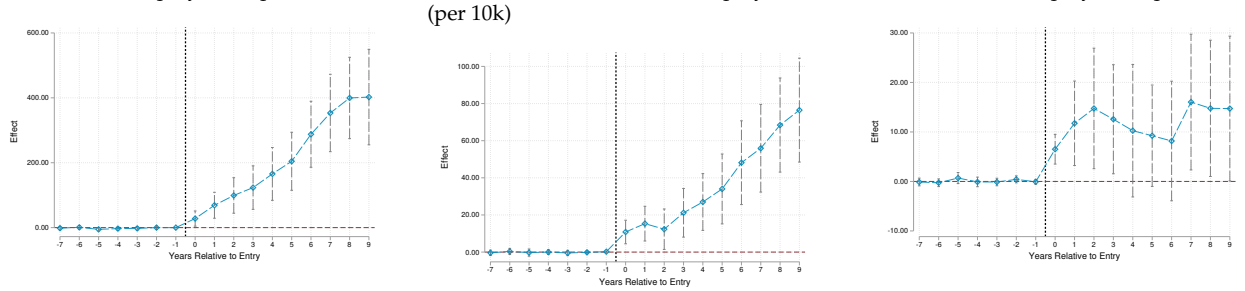
(a) Total Establishments (per 10k) (b) Manufacturing Establishments (per 10k) (c) Information Establishments (per 10k)



Notes: Synthetic Difference-in-Differences (Arkhangelsky et al. 2021) estimates of the effect of first data center entry on establishment counts per 10,000 population, by sector. Control units are reweighted to balance pre-treatment outcome paths. 30 SDID iterations. 95% confidence intervals from bootstrap standard errors (200 replications).

Figure (A.4) SDID: Effect of First DC Entry on Employment

(a) Total Employment (per 10k) (b) Professional Services Employment (per 10k) (c) Information Employment (per 10k)



Notes: Synthetic Difference-in-Differences estimates. Employment per 10,000 population by sector. See Figure A.3 notes for estimation details.

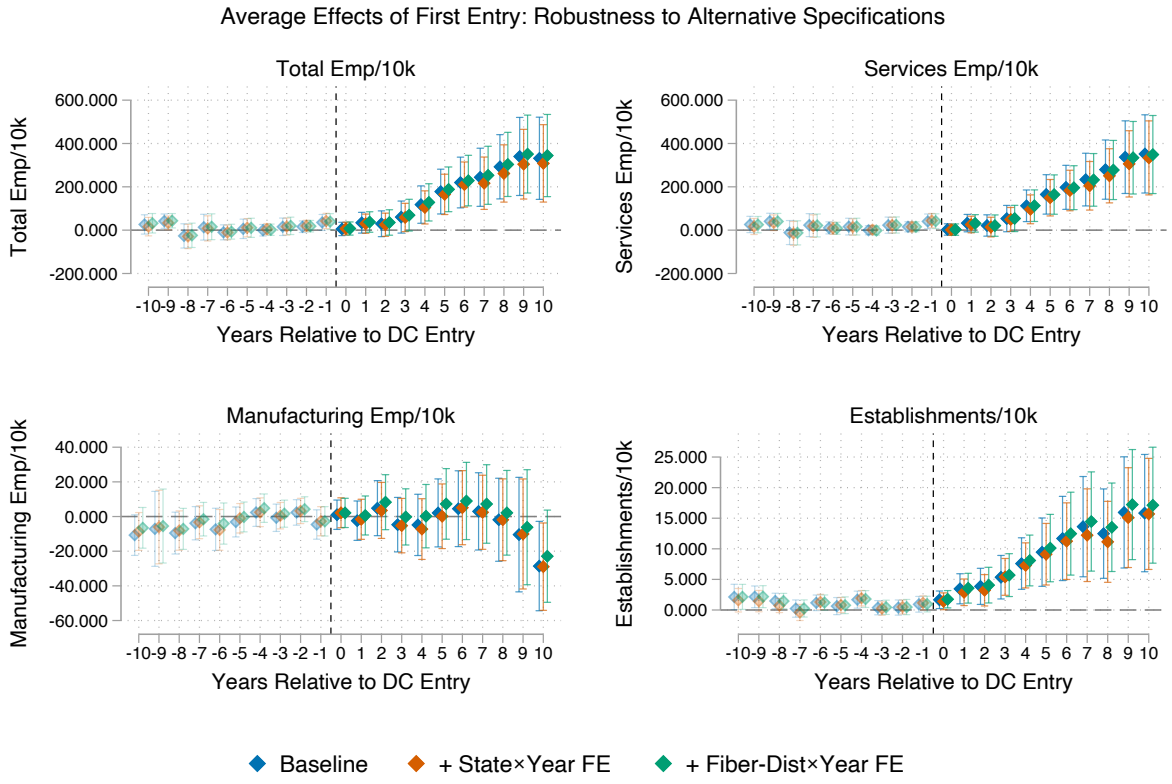
B.2 Robustness to Alternative Trend Controls

Data center entry could coincide with contemporaneous shocks operating at a broader level: a state-wide policy change or industry boom common to every county in a state, or divergent trajectories among places with different *ex ante* digital-infrastructure endowments, since data centers locate close to fiber and connectivity may itself shape local employment. To address these concerns, I re-estimate the pooled first-entry event studies, adding, in turn, (i) state×year fixed effects, which absorb any shock common to a state in a given year, and (ii) fixed effects for the quartile of a county's initial distance to the nearest long-haul fiber node interacted with year, which allow counties with different pre-existing connectivity to follow separate time paths.

Figure A.5 overlays the resulting event studies on the baseline. The dynamics and magnitudes are essentially unchanged across all four outcomes. For total employment the average post-entry effect moves

only from 168.5 additional jobs per 10,000 residents in the baseline to 153.9 with state×year fixed effects and 177.0 with fiber-distance×year fixed effects, and service employment behaves almost identically (162.5, 146.9, and 161.4). Manufacturing employment remains a precise zero under every specification, and the establishment-count response is likewise stable (9.2, 8.6, and 9.8 establishments per 10,000). In every case the pre-entry coefficients stay flat and centred on zero, so the additional controls do not disturb the parallel-trends assumption. The employment and establishment effects of data center entry are thus not an artefact of coincident state-level shocks, nor of differential trends among counties that were better connected to begin with.

Figure (A.5) Robustness to Contemporaneous Shocks: State×Year and Fiber-Distance×Year Fixed Effects



Notes: Each panel overlays Callaway–Sant'Anna event studies of the effect of first data center entry on a per-10,000-resident outcome for the pooled first-entry (*min3*) cohort, with not-yet-treated counties as the comparison group. The baseline (blue) residualises the outcome on 1990 college-share quartile×year and 1990 median household income quartile×year fixed effects; the second specification (vermillion) adds state×year fixed effects; the third (green) instead adds fixed effects for the quartile of a county's initial distance to the nearest fiber node interacted with year. Markers are dynamic coefficients relative to first entry (year -1 omitted) and whiskers are 95% confidence intervals; pre-entry markers are shown faded. The reported "average post-entry effect" averages the event-study coefficients over horizons $h = 0, \ldots, 10$. The vertical scale differs across panels.

B.3 Leave-One-State-Out Estimates

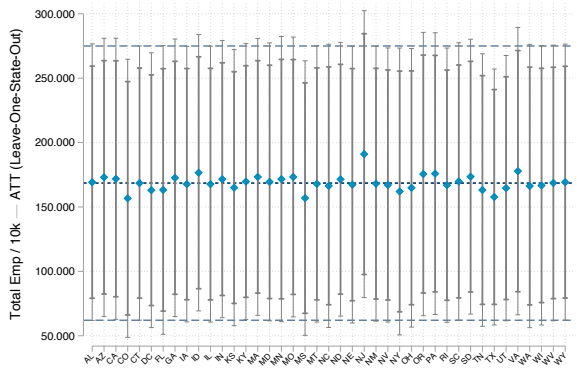
A remaining concern is that the estimated effects are driven by a single state whose data center build-out coincided with other local shocks. Virginia is the natural candidate: Northern Virginia hosts a large share of US data center capacity, and its concurrent expansion of professional and information services could in principle generate the observed employment path without any causal role for data center entry. To assess

this, I re-estimate the pooled first-entry (*min3*) specification 41 times for each outcome, each time dropping all counties in one of the 41 states that contain at least one treated county, and compare each resulting ATT with the full-sample estimate. Figures A.6 and A.7 report the results for employment and establishments respectively.

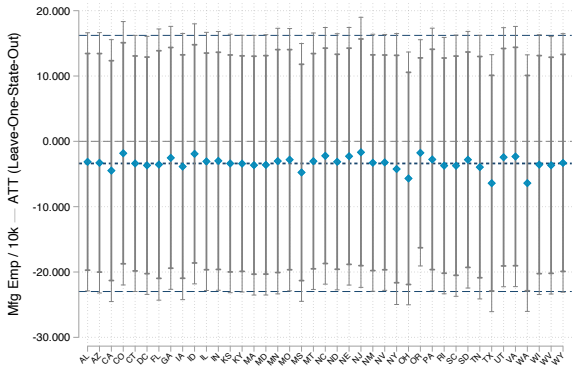
Figure A.6 reports the results for total, manufacturing and service employment. For total employment (Panel a), the full-sample effect of 168.5 additional jobs per 10,000 residents is essentially unchanged by any single exclusion: the estimates range from 156.7 (dropping Colorado) to 191.0 (dropping New Jersey), with a mean of 168.9, and all 41 lie inside the full-sample 95% confidence interval while remaining statistically significant at the 5% level. Service employment (Panel c) behaves almost identically, as expected given that services account for nearly all of the total: the full-sample ATT of 162.5 moves within a range of 151.7 (dropping Texas) to 183.0 (dropping New Jersey), again with no exclusion pushing the estimate outside the full-sample confidence interval or overturning significance. The employment response therefore does not rest on any one state's experience, and in particular is not an artefact of Northern Virginia.

