The Data Center Next Door: Macroeconomic Incidence of Compute

Ryan Wang*
Advisor: Pascual Restrepo†

April 1, 2026

Abstract

As gigawatt-scale data centers are deployed nationwide, the communities that host them face unprecedented energy demand shocks. Applying a TWFE design to a county-year panel, I show that computing facilities effectively act as a localized tax through electricity prices, driving up residential monthly fixed charges by 0.74% per 100 MW in regulated markets. Data centers also induce a unique crowding-out effect by depressing local manufacturing wages by 0.49% without affecting sector employment or output. A two-region DSGE model resolves this empirical puzzle: localized electricity hikes trigger a factor price displacement effect, forcing energy-intensive firms to pass input costs onto labor, while a counteracting wealth effect from larger electricity bills compels households to increase labor supply. Ultimately, electricity market architecture dictates the spatial distribution of this welfare loss. A simulated 100 MW load shock demonstrates that regulated regimes socialize grid investments, incurring a spatially symmetric 0.019% welfare loss. Conversely, deregulated nodal pricing traps physical congestion locally, inflicting a severe 0.150% welfare loss exclusively on host communities.

*Yale College
†Department of Economics, Yale University.

1 Introduction

On April 25, 2025, over two hundred residents of southwest Memphis filed into the Fairley High School auditorium to attend a public hearing on xAI's gas turbine permit request. Located less than seven miles away from the school, the artificial intelligence (AI) company's eight-month-old data center was busy training the latest version of its model—and releasing hazardous levels of nitrogen oxides in the process. At the public hearing, Memphians living near the new data center reported cases of asthma, cancer, and exacerbated health conditions that they attributed to air pollution.[1]

Without a formal voice in the process, the residents of Memphis are now on the frontline of the race to construct AI data centers. In the pursuit to develop smarter models and ultimately achieve artificial general intelligence, technology companies like xAI are investing hundreds of billions of dollars into computing infrastructure across the United States. This competition has captured the suspicious attention of American citizens, who have begun attributing soaring utility costs to their new data center neighbors. Mirroring the reports of Memphians, households living near data centers across the country are reporting quality of life issues related to environmental, noise, and light pollution.[2]

Focus on AI infrastructure has also become a recurring theme among policymakers at all levels of government. Some localities have aggressively courted data centers with tax incentives and streamlined permitting, aiming for tax revenue and job growth. Other local governments have enacted zoning reform, special use permits, and development moratoriums to manage or outright halt construction within their territory. At the state level, the policy momentum has shifted from one of incentivization to regulation. State legislatures introduced more than two hundred data center bills in 2025 alone, with proposals ranging from the elimination of tax exemptions to environmental mandates. The federal government is also actively working to address the new issues posed by the rapid expansion of computing infrastructure. In December 2025, the Federal Energy Regulatory Commission (FERC) ordered a grid operator to establish more transparent guidelines for data centers seeking to connect directly to power plants. In March 2026, Senator Bernie Sanders and Representative Alexandria Ocasio-Cortez introduced legislation that would impose a federal moratorium on all new data centers.

Whether it is taking place in the auditorium of a Tennessee high school or the halls of the Capitol Building, the growing national debate over AI data centers is the result of one fundamental challenge: the intense computation required to train and run large AI models demands an unfathomable amount of energy. At the time of the Memphis public hearing, xAI's data center was estimated to have a demand load of 361 megawatts (MW).[3] If one were to assume that the facility operates at near 100% capacity 24/7, the data center's energy consumption would be nearly equivalent to the residential demand of Seattle, Washington. With the ongoing construction and expansion of similar AI facilities, the American electrical grid faces the daunting task of safely adding a demand load equivalent to that of more than a dozen major American cities combined.

This paper is motivated by two questions plaguing the American public and policymakers: (1) how do data centers affect electricity prices, and (2) what local economic impacts do data centers have on the communities in which they are established? By combining nationwide county-level empirical analysis with a two-region dynamic stochastic general equilibrium (DSGE) model, I show that large-scale computing facilities act as a structural tax on host communities, shifting

[1] For more community testimony regarding the health externalities of the xAI facility, see Ariel Wittenberg, "'How come I can't breathe?': Musk's data company draws a backlash in Memphis," Politico, May 6, 2025.

[2] For a personal account on the experience of living near a data center hotspot in Northern Virginia, see Tess Martinelli and Kay Richards, "I live near 14 data centers. The noise and smell are terrible, but I refuse to leave my dream home," Business Insider, October 24, 2025.

[3] See Epoch AI for granular, satellite and permit-based estimates of AI data center power demand. As of April 2026, Epoch estimates the combined power demand of xAI's Memphis facilities to be 727 MW.

the immense infrastructure costs of the AI race onto captive local ratepayers while crowding out real wages in energy-intensive sectors.