Figure (A.6) Leave-One-State-Out Estimates: Effect of First Data Center Entry on Employment

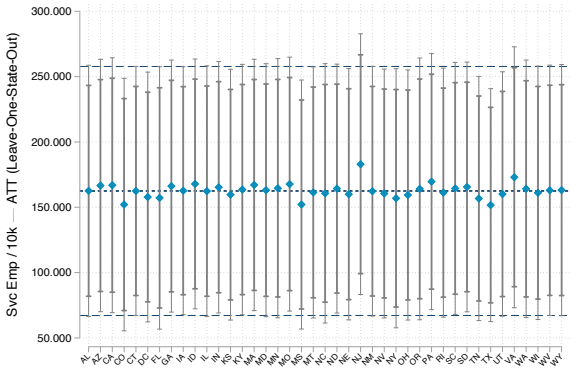
(a) Total Employment



(b) Manufacturing Employment



(c) Service Employment

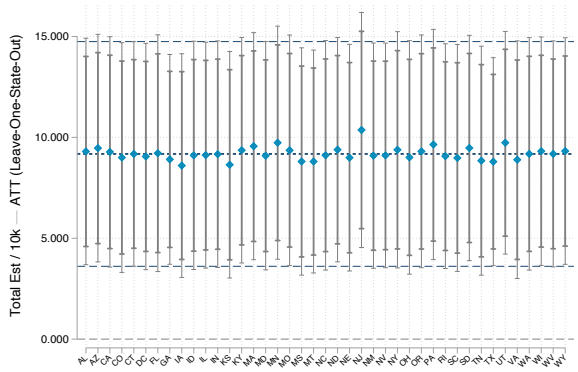


Notes: This figure reports the sensitivity of the baseline employment results to the exclusion of individual states. Each marker is the average post-entry treatment effect (ATT) on employment per 10,000 residents from a Callaway–Sant'Anna event study re-estimated after dropping every county in the state shown on the horizontal axis; the ATT averages the event-study coefficients over horizons $h = 0, \ldots, 10$. Estimation uses the pooled first-entry (*min3*) cohort with not-yet-treated counties as the comparison group, after residualising the outcome on 1990 college-share quartile $\times$ year and 1990 median household income quartile $\times$ year fixed effects. Grey whiskers show 90% (thick) and 95% (thin) confidence intervals. In each panel the solid navy line is the full-sample ATT (168.5 for total, -3.4 for manufacturing, 162.5 for services) and the dashed navy lines its 95% confidence interval; the dashed grey line marks zero. The 41 states shown are those containing at least one treated county. Note that the vertical scale differs across panels.

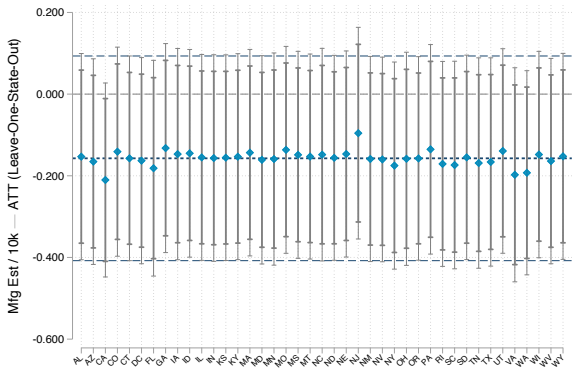
The same exercise for establishment counts is reported in Figure A.7, and the picture is if anything tighter. The full-sample effect on total establishments of 9.2 per 10,000 residents varies only between 8.6 and 10.4 across the 41 exclusions (Panel a), and service establishments move between 7.9 and 9.7 around a full-sample estimate of 8.5 (Panel c); in both cases every leave-one-out estimate remains inside the full-sample 95% confidence interval and statistically significant at the 5% level. Manufacturing establishments (Panel b) mirror the employment result: the full-sample ATT of -0.16 is not distinguishable from zero, and no exclusion moves it outside a narrow band between -0.21 and -0.10 or renders it significant. As with employment, dropping Virginia leaves the establishment estimates essentially unchanged.

Figure (A.7) Leave-One-State-Out Estimates: Effect of First Data Center Entry on Establishments

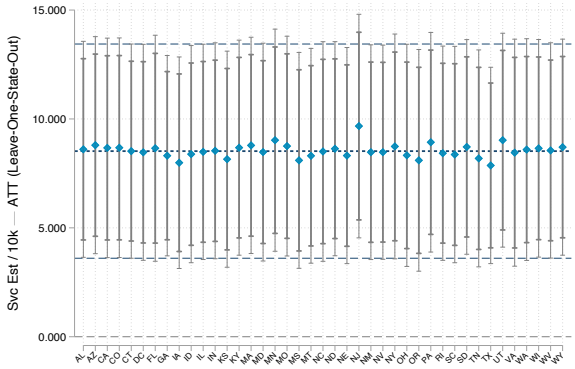
(a) Total Establishments



(b) Manufacturing Establishments



(c) Service Establishments



Notes: This figure repeats the exercise of Figure A.6 for establishment counts per 10,000 residents. Each marker is the average post-entry treatment effect (ATT) from a Callaway–Sant'Anna event study re-estimated after dropping every county in the state shown on the horizontal axis; the ATT averages the event-study coefficients over horizons $h = 0, \ldots, 10$. Estimation, sample and fixed effects are as in Figure A.6. Grey whiskers show 90% (thick) and 95% (thin) confidence intervals. In each panel the solid navy line is the full-sample ATT (9.2 for total, -0.16 for manufacturing, 8.5 for services) and the dashed navy lines its 95% confidence interval; the dashed grey line marks zero. Note that the vertical scale differs across panels.

B.4 More Details on Housing Sector in the Model

This section documents the construction of the housing block's location fundamentals: the supply elasticity η_c and the supply level $\bar{H}_c$;and the resulting equilibrium condition for local housing prices.

Housing Supply Elasticity Imputation. The first step involves imputing county-level housing supply elasticities. I use tract-level supply elasticity estimates from Baum-Snow and Han (2024), aggregated to the county level as housing-unit-weighted means. For around 900 counties covered by BSH, η_c is used directly. For all remaining counties, η_c is imputed from a regression fit on BSH-covered counties:

$$\eta_c = \beta_0 + \beta_1 \cdot FrcDev_c + \beta_2 \cdot FrcDev_c^2 + \beta_3 \cdot FrcFlat_c + \beta_4 \cdot CBDDist_c + \varepsilon_c, \tag{A.1}$$

where $FrcDev_c$ is the fraction of county land developed as of 2001 from NCLD, $FrcFlat_c$ is the fraction of land with slope below 15 degrees from usgs, and $CBDDist_c$ is distance to the nearest urban core. The regression is weighted by $\sqrt{n_{\text{tracts},c}}$. Expected signs are $\beta_1 < 0$ (developed land constrains supply response), $\beta_2 > 0$ (redevelopment at high density), and $\beta_3, \beta_4 > 0$ (flatter, more peripheral land eases construction); fitted coefficients satisfy all four sign restrictions, with in-sample $R^2 = 0.80$. I restrict imputed elasticities at $[0.05, 3.00]$ to prevent pathological extrapolation.

Housing Fundamental Inversion. Given η_c, I invert the housing market-clearing condition to recover the county-level supply fundamental $\bar{H}_c$ exactly at the 2024 baseline:

$$\bar{H}_c = \frac{\alpha_H \cdot w_c L_c}{\left(P_{H,c}^{\text{obs}}\right)^{1+\eta_c}}, \tag{A.2}$$

where $w_c L_c$ is the total baseline wage bill, $\alpha_H = 0.18$ is the Cobb-Douglas housing expenditure share, and $P_{H,c}^{\text{obs}}$ is the observed Zillow home price index, normalised to Cook County, IL (FIPS 17031). Because $\bar{H}_c$ is solved directly from this identity, substituting it back into equation (A.4) below exactly reproduces $P_{H,c}^{\text{obs}}$.

Housing Equilibrium and Local Price Determination. Local housing supply is an increasing, concave function of the quantity of housing consumed in equilibrium, governed by the calibrated fundamentals $(\eta_c, \bar{H}_c)$:

$$HS_c = \bar{H}_c \cdot Q_{H,c}^{\eta_c}. \tag{A.3}$$

On the demand side, each worker spends a fixed Cobb-Douglas share of income on housing. Equating housing demand to supply gives the equilibrium price:

$$P_{H,c} = \alpha_H \left\{ \frac{w_T L_T + w_{S1} L_{S1} + w_{S2} L_{S2}}{HS_c} \right\}^{\frac{1}{1+\eta_c}}. \tag{A.4}$$

In any counterfactual equilibrium, $(\eta_c, \bar{H}_c)$ are held fixed and equation (A.4) is re-solved using counterfactual sector wage bills. County-specific η_c governs how much the potential wage shocks from data center entry pass through to local housing costs versus housing quantities.

B.5 Firm Hosting Location Choice, Latency, and the Elasticity of Substitution

This appendix documents the data and the identification strategy behind η, the elasticity of substitution across compute-source locations that governs the CES price index for compute in Section 6.9. It proceeds in the order the estimate is built: the three underlying datasets (B.5.1), the descriptive distance gravity that motivates the exercise (B.5.2), the structural gravity equation that maps hosting shares into η (B.5.3), the free-flow latency instrument (B.5.4), the main estimates (B.5.5), industry heterogeneity (B.5.6), and robustness to an entirely independent latency measurement from RIPE Atlas (**??**).

B.5.1 Data

Firm hosting locations. The origin side of the panel is the Revelio Labs universe of US firms. For each firm I resolve its website domain to an IPv4 address by DNS and geolocate that address using the MaxMind GeoLite2 City and ASN databases, filling city-level gaps from IP2Location LITE DB11 where MaxMind returns no city. Firm headquarters cities are geocoded once through the ArcGIS geocoder and cached, so origins and destinations are placed on a common footing. Both endpoints are then assigned to 2010 US counties by point-in-polygon.

The critical restriction is on *what a resolved IP address actually locates*. Each observation carries a confidence tier derived from its autonomous-system organisation and HTTP response headers. **Tier 1** covers firms whose IP falls in an identified cloud provider range or which host directly: here the address identifies a genuine hosting site. I exclude the firms which use any-cast Content Distribution Network-platforms — Cloudflare, Fastly, Akamai — where the resolved address is the edge point of presence nearest the resolver, and therefore describes the caching network rather than where the firm buys compute. Table A.1 reports the full sequence of restrictions.

Table (A.1) Construction of the firm hosting-flow sample

Restriction	Observations	% of raw
All firm–IP observations	2,394,143	100.0
Tier 1 hosting (excludes CDN/anycast)	1,320,770	55.2
US, with valid HQ and hosting fields	1,150,578	48.1
Unique firm × location, HQ geocoded	1,151,888	48.1
Cross-county pairs (estimation sample)	1,122,107	46.9
Distinct origin counties	3,052	
Distinct hosting counties	926	
Distinct county pairs	65,601	

One raw observation is a (firm domain, resolved IP) pair. Tier 1 retains firms whose resolved IP identifies an actual hosting site — a cloud provider range or direct hosting — and drops Tier 2, where the IP resolves to an anycast or CDN edge point of presence and therefore locates the caching network rather than the firm's compute. The fourth row rises marginally on the third because a small number of HQ cities match more than one entry in the geocode cache.

Fibre network. Physical network geography comes from Infrapedia's terrestrial cable layer, which after restriction to CONUS leaves 196 cables spanning 1,360,936 km. I build a routable graph by exploding every cable into consecutive vertex pairs, snapping endpoints to a regular 5 km grid, and weighting edges by centre-to-centre distance between snapped nodes — which mechanically guarantees that the routed path is never shorter than the great circle. An all-pairs Dijkstra search yields a routed fibre distance for 8,223,302 county pairs, 85.1% of all ordered pairs. Median route circuity — routed distance divided by great-circle distance — is 1.512.

Latency. Bilateral latency comes from two directly measured sources, compared in Table A.2: WonderNetwork's server-pair ping panel, 2014–2024, and RIPE Atlas, described in **??**.

Table (A.2) Bilateral latency: two directly measured sources

	WonderNetwork	RIPE Atlas
Vantage points	84 servers	152 anchors
Raw observations	44,153	7,493,889
Period	2014–2024	2026
County pairs	5,433	4,860
Origin counties	83	432
Destination counties	67	54
$R^2(\ln \Phi \sim \ln d \mid \text{FE})$	0.84	0.84
Free-flow instrument (Infrapedia)		
County pairs with a route	8,223,302	
Median circuity	1.512	
Median τ^0 (ms)	25.7	

A third, fitted county-level series exists but is unusable for this purpose: constructed as $\exp(\hat{\alpha}_o + \hat{\alpha}_d + \hat{\delta}_t + \hat{\beta} \ln d_{od})$, it is an exact affine function of $\ln d_{od}$ given the fixed effects ($R^2 = 1.0000000000$), so it admits no instrument.

A third, fitted county-level series is available but cannot be used here. It is constructed as $\widehat{\tau}_{od} = \exp(\hat{\alpha}_o + \hat{\alpha}_d + \hat{\delta}_t + \hat{\beta} \ln d_{od})$ and is therefore, conditional on the fixed effects, an *exact* affine function of $\ln d_{od}$ ($R^2 = 1.0000000000$). It carries no information about network quality beyond distance, so no instrument can identify β once distance is controlled — rank deficiency rather than weak instruments. All estimates below use directly measured latency.

B.5.2 Descriptive gravity: hosting demand declines with distance

Before imposing structure it is worth establishing the reduced-form fact. Let i index firms, o the headquarters county, d the hosting county, and N_{od} the number of firms headquartered in o that host in d. The city-pair gravity equation is

$$\ln N_{od} = \alpha_o + \alpha_d + \beta \ln d_{od} + \varepsilon_{od}, \tag{A.5}$$

where d_{od} is great-circle distance and α_o, α_d absorb all origin- and destination-level confounders — city size, income, industrial composition, and the total supply of hosting capacity. Because (A.5) drops zero-flow pairs and zeros are informative here, I also estimate it by Poisson pseudo-maximum likelihood (Santos Silva and Tenreyro, 2006) over the full grid of pairs,

$$\mathbb{E}[N_{od} \mid d_{od}, \alpha_o, \alpha_d] = \exp\big(\alpha_o + \alpha_d + \beta \ln d_{od}\big). \tag{A.6}$$

Both are estimated with `fixest` (Bergé, 2018).

Table (A.3) Distance Elasticity of Hosting Demand: City-Pair Gravity

	(1) OLS	(2) PPML
$\hat{\beta}$ (log distance)	-0.074	-0.114
	(0.002)	(0.007)
R^2 / Pseudo-R^2	0.582	NA
N	194,010	8,934,617
Origin FE	✓	✓
Destination FE	✓	✓
Includes zero flows	No	Yes

Distance is the Haversine distance between the geocoded firm HQ city centroid and the IP-geolocated hosting location. Both columns use city-to-city origin–destination flows with origin and destination fixed effects; column (2) estimates equation (A.6) by PPML over the full grid of pairs, including those with zero flow.

In Table A.3 the OLS elasticity is -0.074 and the PPML elasticity, estimated over the 8.9 million zero-inclusive pairs, is -0.114. Distance decay is real but modest — which is exactly what motivates the structural exercise: distance is a proxy, and what should matter economically is the network latency it stands in for.

B.5.3 From distance to latency: identifying η

In the model of Section 6.9, firms in county c buy compute from data centres in county d through a CES aggregator with elasticity η over source locations. Compute delivered from d to c is discounted by network quality Φ_{dc}. The share of c's compute purchased from d is then log-linear in quality, which gives the estimating equation

$$\ln s_{od} = \beta \ln \Phi_{od} + \alpha_o + \alpha_d + \varepsilon_{od}, \qquad \eta = \beta + 1. \tag{A.7}$$

Here s_{od} is the share of firms headquartered in o that host in d; network quality is the reciprocal of latency, $\Phi_{od} = 1/\tau_{od}$; and the fixed effects absorb the origin's total compute demand and the destination's price and capacity.

Why OLS is biased toward Cobb–Douglas. Two forces attenuate $\hat{\beta}$: latency is measured with error, and realised latency is endogenous to the flows being explained, because carriers provision capacity and optimise routing on exactly the paths that carry the most traffic. Both push $\hat{\eta}$ toward unity, where the CES collapses to Cobb–Douglas and the compute channel goes inert.

B.5.4 The free-flow latency instrument

I instrument measured latency with the latency the physical network would deliver if it were uncongested and optimally routed:

$$\tau^0_{od} \;=\; \frac{2\, D^{\text{fibre}}_{od}}{v}, \qquad v = 150 \text{ km/ms}, \tag{A.8}$$

where D^{fibre}_{od} is the Dijkstra shortest path through the Infrapedia graph and the factor of two converts to round trip. **Equation (A.8) contains no estimated parameters.** It is a physical property of where cable was laid, which is what distinguishes it from the fitted series discussed above.

Regressing minimum observed round-trip time on routed fibre distance across 5,203 directly measured server pairs implies a propagation speed of 174.2 km/ms ($R^2 = 0.843$). Since light in fibre travels at roughly 200 km/ms, the assumed 150 km/ms is conservative and τ^0 is an upper bound on true free-flow latency.

Exclusion. The identifying assumption is that the physical cable network was laid for reasons unrelated to today's firm-level hosting choices. The vintage of the network supports this directly: 74.5% of terrestrial cables, and 75.8% of terrestrial route length, were already in service by 1999 — roughly fifteen years before the data-centre build-out this paper studies. The backbone is essentially predetermined. This is also why I build a single instrument rather than a pre-boom subgraph: a network restricted to pre-2000 cable would be numerically almost identical to the full one.

Figure A.8 plots the residualised first-stage plot, relating measured compute quality against the predicted quality from circuity-driven variation due to the structure of terrestrial fiber networks.

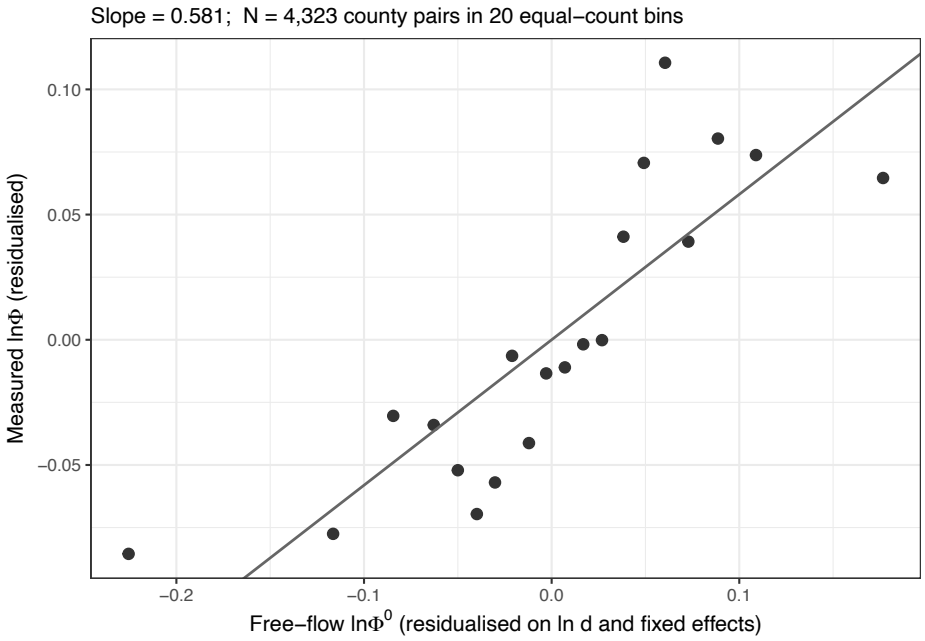


Figure (A.8) First stage, identifying variation only. Both axes are residualised on $\ln d_{od}$ and on origin and destination fixed effects, so the figure shows the circuity-driven variation that identifies column (4) of Table A.4. Twenty equal-count bins.

B.5.5 Main estimates

Table A.4 reports a four-column ladder for each latency source and each level of geographic aggregation.

Table (A.4) Elasticity of substitution across compute-source locations

	(1) OLS	(2) OLS	(3) 2SLS	(4) **2SLS**
Controls for $\ln d_{od}$	No	Yes	No	**Yes**
Panel A: County pairs				
WonderNetwork				
$\hat{\beta}$	0.313***	0.072	0.370***	0.549*
	(0.022)	(0.063)	(0.027)	(0.297)
$\hat{\eta} = \hat{\beta} + 1$	1.313	1.072	1.370	1.549
AR 95% CI for η	—	—	—	[1.00, 2.15]
First-stage F	—	—	41.4	41.4
OD pairs			2,659	
Origins / destinations			82 / 62	
RIPE Atlas				
$\hat{\beta}$	0.333***	-0.004	0.465***	0.392*
	(0.024)	(0.043)	(0.032)	(0.222)
$\hat{\eta} = \hat{\beta} + 1$	1.333	0.996	1.465	1.392
AR 95% CI for η	—	—	—	[1.00, 1.85]
First-stage F	—	—	50.4	50.4
OD pairs			2,354	
Origins / destinations			391 / 49	
Pooled				
$\hat{\beta}$	0.316***	-0.001	0.427***	0.399**
	(0.019)	(0.036)	(0.024)	(0.200)
$\hat{\eta} = \hat{\beta} + 1$	1.316	0.999	1.427	1.399
AR 95% CI for η	—	—	—	[1.05, 1.80]
First-stage F	—	—	78.0	78.0
OD pairs			4,421	
Origins / destinations			406 / 88	

The dependent variable is $\ln s_{od}$, the log share of firms headquartered in o that host compute in d. All columns include origin and destination fixed effects; standard errors are clustered on origin and reported in parentheses.
Compute quality is $\Phi_{od} = 1/\tau_{od}$, so $\ln \Phi = -\ln \tau$ and the CES elasticity is $\eta = \beta + 1$.
Latency is the mean across measured pairs of the *minimum* round-trip time, the least-congested observation. *Pooled* averages the two sources in latency space over the union of their county pairs.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Column (1) is the OLS analogue of the specification underlying the $\eta = 1.295$ baseline. Column (2) adds $\ln d_{of}$. Columns (3) and (4) instrument. Column (4) uses the free-flow latency to instrument for observed latency and finds very similar estimates. The central estimate of the elasticity of substitution is $\hat{\eta} = 1.399$ with a standard error of 0.200 on the pooled latency panel, with a first-stage F of 78.0. The two independent measurement sources align closely: 1.549 on WonderNetwork and 1.392 on RIPE Atlas

B.5.6 Heterogeneity by industry

The model in Section 6.9 treats compute as a single good, but latency sensitivity plausibly differs across industries: a firm serving interactive applications should weight network quality more heavily than one

running batch analytics. Figure A.9 re-estimates the main specification separately by 3-digit NAICS and orders industries by the point estimate.