In my empirical analysis, I apply a two-way fixed effects (TWFE) regression on county-year panel data with data center demand load as a continuous treatment. I find that each additional 100 MW in county data center demand load increases the probability of a residential rate case approval in regulated states by 0.65 percentage points, on average, driving up monthly fixed charges for households by 0.74%. I also find that each 100 MW increment in demand load outside the county but within the balancing authority (BA) region drives up residential fixed charges by 0.06%, indicating a spillover effect through the federally administered socialization of grid investment costs. While I find no evidence of data centers affecting local utility behavior in deregulated states, similar residential price hikes are likely occurring through generation-exposed retail energy providers unaccounted for in the Utility Rate Database.[4]

Most remarkably, I find limited evidence of data center demand loads spurring utilities to file rate cases to raise tariffs for commercial and industrial users across either regulatory regime. Counterintuitively, a localized 100 MW demand shock appears to drive down industrial rates in deregulated counties by 1.04%. These empirical observations reveal a deliberate asymmetry in utility cost-recovery: regulated utilities strategically shield highly elastic commercial entities by passing capacity burdens onto captive households, while deregulated utilities actively secure massive computing loads by offering bespoke, below-market volume discounts.

While the technology sector and host localities argue that data centers create a windfall for the local economy, my empirical analysis reveals a more nuanced story. I find that each additional 100 MW in demand load produces a 0.78% increase in aggregate county GDP. However, a by-sector breakdown indicates that this headline growth is driven almost entirely by the direct contribution of data center construction and operation. Furthermore, while this 100 MW increment is associated with a 0.44% increase in population and a 0.54% increase in employed residents, host counties experience no growth in aggregate wages or across an array of welfare metrics. Instead, the demand shock corresponds with a 0.69% increase in unemployment and a 0.49% decline in manufacturing wages. This presents a critical empirical puzzle: why do local manufacturing wages severely erode while the sector's aggregate output and employment remain statistically unaffected?

Drawing from my reduced-form results and the mechanics of electricity markets, I create and calibrate a novel two-region DSGE model to resolve this puzzle. The framework captures the macroeconomic impact of data centers on a host region and a non-host region under the same BA. Because empirical evidence shows that data center economic activity is largely isolated from the rest of the county economy, the model applies an exogenous data center load shock that affects households and firms solely through the electricity channel. To account for the structural divergence between regulated and deregulated states, I design distinct electricity price mechanisms for each regulatory regime.

The DSGE model demonstrates that localized computing expansion imposes a structural tax on the host community, driven by two counteracting mechanisms. First, an exogenous increase in electricity input costs triggers a factor price displacement effect. Because energy-intensive firms operate in competitive national markets with rigid production technologies, they cannot pass elevated energy costs onto consumers. Instead, these firms are mechanically forced to extract the penalty entirely from local labor, driving the observed erosion of real wages. This income collapse, compounded by surging residential utility bills, subsequently triggers a wealth effect. Compelled to smooth their rapidly compressing consumption, households must increase their labor supply, effectively stabilizing aggregate employment and output at the severe expense of regional welfare.

The model reveals that the spatial distribution of this welfare tax is strictly determined by

[4]See Section 2.3 for an overview of electricity market structure, and see Section 3.1 for data limitations of the Utility Rate Database.

the regulatory regime. Under the regulated model variant, the capital burdens of infrastructure expansion are systematically socialized. Both host and non-host regions face a unified grid variable rate shock, and non-host households absorb transmission cost adjustments. Consequently, a baseline 100 MW data center demand shock yields a relatively symmetric Consumption Equivalent Variation (CEV) of -0.019% in the host region and -0.016% in the non-host region.

Under the deregulated model variant, the unbundled market architecture structurally isolates the physical frictions of the load shock. Driven by locational marginal pricing (LMP), severe transmission congestion traps wholesale volatility within the host municipality. While this structure protects the broader BA from meaningful spillover, it amplifies the localized factor price displacement effect, resulting in a catastrophic host CEV of -0.150% compared to just -0.013% for the non-host region.

I finally apply the model to three high-fidelity, real-world case studies: (1) the historical build-out of "Data Center Alley" in Loudoun County, Virginia, (2) the development of xAI's Colossus facilities in Memphis, Tennessee, and (3) the deployment of OpenAI's Stargate infrastructure in Abilene, Texas. I find that the absorption of 7.0 gigawatts (GW) in Loudoun between 2014 and 2025 resulted in a CEV of -0.117% for the host county and -0.048% for the rest of the PJM region, translating to an annual consumption loss of \$98.68 and \$40.77 per household, respectively. In Memphis, the 1.9 GW AI-focused rollout is projected to produce a more symmetric toll across the heavily socialized TVA footprint, equating to an annual burden of \$129.34 for local Shelby County residents and \$121.18 for broader grid ratepayers. Conversely, the simulation of the 4.1 GW Stargate project within ERCOT's deregulated market reveals a catastrophic localized penalty. Lacking regional cost-sharing, the unprecedented load saturates local transmission capacity and triggers a CEV of -10.5%—equivalent to an overwhelming \$8,855.63 in lost annual consumption per local household.

My paper contributes to emerging evidence of the localized macroeconomic externalities associated with the AI race, demonstrating that households bear the cost burden of supporting their computing neighbors on the grid. Section 2 provides background on the relevant literature, the evolving nature of data center development, and the structure of American electricity markets. Section 3 details the data used in empirical analysis, its limitations, and evaluates the baseline comparability of the treatment and control cohorts. Section 4 presents the TWFE methodology and its empirical results. Section 5 sets up the structural model and the mechanics of the counteracting wealth and factor price displacement effects. Section 6 demonstrates the direct analytical calibration of electricity, household, and firm parameters. Section 7 applies the model to analyze a 100 MW shock and the three real-world scenarios. Section 8 concludes, highlighting limits to the paper's methodology and proposing the direction of future research.

2 Background

2.1 Relevant Literature

The growing body of economic literature on AI can generally be classified into four broad categories. The first branch of literature is concerned with how AI affects micro-level productivity and automation. Acemoglu and Restrepo (2018a) introduce a task-based framework in which automation reshapes labor demand through two counteracting forces: a displacement effect that reduces labor demand in automated tasks, and a productivity effect that raises labor demand in the remaining, non-automated tasks. Empirical studies investigating the displacement effect of AI corroborate this framework, suggesting early signs of falling labor demand in occupations with tasks most exposed to AI (Eloundou et al. 2024; Brynjolfsson et al. 2025; Hui et al. 2024). Meanwhile, randomized and quasi-experimental designs have revealed AI-driven productivity gains in specific environments (Brynjolfsson et al. 2023; Noy & Zhang 2023; Peng et al. 2023; Dell'Acqua 2023).

More recently, this branch of literature has explored the effect of AI at the level of an individual's workflow. Humlum and Vestergaard (2025) find that the wide-scale adoption of generative AI in Denmark has led to a substantial reorganization of work, with workers redirecting saved time from AI productivity gains to new job functions. The theoretical mechanics of this workflow-level restructuring are formalized by Demirer et al. (2026) through the concept of "AI chaining." By modeling production as a sequence of contiguous steps that can be manually executed, AI-augmented, or fully automated, Demirer et al. demonstrate the systemic redesign of a firm's job structure as AI quality improves and longer chains of tasks are automated.

The second branch of the literature studies the role that AI plays in reshaping the frontier of micro-level knowledge production. Theoretical work has envisioned the relationship between AI and knowledge in diverse ways. Cockburn et al. (2018) conceptualize AI as a "general-purpose method of invention" that can transform the process of innovation. More recently, Gans (2025) formalizes AI as an "interpolative technology" that reduces the uncertainty of forming conjectures to bridge gaps in existing knowledge, while Athey and Scott Morton (2025) model AI as a priced, imported factor of production within an open-economy general equilibrium. On the empirical front, emerging evidence documents how AI adoption may complement product innovation and restructure the R&D process at the firm level (Babina et al. 2024; Arenas Días et al. 2025; Dell'Acqua et al. 2025). These empirical findings remain preliminary, reflecting the early stage of AI diffusion in commercial innovation and the limited availability of comprehensive data.

The third branch of literature seeks to aggregate the micro-level effects of AI to characterize the implications for long-run growth and structural macroeconomic change. Acemoglu and Restrepo both build on their task-based framework in a macro context, deriving insights on changes to total factor productivity, wage inequality, and factor income shares (Acemoglu & Restrepo 2018b, 2022; Acemoglu 2024; Restrepo 2025). Meanwhile, Aghion et al. (2017) incorporate both AI-driven productivity and idea production into a growth model, exploring conditions under which AI can generate explosive growth and identifying the bottlenecks that would prevent economic singularity. Jones (2026) expands on the nature of barriers to singularity, extending task-based models with the concept of "weak link" tasks that will forever require human labor and thus become binding constraints on aggregate output.

While the first three branches of literature focus on the direct consequences of AI adoption, the fourth branch has recently emerged to address the systemic externalities generated by AI's physical and financial deployment. Benetton et al. (2025) document severe electricity-channel negative spillovers from data centers dedicated to cryptomining operations in Upstate New York, calculating an impact of $1.5 billion in extra costs for households in counties with the largest data center footprints. Kay, Reaser, and Taylor (2026) build a granular model of U.S. wholesale electricity markets, which demonstrates existing data centers driving a 2–6% increase in average wholesale electricity prices nationwide. The model also produces localized effects at the

eGRID subregion level, finding a greater than 10% increase in wholesale prices in the subregion hosting "Data Center Alley" in northern Virginia. Knittel, Senga, and Wang (2025) explore the environmental aspect of data center development, concluding that compute infrastructure with temporally inflexible demand loads risks crowding out broader de-carbonization efforts by forcing grid reliance on legacy thermal baseload generation, thereby driving up regional emissions.

Beyond physical infrastructure, the sheer scale of capital required for training and running AI is actively reshaping financial markets. While academic literature on such financial channels is sparse, De Vere, Ramaswamy, and Searls (2026) of the Dallas Federal Reserve argue that the massive issuance of long-term, investment-grade corporate bonds to fund AI data center projects introduces a profound crowding-out effect on traditional financial issuers.