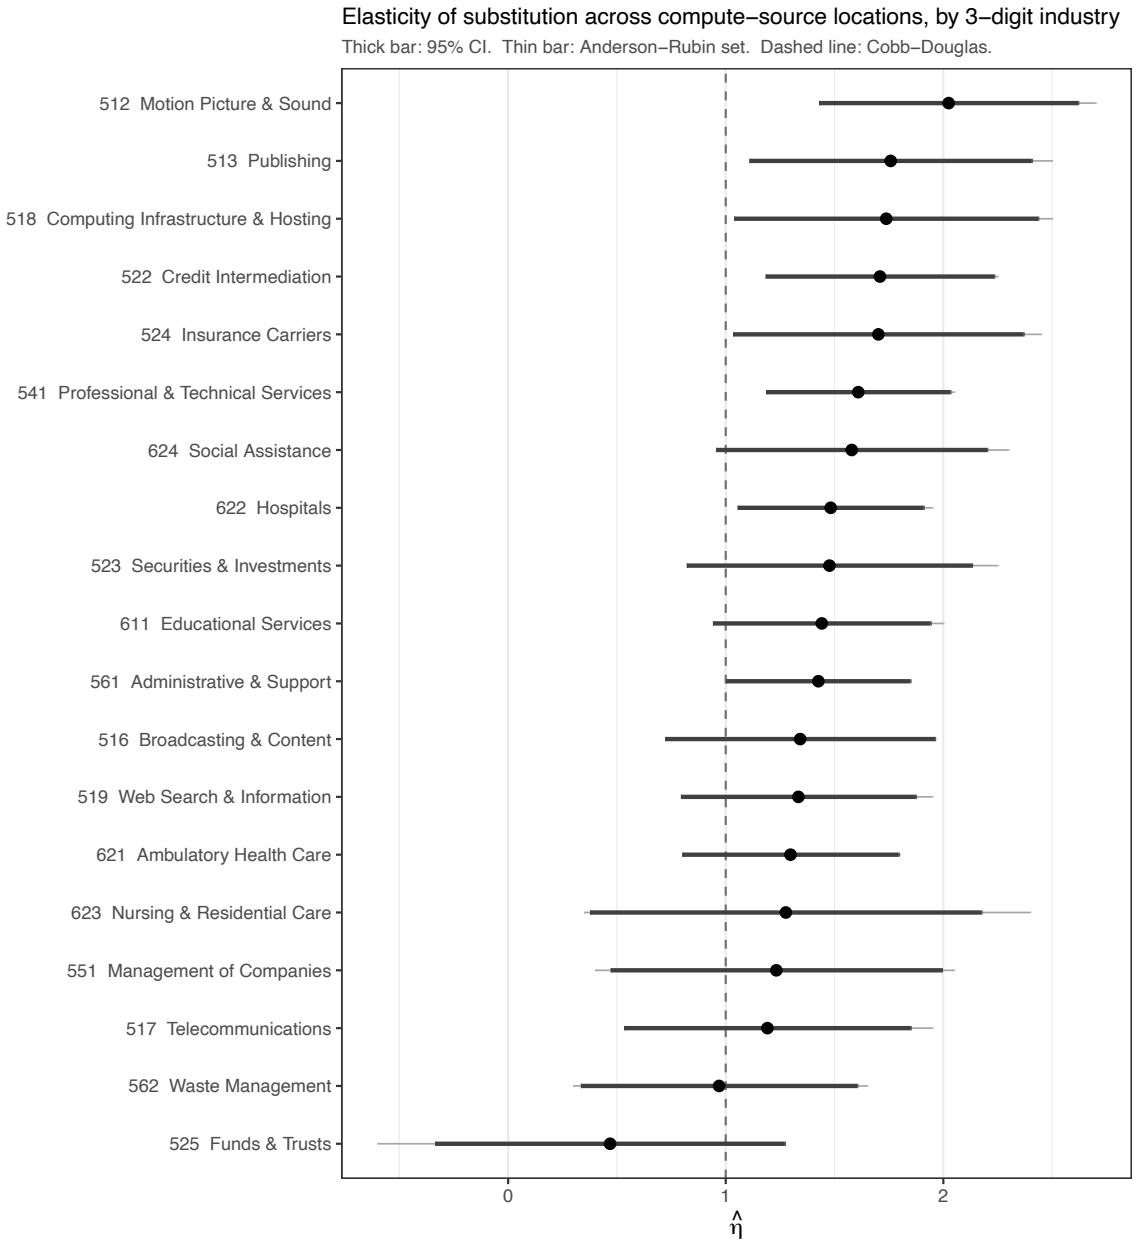


Figure (A.9) Elasticity of substitution by 3-digit industry. Each row re-estimates column (4) of Table A.4 — 2SLS instrumenting $\ln \Phi_{od}$ with free-flow latency and conditioning on $\ln d_{od}$ — on that industry's firms, at county-pair level on the pooled latency panel. Industries with at least 30 firms and 10 OD pairs. Thick bar: 95% confidence interval, clustered on origin. Thin bar: Anderson–Rubin set. Dashed line marks the Cobb–Douglas case $\eta = 1$.

The ordering is broadly consistent with latency sensitivity. Motion picture and sound recording (NAICS 512, $\hat{\eta} = 2.03$), publishing (513) and computing infrastructure and hosting (518, 1.74) sit at the top, alongside credit intermediation (522) and insurance carriers (524), where execution latency is directly monetised; telecommunications carriers (517, 1.19) and waste management (562) sit at the bottom.

C The Power-Plant Dispatch Model

One key innovation of the framework is to embed an integrated electricity generation and pollution dispersal model into the existing Quantitative Spatial Model.

Similar to Arkolakis and Walsh (2023), electricity in each location c is generated from a set of power plants, with fixed locations and capacity, through an hourly merit-order dispatch model of thermal generating plants.

The dispatch model is solved on 45 balancing authorities and produce empirically grounded Balancing-Authority level electricity supply curves, which governs how fast predicted wholesale electricity tariffs change following more generation.

C.1 Marginal cost

The first-step of the model constructs plant-hour marginal costs from fuel costs, heat rates and emission-permit prices.

Each generating unit i in hour t is assigned a short-run marginal cost:

$$mc_{i,t} = \underbrace{vom_i}_{\text{variable O\&M}} + \underbrace{hr_i}_{\text{heat rate}} \times \underbrace{fp_{f(i),s(i),t}}_{\text{delivered fuel price}} + \underbrace{ec_{i,t}}_{\text{emission permits}} \tag{A.9}$$

Heat rates are measured per unit-year from EPA continuous-emissions-monitoring (CEMS) data as reported heat input divided by gross generation. Fuel prices vary by fuel, state and month. Permit costs price SO_2 and NO_x allowances at market rates. Marginal costs are winsorised at the 1st and 99th percentiles *within fuel*, so a single implausible heat rate cannot move a fuel's position in the stack.

Table (A.5) Marginal cost by fuel and technology, 2024

Fuel	Technology	Units	Capacity (GW)	Marginal cost ($/MWh) p10	median	mean	p90
Natural gas	Combined cycle	1,073	291.7	16.20	24.32	26.18	37.44
Coal	Boiler	343	165.2	20.57	29.96	30.98	44.19
Natural gas	Combustion turbine	1,272	110.3	21.50	34.14	37.38	56.37
Natural gas	Boiler	204	53.4	13.20	24.38	29.94	50.50
Petroleum	Boiler	9	4.7	251.88	264.75	306.98	441.30
Petroleum	Combustion turbine	48	3.7	232.71	263.56	355.95	713.03
Coal	Combined cycle	2	0.8	23.31	25.33	25.33	27.36
Petroleum coke	Boiler	2	0.4	32.54	32.66	32.66	32.79
Petroleum	Combined cycle	2	0.2	202.11	204.34	204.34	206.58

Table A.5 is why the supply curve has the shape it does. Gas combined cycle, 291.7 GW and the largest single block of capacity, sits at a median of $24.32 per MWh; coal boilers at $29.96; gas combustion turbines at $34.14; and oil-fired units an order of magnitude above everything else, with a median near $264 and a p90 above $713.

C.2 The Set of Power Plants in Optimisation Problem

Let $\mathcal{I}$ denote the *dispatchable fleet*: every generating unit that reports to CEMS in the year in question, burns a fossil fuel, and can be assigned to a balancing authority. We focus on power plants that satisfy the following properties:

1. **Fossil-fired and CEMS-reporting.** CEMS covers units serving a generator larger than 25 MW that burn fossil fuel. This results in the exclusion of 1,282 small-scale power plants that do not report their emissions. We take their hourly annual generation and construct their emissions using fuel-source average emissions factor. These plants account for 2.8% of metered demand in 2024.
2. **Nuclear, hydro and renewables are excluded by technology.** The model dispatches the fossil fleet and treats everything else as given: nuclear and hydro are scheduled on considerations the model does not represent, and wind and solar are not dispatchable at all. Biomass and waste units do report to CEMS but are excluded by the same fuel screen. To account for them we net their generation out of the balancing-authority load before solving, so the merit order responds only to the residual thermal target. These plants account for 1,497.1 TWh of 2024 generation, 36.7% of metered US demand.

This results in 3,222 units at 1,129 plants, generating 2,430.9 TWh in 2024, across 45 balancing authorities (the dispatch carries hourly constraints for 50 region-series, of which 45 hold a priced fossil fleet). The aggregate dispatchable thermal fleet capacity is around 734.8 GW. The remaining plants are in the overall power plant sample but their generation decisions are not directly modeled in the dispatch framework, and we set their observed generation at eGRID's measured annual values. We cap each power plant's maximum generation at the larger of its nameplate capacity and the 99th percentile of its observed hourly generation. Nameplate exceeds observed p99 for 2,773 of the 3,222 units, and for those the cap is raised — the gap is widest for peaking plants, whose observed p99 understates what they can produce because they run so rarely. The reason for taking a maximum rather than simply switching to nameplate is the other tail: 131 plants record a nameplate *below* output they have actually been observed to produce, which is a data conflict rather than evidence of a lower limit, and the maximum declines to discard the observed capability.

The dispatch solves 2,430.9 TWh, 59.6% of metered US demand of 4,079.0 TWh, and about 95% of *operating fossil* generation. Table A.6 classifies every plant in the frame by the reason it is or is not solved hourly. The reasons are mutually exclusive and applied in the order shown, so a retired nuclear plant is recorded as having no 2024 output rather than as non-dispatchable. The 40.4% of demand the merit order does not solve is exactly three things: 36.7% is non-dispatchable technology, netted out of the target by construction; 2.8% is dispatchable fossil plant with no hourly monitoring; and 1.3% is plant that generated nothing in 2024. Only the middle group is a limitation of the data rather than a modelling choice.

Two notes so the columns reconcile. The frame attributes 2,290.7 TWh to dispatched plants against the 2,430.9 TWh the dispatch solves; the 195.1 TWh difference is solved output on units the plant frame cannot match to a plant identifier, and is the same generation that cannot be placed in a county for the damage channel. And the third group carries positive generation in the table's first column only because the frame's output figure is the eGRID 2023 vintage while the dispatch year is 2024, so a plant that ran in 2023 and stopped appears with output and no base-year monitoring record; its 2024 contribution is zero.