This paper contributes to the fourth and youngest branch of AI literature by providing two novel advancements. To my knowledge, no prior research has used nationwide data to capture the effect of data centers on retail electricity prices. While existing work provides the direct cost to households and firms from higher utility bills, I also find no literature that calculates the total welfare costs that result from general equilibrium effects. This paper (1) provides empirical observations from a constructed national dataset with county-year granularity on data center demand load, retail electricity prices, and economic conditions, and (2) proposes a structural model that reveals the general equilibrium effects of data centers through the electricity channel.

2.2 Data Center Development

In essence, data centers are computers scaled up to the size of a building. Servers handle computation, storage arrays manage and retain data used by the servers, and networking equipment directs rapid communication between servers, storage systems, and the outside world. Just as individual computers rely on power supply units, fans, and heat sinks, data centers have power distribution and cooling systems as critical supporting infrastructure.

The development of data centers can be characterized by three waves. The first wave took place during the early internet and dot-com era. Advancements in internet connectivity, web-hosting, and corporate IT drove companies to build and operate their own private data centers. These computing facilities were designed to centralize critical business data and support corporate internet presence. Facility demand loads were generally less than 5 MW and easily integrated into existing commercial grid infrastructure.

The launch of Amazon Web Services (AWS) in 2006 marked the beginning of the second wave in data center development. By offering to host and process corporate data on Amazon servers, the growing e-commerce retailer allowed other companies to abandon the massive capital expenditures of building private facilities in favor of flexible pay-as-you-go operational expenditures. The mass migration of corporate data to "the cloud" managed by Amazon and other technology companies gave rise to the "hyperscale" facilities that began demanding over 50 MW in power. Hyperscaler power demands grew alongside American companies from the mid-2000s to present day, reflecting the energy required to manage the data of flourishing enterprises. As a result, hyperscale facilities began requiring dedicated substations and sparked tension with local utilities.

OpenAI's release of ChatGPT in late 2022 sounded the starting pistol of the AI race. In this third wave of development, participants are aggressively seeking computing power to train multi-trillion-parameter foundational models. AI companies also require computing facilities to run deployed models and provide rapid and high-quality responses to end users—a process known as inference. These new data centers dedicated to AI demand between 100 MW to over 1 GW. Whereas private corporate and cloud data centers have diurnal demand load profiles based on human internet traffic, AI training facilities run at near 100% utilization for months at a time. As discussed in Section 1, the arrival of these demand loads has sparked immense public backlash and political turmoil as utilities, grid operators, and policymakers navigate the

Figure 1: Operational Data Center Load in the U.S. (1990–2025)

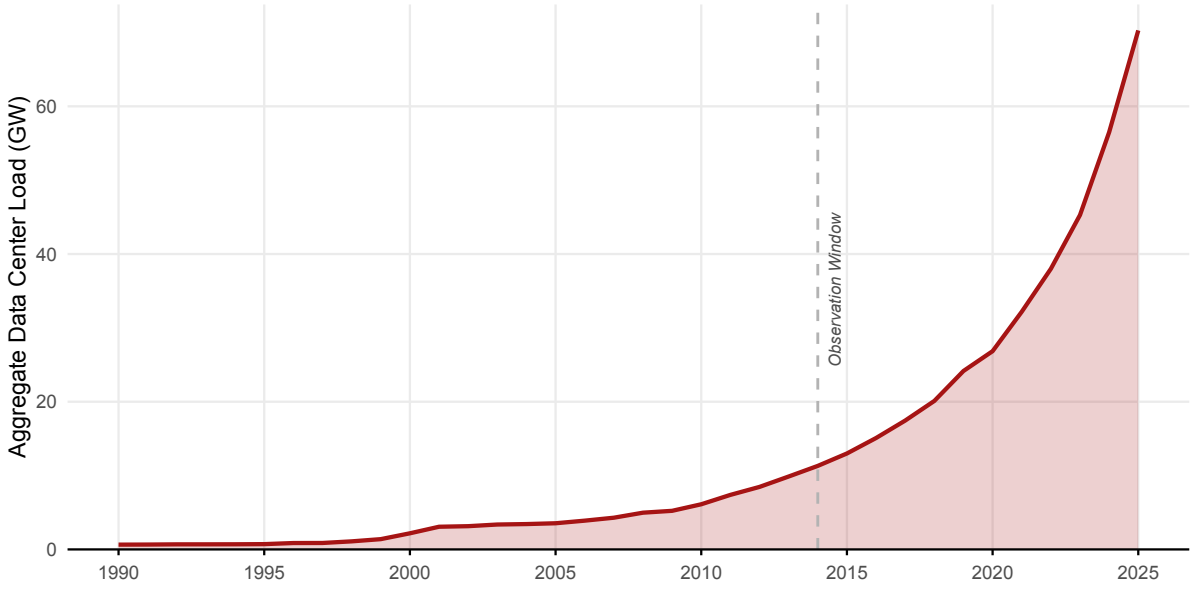

massive grid infrastructure investments required to accommodate the sudden influx of energy demand.

Data centers offer two economic promises to their host communities: tax revenue and job creation. Hyperscale and AI facilities can create a revenue windfall for localities through the collection of high-value property taxes, sales taxes on rapidly depreciating equipment, and corporate income taxes. In addition, data centers generally require minimal public services compared to traditional commercial and industrial businesses, making them a low-cost taxpayer from the perspective of local governments. However, many states and localities offer lucrative incentives (e.g., exempting server hardware from sales tax) to attract data centers, creating a regulatory "race to the bottom" that greatly diminishes the actual tax revenue generated by facilities.[5]

Although tax revenues constitute the primary fiscal incentive for local governments, the promise of job creation often serves as the central justification for garnering community support. The labor market impact of a data center can be bifurcated into two distinct phases. The initial construction phase generates a massive, albeit temporary, surge in local employment. A hyperscale or AI project requires a mobilization of labor that often exceeds 3,000 personnel. Key workers include specialized construction crews, electrical/mechanical engineers, master electricians, HVAC technicians, and fiber optic specialists. These highly specialized jobs are rarely sourced from the host community and are brought in from established firms elsewhere. While the primary labor income from construction may leave the host community, the massive transient workforce likely provides a short-term stimulus to the local service and hospitality sectors.

After the completion of construction, the data center enters the operational phase. Once online, expansive facilities are often maintained by a remarkably lean permanent staff of fewer than a hundred individuals. The limited employment growth is driven by contracted cybersecurity, systems architecture, and general facility management. Public discourse on data centers often centers on the limited permanent job creation that computing facilities bring to host communities, often overlooking the greater threat to the local economy: the possibility of data centers raising electricity costs for other sectors.

[5]The Memphis Chamber of Commerce claims that xAI will likely pay over $100 million in local tax revenue.

2.3 The American Electricity Market

Electricity is unique among market goods in requiring the instantaneous, real-time matching of supply and demand. Without a way to store energy at scale, the grid requires that the aggregate energy generated by power plants exactly meets the energy demands of all end users at any given moment. Failure to maintain this delicate balancing act risks infrastructure damage and grid failure. This physical constraint has produced an idiosyncratic and fragmented market structure.

The physical grid is separated into three distinct segments: (1) generation by power plants, (2) transmission over long distances via high-voltage power lines, and (3) distribution to end-users via low-voltage power lines. While electricity is generated and distributed at lower voltages, long-distance transmission uses high voltage to minimize power loss. While the physical grid is interconnected, the method through which electricity is managed and sold is highly fractured across the country. Currently, three distinct regulatory regimes exist.

The first regime is the traditionally regulated market (TRM), which services about one third of U.S. electricity demand.[6] In regions under TRMs, the entire three-segment electricity supply chain is owned by vertically integrated utilities that maintain monopolies within their respective territories. While these utilities may engage in bilateral wholesale trading among each other, no centralized wholesale electricity market exists. Rather, these regional monopolies operate as independent balancing authorities, bearing the responsibility of matching generation to load in real-time across their footprints. To maintain reasonable electricity prices for end users, states operate a public utilities commission (PUC) that oversees the utility price setting process. If a vertically integrated utility would like to raise its variable rate or fixed charge schedule, it must request permission from its PUC in a formal regulatory process known as a rate case.[7] To get a rate case approved, utilities must demonstrate that their requested increases are needed to cover operating costs and investments in grid infrastructure, allowing them to recover costs and earn a reasonable rate of return. Memphis, the home to xAI's new AI data centers, is located in a uniquely structured TRM. While Memphis Light, Gas and Water is the monopoly utility of the city, it cedes generation and balancing authority to the federally owned Tennessee Valley Authority (TVA).

The second regime, which I term a restructured-regulated market (RRM), represents the dominant regulatory architecture for most of the country. Under RRM regions, utilities are not allowed to own generating assets. Instead, power plants are owned by distinct private entities that sell to utilities on a centralized, competitive wholesale market. The wholesale market is run by independent, non-profit entities known as Independent System Operators (ISOs) or Regional Transmission Organizations (RTOs). Under RRMs, ISO/RTOs take on the role of the balancing authority. Critically, utilities still own all transmission/distribution infrastructure and operate as retail monopolies within their respective territories. Like their TRM counterparts, RRM utilities must still file rate cases with PUCs. Loudoun County, the host of "Data Center Alley," is located in an RRM. While the county lies in the territory of the RTO PJM (Pennsylvania-New Jersey Maryland Interconnection), the state of Virginia retains a predominantly regulated retail market.

The third regulatory regime is the deregulated market (DM). This type of market is formed when a state opts to deregulate its retail electricity market, eliminating the utility monopoly in the retail market. Under DMs, ISO/RTOs take on balancing authority and operate a competitive wholesale market, playing the exact same role they do in RRMs. However, DMs open the wholesale market to an array of demand-side buyers alongside traditional utilities. Retail Energy Providers (REPs) are the dominant wholesale buyers in DMs. These private entities purchase wholesale electricity from generators and offer retail plans to end-users, typically residential

[6] For a more expansive overview of electricity market structures, see Kathryne Cleary and Karen Palmer, "US Electricity Markets 101," Resources for the Future, March 17, 2022.