C.3 The hourly dispatch problem

In each hour t, a generator unit $i \in \mathcal{I}$ has a marginal cost $mc_{i,t}$ and an available capacity $\bar{C}_{i,t}$, its capacity net of an outage derating. Let $r(i)$ be the region containing unit i and $\mathcal{I}_r = \{i : r(i) = r\}$. We solve for two programs every hour, differing in how many energy-balance constraints they impose.

The constrained problem imposes one balance per region, equal to that region's observed thermal generation $G_{r,t}$:

$$\min_{\{g_{i,t}\}} \sum_i mc_{i,t}\, g_{i,t} \qquad \text{s.t.} \qquad \sum_{i \in \mathcal{I}_r} g_{i,t} = G_{r,t} \;\; \forall r, \qquad 0 \le g_{i,t} \le \bar{C}_{i,t} \;\; \forall i \tag{A.10}$$

Table (A.6) Every plant in the frame, by the reason it is or is not dispatched, 2024

Category	Plants	Capacity	Generation (TWh)	
		(GW)	eGRID 2023	In 2024
Dispatched hourly				
Natural gas	886	496.3	1,590.0	1,590.0
Coal	177	207.1	676.0	676.0
Nuclear	1	2.9	20.1	20.1
Petroleum	23	9.5	4.4	4.4
Biomass	4	0.5	0.2	0.2
Subtotal	**1,091**	**716.3**	**2,290.7**	**2,290.7**
Not a dispatchable technology				
Nuclear	56	99.0	743.0	743.0
Wind	1,227	120.5	341.4	341.4
Hydro	1,456	102.0	244.7	244.7
Solar	3,924	48.2	95.9	95.9
Biomass	600	10.4	40.1	40.1
Geothermal	64	3.8	15.9	15.9
Waste	64	2.5	12.4	12.4
Other	157	2.5	3.7	3.7
Subtotal	**7,548**	**389.0**	**1,497.1**	**1,497.1**
No hourly generation data				
Natural gas	728	30.2	100.6	100.6
Petroleum	526	8.6	9.3	9.3
Coal	28	1.1	4.1	4.1
Subtotal	**1,282**	**39.9**	**114.1**	**114.1**
No generation in the dispatch year				
Nuclear	1	2.4	18.4	0.0
Coal	73	40.8	15.7	0.0
Natural gas	265	33.2	10.1	0.0
Biomass	26	2.0	6.6	0.0
Petroleum	298	8.0	1.2	0.0
Solar	3	0.4	0.6	0.0
Waste	1	0.1	0.2	0.0
Hydro	2	0.1	0.1	0.0
Subtotal	**669**	**87.0**	**52.8**	**0.0**

The unconstrained problem imposes a single system-wide balance equal to the sum of the regional ones, so total thermal output is identical and only its *allocation* across regions can differ:

$$\min_{\{g_{i,t}\}} \sum_i mc_{i,t}\, g_{i,t} \qquad \text{s.t.} \qquad \sum_i g_{i,t} = \sum_r G_{r,t}, \qquad 0 \le g_{i,t} \le \bar{C}_{i,t} \;\; \forall i \tag{A.11}$$

The unconstrained problem is a relaxation of the constrained one, so its cost is weakly lower and the difference

$$W_t \;=\; \sum_i mc_{i,t}\, g_{i,t}^{\text{con}} \;-\; \sum_i mc_{i,t}\, g_{i,t}^{\text{unc}} \;\ge\; 0 \tag{A.12}$$

is the allocative-inefficiency wedge: the extra cost of meeting each region's load largely from inside that region rather than from the cheapest units nationally.

C.3.1 From the dispatch to a price the spatial model can use

From the hourly dispatched plants, we construct the predicted hourly wholesale price as the marginal cost of the most expensive dispatched unit, $p_{b,t} = \max_i\{mc_{i,t} : g_{i,t} > 0\}$. Next, we construct the annual averaged predicted wholesale price, weighting each hourly observation by the observed Balancing-Authority level load. Finally, to allow the model to match the observed *retail* electricity tariffs at the BA level, we add a BA-level wedge μ_b fixed in 2024, so that $\bar{p}_b + \mu_b$ reproduces the observed industrial tariff, held constant thereafter. All counterfactual price *changes* therefore come from the dispatch, while the base-year level matches the data by construction.

C.4 Merit-order supply curves

Figure A.10 draws the solved supply curve for each major ISO: one rectangle per unit, width equal to its available capacity, height equal to its marginal cost, shaded by fuel. The area under the staircase is total variable cost, and the height at which cumulative capacity crosses the thermal target is the wholesale price — so these are the model, not an illustration of it. There is tremendous spatial heterogeneity in the steepness of the supply curves, leading to large differences in the potential impacts of large electricity demand shocks (such as data center eentry) on equilibrium electricity tariffs. In PJM the median thermal target of 55.0 GW sits well inside cheap coal and combined-cycle gas, and the curve is nearly flat there; the peak target of 111.2 GW pushes into single-cycle gas and oil, past the fleet's median marginal cost of $34.02 and into a tail that reaches $465 per MWh across a few gigawatts. The same increment of new load barely moves price in the first case and moves it a great deal in the second, which is why an elasticity estimated at average load cannot be used to price a data-centre buildout.

C.5 Emissions of the modelled sample

Table A.7 reports what the fleet emits and what it is worth. Total monetised damage is $277.0 billion a year, of which $79.9 billion is local air pollution and $197.1 billion is carbon at a social cost of $145 per tonne.

The last column is the share of each pollutant's tonnage that the merit order controls, and it varies far more than the generation share does: 89.6% of SO_2 and 91.0% of NH_3 against 66.2% of NO_x. A demand shock therefore moves the pollutants very unevenly, and a fleet-average emission factor applied to a demand change would misstate the composition badly. This is why the damage channel is built on dispatched unit generation rather than an average rate — and why ammonia matters more than its 2.7% share of the damage *level* suggests, since almost all of it sits on units the merit order actually moves. (Ammonia is 1.7% of the damage total on the per-capita basis of Table A.7, or 2.7% once damages are population-weighted.)

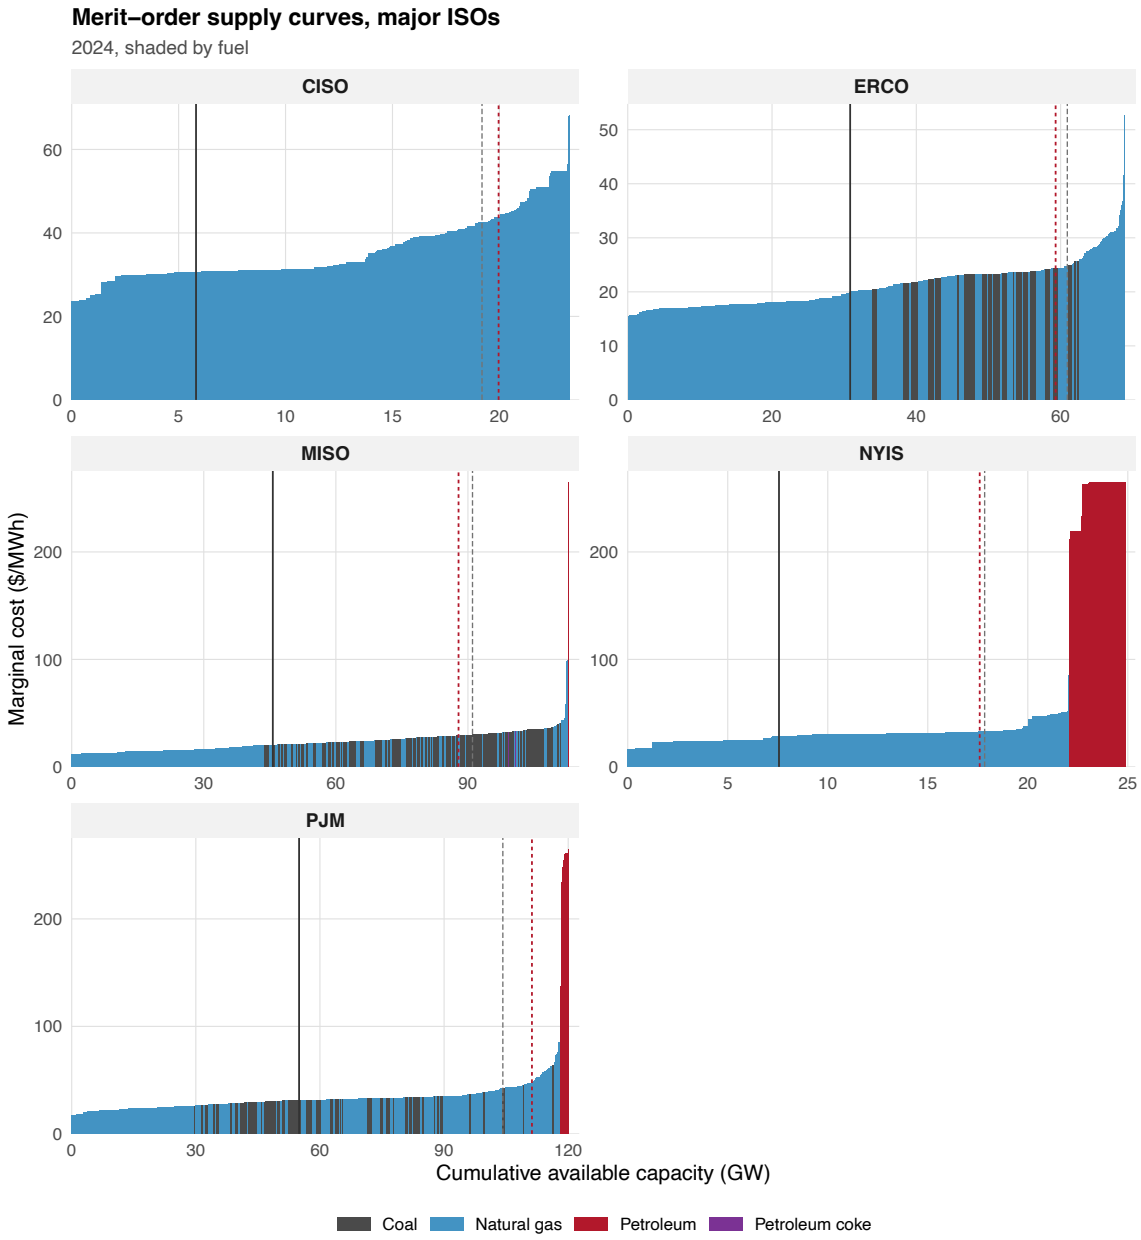


Figure (A.10) Merit-order supply curves, major ISOs, shaded by fuel. Solid grey: median thermal target. Dashed red: peak thermal target. Dashed grey: end of the incumbent p99 fleet.