[7] If a regulated utility operates in multiple states, an approved rate case from one state's PUC will only allow the utility to raise prices in that state.

households and standard commercial businesses. REPs commonly offer fixed and variable-rate plans, with contract lengths of 12, 24, and 36-month terms.[8] Large commercial and industrial (C&I) customers (e.g. data centers, manufacturing plants, chemical refineries) are also allowed to participate directly in the wholesale market. C&I entities often sign direct power purchase agreements (PPAs) with generators to secure long-term, below-market rates. Importantly, traditional utility companies still own all transmission/distribution infrastructure. Often referred to as providers of last resort or default service providers, utility companies are generally required to buy wholesale power to supply end-users that do not select a REP.[9] As a result, retail electricity bills in DMs are comprised of two components: paying (1) the REP for procuring electricity from the wholesale market and (2) the local utility company for delivering the electricity. Despite the label of "deregulated market," utilities in DMs must still seek rate case approval with PUCs to raise prices related to their component of the bill. Abilene, the home of OpenAI's ongoing Stargate development, resides in the deregulated state of Texas under the balancing authority of the ISO ERCOT (Electric Reliability Council of Texas).

For the purpose of this paper's macroeconomic framework, the critical distinction between the three regulatory regimes lies in the structural mechanism by which wholesale grid costs are passed through to retail users. Under TRMs and RRMs, both the generation and delivery components of electricity bills are bundled and administered by local utility monopolies under the strict oversight of state PUCs. While RRM utilities participate in a competitive wholesale market, they still function as financial buffers that aggregate localized wholesale price volatility and socialize those costs into a blended, regulated retail price for end-users. Two crucial mechanisms allow these regulated utilities to pass certain wholesale grid costs to end-users without initiating a formal, protracted rate case process. First, the Power Cost Adjustment (PCA) is a variable rate rider that passes through fluctuations in fuel and wholesale electricity purchasing costs directly to retail customers. Second, the Transmission Cost Adjustment (TCA) is typically an additional fixed charge that passes through changes in the cost of transporting bulk electricity. These adjustments allow regulated utilities to mechanically recover costs incurred by wholesale market volatility (via the PCA) and regional infrastructure upgrades (via the TCA).

While the pass-through mechanisms under the two regulated regime types involve a carefully orchestrated legal process with specific PUC exceptions, retail customers in DMs face a bifurcated bill structure. The competitive Retail Electric Provider (REP) assumes the financial risk of wholesale energy procurement and charges market-driven retail prices for electricity generation. Meanwhile, the local utility remains subject to state PUC oversight and levies a regulated tariff strictly to recover capital investments in transmission and distribution infrastructure. While DM utilities are not exposed to generation risk (i.e., no PCA), they retain TCAs to mechanically pass through capital expenditures on local transmission and distribution. As the paper will demonstrate in subsequent sections, this structural bifurcation in DMs creates an asymmetric exposure to data center load shocks compared to TRM and RRM environments. While institutional differences between the TRM and RRM regime structures must be acknowledged, I opt to group TRM and RRM regions under a single "regulated market" label for the remainder of the paper to reflect their shared macroeconomic incidence and stark contrast to deregulated markets.

[8] Readers may confuse the terms fixed and variable-rate plans with fixed charges and variable rates. Fixed and variable-rate plans refer only to the variable rate element of the bill, or the price per kilowatt-hour that an end-user must pay. Fixed-rate plans provided by REPs generally still have varying fixed monthly charges.

[9] Utility companies will also supply end-users whose selected REPs go out of business or fail to deliver electricity.

Figure 2: County-Level Operational Data Center Load

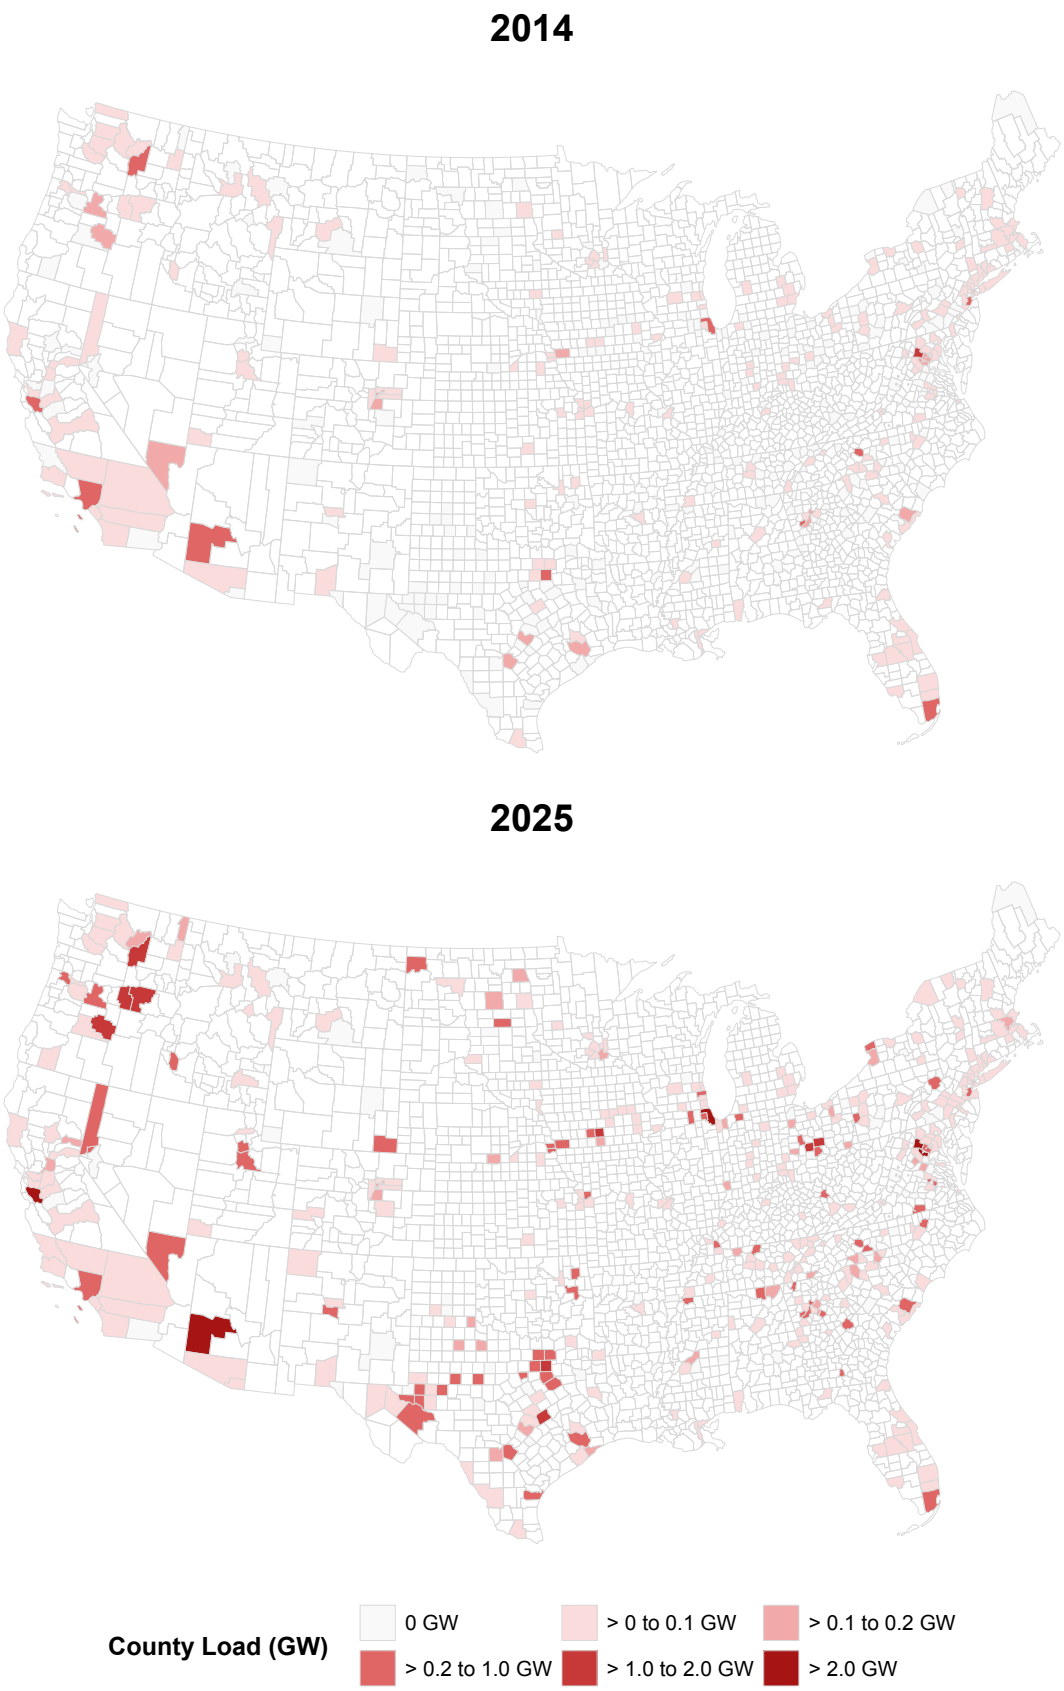

3 Data

3.1 Sources and Sample Construction

To evaluate the causal effect of data centers on local electricity markets and host economies, I construct a balanced county-year panel linking computing infrastructure deployment with retail electricity tariffs and macroeconomic outcomes. I acquire facility-level computing infrastructure data from the S&P Global 451 Research Datacenter KnowledgeBase (DCKB), which tracks 4,688 U.S. data centers using primary research from site visits, corporate briefings, and surveys. While comprehensive, the DCKB presents two temporal limitations. First, data on construction start times are sparse, precluding an analysis of anticipatory effects or the short-term economic impacts of the construction phase itself. Second, the dataset is primarily resolved at an annual frequency. Although some facility observations contain quarterly data, this granularity is not uniform, preventing the study of higher-frequency, month-to-month price volatility following a load shock. Consequently, I utilize each facility's geospatial coordinates, operational timeline (year of operation and decommission), and maximum utility power demand to aggregate the total operational data center load (in MW) for each county at an annual frequency. To simulate the general equilibrium effects of the xAI Colossus and OpenAI Stargate developments in Section 7.2, I use facility-level forecasts from Epoch AI.