Table (A.7) Emissions and monetised damage over the modelled sample

Pollutant	Emissions (kt)	Damage ($bn/yr)	% of local damage	Dispatch-controlled
Local air pollutants, priced through the AP3 source–receptor matrices				
SO_2	526.9	36.86	46.16%	89.6%
NO_x	827.5	23.44	29.36%	66.2%
$PM_{2.5}$	75.2	17.19	21.53%	88.5%
NH_3	17.2	2.14	2.67%	91.0%
VOC	22.5	0.22	0.28%	78.5%
Total local		**79.85**	**100.00%**	
Carbon, priced at the social cost of carbon				
CO_2	1,359,650	197.1	—	90.9%
Total monetised damage		**277.0**		

C.6 Which plants respond to load

Figure A.11 locates the plants that respond as load rises. For each of the five major ISOs the dispatch is solved at the hours sitting at the 25th, 50th, 75th and 90th percentile of the thermal target, and each plant is coloured by the lowest percentile at which it is dispatched. A plant shaded for the 25th percentile is effectively baseload; one appearing only at the 90th is part of the marginal fleet whose emissions a demand shock actually moves.

The asymmetry is the point. In PJM, 50.3 GW enters at or below the 25th percentile while 11.1 GW appears only at the 90th, and the median marginal cost of that last tier is $28.39 per MWh against $21.50 for the baseload tier, and it contains 53 gas units and no coal. The plants a new load brings on are therefore both markedly more expensive and, being entirely single-cycle gas, differently sited from the fleet serving average load — so the geography of the damage a data centre causes is not the geography of the existing fleet.

C.7 Validation

This is the central test, and the model never sees a price. Modelled hourly prices are confronted with observed day-ahead wholesale prices for five ISOs, each reduced to a comparable balancing-authority series using that market's own convention: PJM's RTO-wide total, already load-weighted; the mean over NYISO's eleven load zones; the mean over CAISO's NP-15, SP-15 and ZP-26; the mean over MISO's hub and load-zone nodes; and the mean over ERCOT's eight load zones. An ISO-year is usable only if the observed series covers at least 80% of the calendar, which 21 of 28 ISO-years do.

The headline fit is 2024, and only 2024. Every other quantity in this appendix — the fleet, the fuel and permit prices, the heat rates, the capacity ceiling, the merit orders, the emission rates — is built for 2024, so a statistic pooled over 2019–2024 is not a test of the model the paper actually uses. Pooling also imports years the model was never meant to price: 2022, when fuel costs moved faster within the month than the monthly delivered-price series the model reads, and 2021, whose ERCOT February episode is administrative scarcity pricing rather than marginal cost. All 21 usable ISO-years are still reported, in Table **??**, so the restriction is visible and can be audited rather than taken on trust. All five ISOs clear the coverage gate in 2024.

ERCOT is newly usable. The previously assembled panel held 50 delivery dates of bus-level locational prices, far too few to validate on, and ERCOT was excluded from every published fit statistic. Re-parsing ERCOT's day-ahead settlement-point archives gives 5 complete delivery years at 99.7% to 100.0% of calendar days, so ERCOT enters as a fifth validated market rather than a caveat.

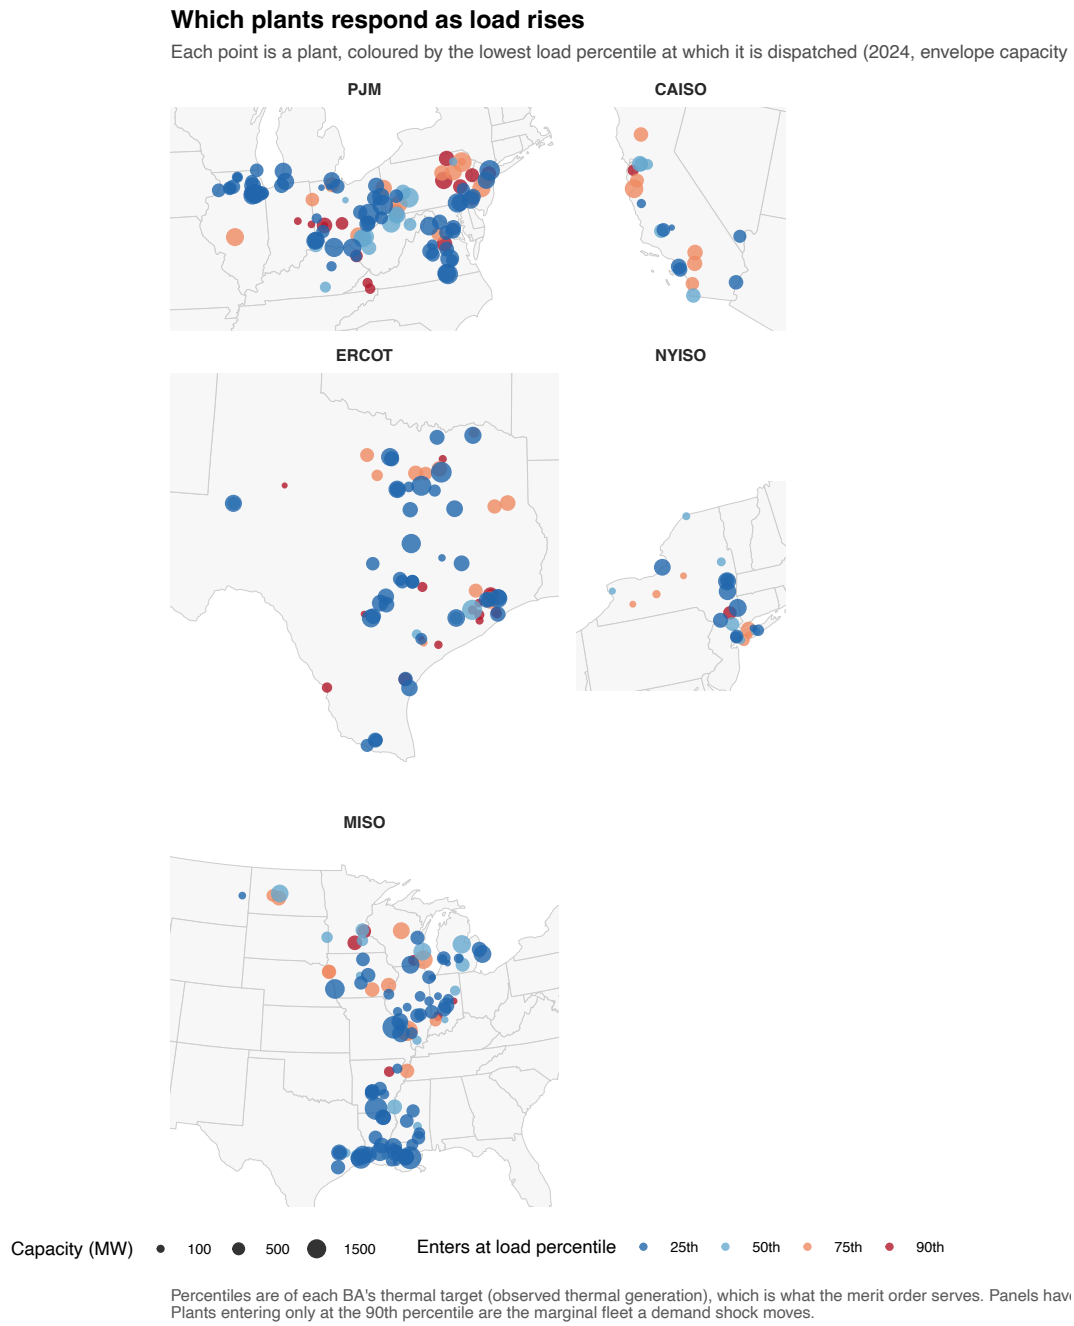


Figure (A.11) Plants responding to load, by the lowest load percentile at which they are dispatched. Panels have independent extents.

Table (A.8) Units and capacity by the load percentile at which they enter

Entry load percentile	Units	Capacity (GW)	Median MC ($/MWh)	Share of ISO capacity
PJM				
25th percentile	178	50.3	21.50	57.1%
50th percentile	32	14.2	29.44	16.2%
75th percentile	43	12.4	27.68	14.1%
90th percentile	53	11.1	28.39	12.6%
CAISO				
25th percentile	12	3.6	25.33	29.9%
50th percentile	8	2.6	30.42	21.4%
75th percentile	12	3.7	65.31	30.6%
90th percentile	7	2.2	108.90	18.1%
ERCOT				
25th percentile	85	27.6	15.79	58.3%
50th percentile	19	4.3	17.72	9.1%
75th percentile	28	7.7	17.51	16.4%
90th percentile	66	7.6	21.25	16.1%
NYISO				
25th percentile	19	6.7	18.23	62.3%
50th percentile	11	2.1	20.73	19.1%
75th percentile	8	1.4	20.00	13.4%
90th percentile	1	0.6	18.49	5.3%
MISO				
25th percentile	148	41.1	17.57	59.1%
50th percentile	55	10.1	21.09	14.5%
75th percentile	25	15.2	22.52	21.8%
90th percentile	19	3.2	24.24	4.5%

C.7.1 Against the published benchmark

The only published figures a model of this class can be held against are Hausman's: an hourly correlation of 0.21 raw and 0.55 winsorised at the 99th percentile, with a modelled mean within 2% of observed. Table A.9 reports ours beside them, for 2024.

Table (A.9) Price goodness of fit by ISO, 2024, against the published benchmark

ISO	Hours	Mean $/MWh		Correlation			MAPE	Level
		Obs.	Mod.	hourly	wins.	monthly	(%)	err. (%)
CAISO	33,312	34.74	30.35	0.558	0.600	0.881	31.8	−12.6
MISO	35,116	27.96	25.40	0.506	0.544	0.977	7.9	−9.2
NYISO	35,136	36.67	21.78	0.390	0.425	0.876	39.0	−40.6
PJM	41,165	31.82	30.46	0.339	0.355	0.873	9.7	−4.3
ERCOT	35,032	30.14	20.39	0.186	0.286	0.656	29.4	−32.4
Pooled (mean of ISOs)	**179,761**	**32.27**	**25.68**	**0.396**	**0.442**	**0.853**	**23.6**	**−19.8**
Pooled (all ISO-hours)	**179,761**	**32.23**	**25.79**	**0.280**	**0.403**	**0.628**	**23.6**	**−20.0**
Hausman benchmark	—	—	—	*0.21*	*0.55*	—	—	*within 2.0*

The table carries two pooled rows because they answer different questions. Weighting each market equally gives an hourly correlation of 0.396 raw and 0.442 winsorised; stacking all 179,761 ISO-hours and computing one correlation gives 0.280 and 0.403. The second is the like-for-like analogue of a pooled published statistic and is the lower of the two, because the markets differ in price variance and the high-variance ones fit worst. The monthly correlation is 0.853 and 0.628 respectively. The raw benchmark of 0.21 is cleared by 4 of the five ISOs; the winsorised benchmark of 0.55 by 1.

Restricting to the base year moves the two criteria in opposite directions, and both movements are worth stating. The hourly correlation falls, from 0.473 pooled over all 21 usable ISO-years to 0.396 in 2024 alone; the monthly correlation rises sharply, from 0.697 to 0.853.

Figure A.12 shows the correlation between monthly mean price in the model and in the observed wholesale markets, showing high level of correlations and reasonable fit.

Figure A.13 shows that the model is able to track the seasonal and inter-annual profile closely in every market, in 2024 the monthly correlation is 0.853 across all markets.

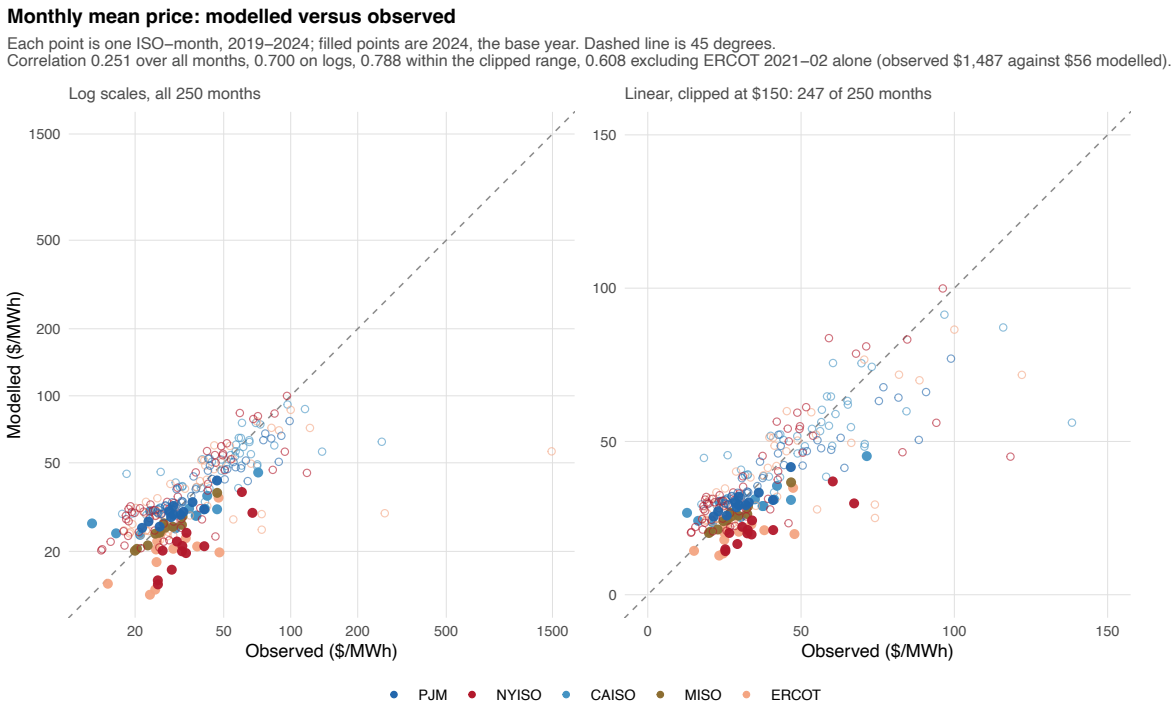


Figure (A.12) Monthly mean price, modelled against observed. Each point is one ISO-month, 2019–2024; filled points are the base year. Left panel is log–log so every month is visible, including ERCOT 2021–02 at $1,487 observed; right panel is linear and clipped at $150.

C.8 Solution Algorithm

The model is solved numerically via a damped Gauss-Seidel iteration. Within each outer iteration, the solver executes the following steps in sequence:

1. **Gravity:** Compute FMA and CMA indices. Extract sector-level price indices $P_{T,c}, P_{S1,c}, P_{S2,c}$ from the CMA.
2. **Environmental damage:** Compute county-level electricity demand, aggregate to NERC regions, run the plant dispatch model, compute $PM_{2.5}$ emissions, apply the AP3 source-receptor matrices, and form g_c and $\hat{B}_c$.
3. **Migration:** Given $\hat{B}_c$, $\hat{w}_{s,c}$, and $\hat{P}_c$, update bilateral migration shares and sectoral employment $L_{s,c}$ via the Fréchet hat-algebra perturbation of π^{obs}.
4. **Wages:** Update sector-specific wages via the FMA equation (32), using level-scale FMA and inverted productivities.
5. **DC entry and electricity:** Update data center allocation via the Berry logit (in counterfactual mode), then update BA-level electricity prices via demand hat algebra.
6. **Convergence check:** Track the maximum absolute change in wages and sectoral employment across counties until convergence.

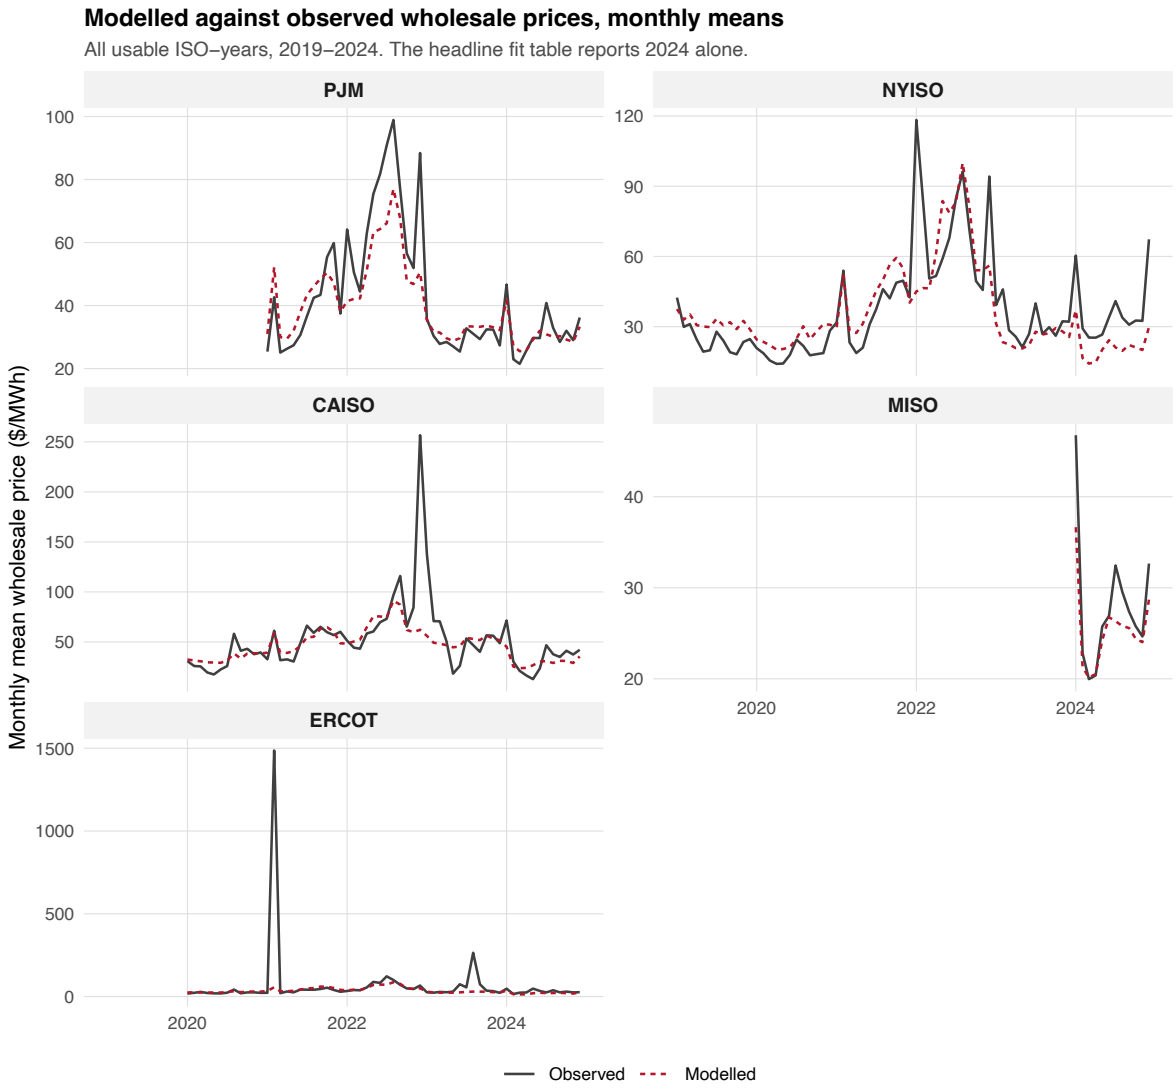
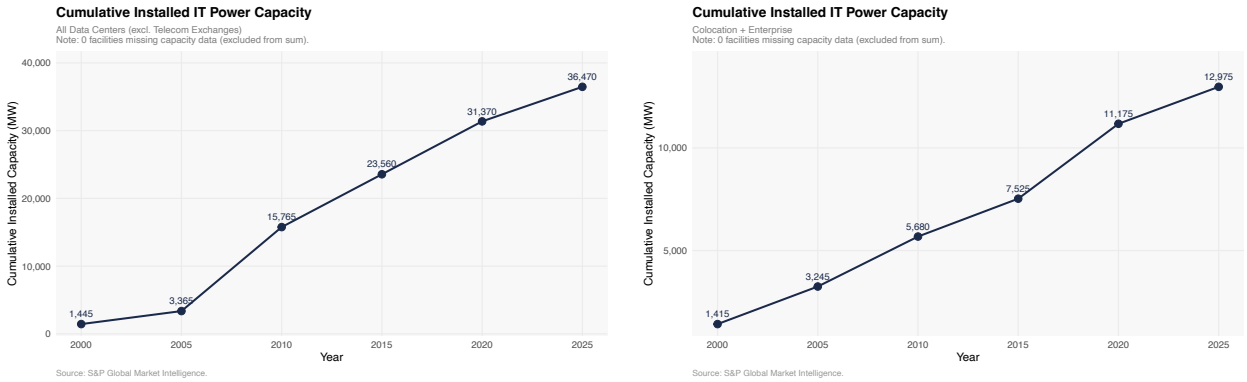


Figure (A.13) The same comparison over all usable ISO-years, so the base year can be read in context.

C.9 Appendix Figures

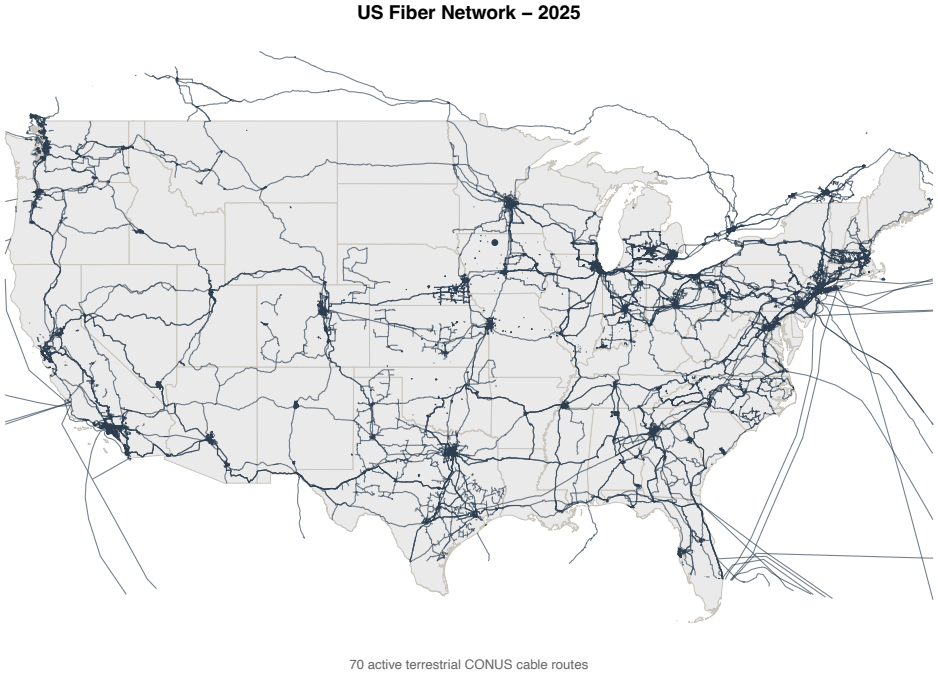
Figure (A.14) Installed Data Center Power Capacity Over Time, by Facility Type

(a) All Facilities (excl. Telecom) (b) Colocation and Enterprise



Notes: Cumulative installed IT power capacity (MW) of operational US data centers, by year. Panel (a) covers all facility types except telecom exchanges; Panel (b) covers colocation and enterprise facilities. The gap between the two panels corresponds to hyperscale capacity, which accounts for most of the growth after 2015. Source: S&P Global Market Intelligence.

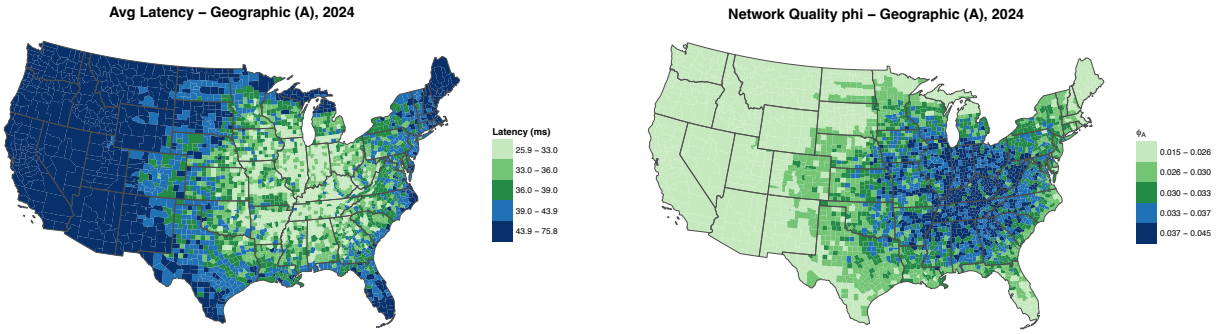
Figure (A.15) US Terrestrial Fibre Network



Notes: US long-haul terrestrial fibre-optic cable network from Infrapedia. The shortest-path distance between counties through this network, together with the distance to the nearest fibre node from the Internet Topology Zoo (Knight et al. 2011), are the connectivity measures used in the site-selection analysis (Section 4.1) and in the computation-quality index of the quantitative model (Section 6.9.1).

Figure (A.16) Network Latency and Computation Quality Index

(a) Network Latency (ms) (b) Computation Quality Index $\hat{\Phi}$



Notes: Panel (a) plots predicted county-level round-trip network latency (milliseconds) in 2024. Panel (b) plots the implied computation-quality index $\hat{\Phi}$, constructed from the latency gravity model as described in Section 6.9.1. Higher values of $\hat{\Phi}$ indicate lower effective latency and higher-quality access to computation.

C.10 Appendix Tables

Table (A.10) Summary Statistics, County-Year Panel, Firm & Worker Outcomes

Panel A: Overall Summary Statistics

	Mean	SD	p10	p90	N_counties	N_obs
Information Estabs per 10k	4.583	3.144	1.986	7.768	3057	66636
Manufacturing Estabs per 10k	11.854	6.567	5.043	19.746	3108	73865
Prof Service Estabs per 10k	18.891	17.188	7.258	34.150	3116	75196
Finance Estabs per 10k	15.772	7.711	8.226	24.258	3099	75777
Real Estate Estabs per 10k	9.732	10.037	3.043	17.305	3116	69979
Management Estabs per 10k	1.082	1.981	0.000	2.571	3003	55229
Total Estabs per 10k	241.574	120.298	136.376	356.116	3129	78225
Non-Fin Serv Estabs per 10k	170.801	82.171	97.241	251.773	3129	78225
Information Employment per 10k	63.255	88.361	10.682	133.330	3057	66636
Manufacturing Employment per 10k	550.430	510.582	70.764	1167.411	3108	73865
Prof Service Employment per 10k	129.509	200.739	24.897	277.380	3116	75196
Finance Employment per 10k	129.779	170.181	46.067	224.188	3099	75777
Real Estate Employment per 10k	46.066	64.241	7.330	95.460	3116	69979
Management Employment per 10k	50.984	139.795	0.000	135.119	3003	55229
Total Employment per 10k	3196.647	1940.268	1283.075	5448.480	3129	78225
Non-Fin Serv Employment per 10k	2187.888	1478.698	843.801	3880.924	3129	78225
Avg Broadband Providers per ZIP	3.841	2.326	2.000	7.125	3028	26353
N(Broadband Providers)	42.665	105.865	2.000	89.000	3028	27252
N(ZIPs with Provider)	8.505	11.843	1.000	18.000	3028	27252
Total High-Speed Connections	34.936	110.768	2.000	73.000	3078	43428
Employment 1990	0.680	0.078	0.575	0.770	3129	78225
Median HH Income 1990	58846.921	16113.937	42030.629	77837.773	3129	78225
Population 1990	78488.813	2.63e+05	5169.000	1.49e+05	3129	78225
College Fraction 1990	0.135	0.066	0.075	0.217	3129	78225

Panel B: Treated vs Control Counties

	Mean_Treat	SD_Treat	Mean_Control	SD_Control	Diff	p-value
Information Estabs per 10k	5.891	3.834	4.426	3.012	0.386	0.0069
Manufacturing Estabs per 10k	12.520	5.965	11.782	6.625	0.548	0.1734
Prof Service Estabs per 10k	33.929	25.922	17.306	15.145	7.704	0.0000
Finance Estabs per 10k	20.400	10.701	15.287	7.156	3.134	0.0000
Real Estate Estabs per 10k	14.553	9.393	9.181	9.961	2.623	0.0000
Management Estabs per 10k	2.519	2.413	0.885	1.828	1.111	0.0000
Total Estabs per 10k	307.250	134.026	234.866	116.746	35.957	0.0000
Non-Fin Serv Estabs per 10k	221.142	94.952	165.658	78.966	26.765	0.0000
Information Employment per 10k	157.664	177.868	51.913	61.203	56.070	0.0000
Manufacturing Employment per 10k	549.981	386.285	550.478	522.257	18.227	0.5249
Prof Service Employment per 10k	356.847	367.894	105.540	155.733	122.900	0.0000
Finance Employment per 10k	292.116	348.635	112.756	127.285	97.293	0.0000
Real Estate Employment per 10k	90.843	66.067	40.957	62.007	30.418	0.0000
Management Employment per 10k	153.892	177.679	36.868	127.442	70.123	0.0000
Total Employment per 10k	5278.641	2446.363	2983.974	1746.403	1399.024	0.0000
Non-Fin Serv Employment per 10k	3805.028	1847.110	2022.700	1329.290	1057.952	0.0000
Avg Broadband Providers per ZIP	6.185	3.395	3.588	2.024	1.845	0.0000
N(Broadband Providers)	182.748	285.054	27.997	40.075	92.956	0.0000
N(ZIPs with Provider)	24.316	28.216	6.849	6.534	11.737	0.0000
Total High-Speed Connections	193.080	295.270	18.017	35.988	53.958	0.0000
Employment 1990	0.717	0.065	0.676	0.078	0.027	0.0000
Median HH Income 1990	75170.072	20958.512	57179.534	14536.896	11028.659	0.0000
Population 1990	4.26e+05	7.40e+05	42947.857	82457.034	2.03e+05	0.0000
College Fraction 1990	0.208	0.086	0.128	0.058	0.051	0.0000

Summary statistics for US counties, 1999-2023. Panel A shows overall statistics. Panel B compares treated vs control counties.

Table (A.11) Summary Statistics, County-Year Panel, Electricity Market Outcomes

Panel A: Overall Summary Statistics

	Mean	SD	p10	p90	N_counties	N_obs
Industrial Rate ($/MWh)	72.19	41.43	45.45	97.34	2936	72706
Commercial Rate ($/MWh)	97.35	29.45	65.41	126.35	2944	73597
Residential Rate ($/MWh)	110.41	32.49	74.30	143.90	2944	72693
IOU Industrial Rate ($/MWh)	65.25	27.41	41.20	88.80	2335	56844
IOU Commercial Rate ($/MWh)	91.06	30.54	59.80	124.40	2340	57214
IOU Residential Rate ($/MWh)	109.34	34.95	72.90	147.18	2340	56206
Non-IOU Industrial Rate ($/MWh)	76.55	68.62	47.40	103.05	2349	56809
Non-IOU Commercial Rate ($/MWh)	100.28	29.77	68.04	129.00	2390	59282
Non-IOU Residential Rate ($/MWh)	109.12	28.53	73.94	141.51	2389	58670
(mean) num_ind_customers	2286.67	6836.33	13.50	4649.00	2944	73600
(mean) num_comm_customers	55612.68	82564.36	2807.00	123019.62	2944	73600
(mean) num_res_customers	396001.46	614147.08	14855.25	956854.25	2944	73600

Panel B: Treated vs Control Counties

	Mean_Treat	SD_Treat	Mean_Control	SD_Control	Diff	p-value
Industrial Rate ($/MWh)	74.57	38.94	71.97	41.65	2.60	0.0787
Commercial Rate ($/MWh)	100.34	30.81	97.06	29.30	3.28	0.0338
Residential Rate ($/MWh)	115.10	36.28	109.96	32.07	5.15	0.0049
IOU Industrial Rate ($/MWh)	71.63	35.85	64.62	26.36	7.01	0.0003
IOU Commercial Rate ($/MWh)	97.21	36.07	90.46	29.88	6.75	0.0012
IOU Residential Rate ($/MWh)	118.24	42.28	108.47	34.02	9.77	0.0000
Non-IOU Industrial Rate ($/MWh)	75.53	48.20	76.64	70.15	-1.11	0.5037
Non-IOU Commercial Rate ($/MWh)	100.47	27.24	100.26	29.99	0.21	0.8850
Non-IOU Residential Rate ($/MWh)	109.44	28.43	109.09	28.54	0.35	0.8202
(mean) num_ind_customers	2077.25	5677.04	2306.79	6937.20	-229.53	0.4009
(mean) num_comm_customers	68318.65	81952.10	54392.22	82520.62	13926.42	0.0065
(mean) num_res_customers	512921.96	629623.64	384770.82	611468.97	128151.14	0.0011

Summary statistics for US counties with electricity data, 2000-2024. Balanced panel only. Panel A shows overall statistics. Panel B compares treated vs control counties.