To construct the county-year retail electricity panel, I combine tariff data from the Open Energy Information's Utility Rate Database (URDB) with utility service territory crosswalks from the U.S. Energy Information Administration's (EIA) Form 861. The URDB provides detailed retail rate structure information for over 3,700 utilities nationwide, with each observation detailing the pricing components of an approved case. Because utilities do not strictly adhere to county boundaries, I utilize EIA-861 territory data to map these utility-level tariffs to their respective counties. By synthesizing these sources, I construct the retail electricity market

Table 1: Retail Electricity Variables at County-Year Level

Variable Name	**Unit**	**Methodology**
Rate Case Approved	Indicator (0/1)	Identifies if a new tariff went into effect during the observation year. Aggregated to the county level (equals 1 if any utility serving the county acquires an approved rate case).
Variable Rate	Cents/kWh	Averages the Tier 0 energy charge (base rate + adjustment rate) across all periods. Aggregated to the county level via median across local utilities.
Fixed Charge	Dollars/month	Directly extracts the fixed charge for the first meter. Aggregated to the county level via median across local utilities.
Primary BA	Categorical	Assigns a single BA by mapping utilities to their corresponding BAs using EIA territory data. If utilities operating in the same county-year report to different BAs, the BA of the utility with the highest sales in MWh is selected.

Note: All metrics are calculated separately for the residential, commercial, and industrial rate class. The calculated rate and charge variables explicitly exclude EV, solar, net metering, time-of-use (TOU), low-income, green, and standby tariffs. Non-positive values and top 1% outliers can largely be attributed to misreported data and are also removed.

Table 2: Economic and Environmental Data at County-Year Level

Source	County-Year Variable	Years Available
Bureau of Economic Analysis - BEA	Aggregate GDP	2014–2024
	Sectoral GDP	2014–2024
Bureau of Labor Statistics - BLS	Labor Force	2014–2025
(Local Area Unemployment Statistics)	Employed	2014–2025
	Unemployed	2014–2025
	Unemployment Rate	2014–2025
U.S. Census Bureau	Sectoral Establishments	2014–2023
(County Business Patterns - CBP)	Sectoral Employment	2014–2023
	Sectoral Derived Wages	2014–2023
U.S. Census Bureau	Median HHI	2014–2024
(American Community Survey - ACS)	Poverty Individual Count	2014–2024
	SNAP HH Count	2014–2024
Zillow	Home Value Index (ZHVI)	2014–2025
	Observed Rent Index (ZORI)	2014–2025
National Oceanic and Atmospheric Admin.	Annual HDD	2014–2025
	Annual CDD	2014–2025

Note: Years available reflect the limits of institutional bulk data releases at the time of dataset construction. Sectoral variables (GDP, establishments, employment, and derived wages) cover specific comparative industries (NAICS Code): Information/Data Centers (518), Construction (23), Professional & Business Services (54-56), Manufacturing (31-33), Agriculture (11), and Service (71, 72, 81). Derived wages from the CBP are calculated by dividing total annual payroll by total employment for each respective NAICS grouping and are winsorized at the 1st and 99th percentiles to mitigate the impact of outliers.

variables detailed in Table 1, calculating metrics separately for the residential, commercial, and industrial rate classes. The URDB offers comprehensive retail rate structure data covering most counties from 2014 onward, effectively establishing 2014–2025 as the primary observation window for the paper's empirical analysis. This period captures a massive structural shift in computing infrastructure, with nationwide data center loads scaling from 11 GW to over 70 GW. This observation window allows my analysis to observe the transition from the peak cloud computing era into the nascent AI boom, as highlighted by Figures 1 and 2.

While comprehensive in its coverage of utility rate cases, the URDB fails to cover two critical components relevant to retail electricity pricing. In regulated states, rate case data do not capture the PCA and TCA that utilities pass directly onto end-users. In deregulated states, rate case data only capture the bill component charged by traditional utilities, as REP pricing is determined by market competition outside the purview of PUCs. Consequently, the strict reliance of my empirical analysis on approved base rates likely attenuates the estimated treatment effect of computing infrastructure on electricity costs. Because wholesale price spikes and grid investment induced by major data center load shocks are typically recovered through these unobserved pass-through mechanisms, the current panel only represents base tariffs and may systematically underestimate the true financial burden shifted onto host communities.

To control for different electricity market regulatory regimes, I integrate state-level deregulation statuses sourced from Electric Choice. While Section 2.3 broadly describes retail deregulation as a policy granting Retail Electric Provider (REP) selection to consumers, the implementation of these markets is heavily fragmented. Currently, fourteen states operate fully deregulated retail markets across all rate classes, while five states have adopted partial deregulation—typically restricting residential choice while granting C&I entities greater procurement flexibility. To accurately capture this heterogeneity, I construct separate deregulation indicators

for the residential, commercial, and industrial sectors, assigning these flags to counties based on their state jurisdiction. Because no major state-level shifts in retail deregulation policy occurred between 2014 and 2025, these indicators are time-invariant within the panel.

Finally, to comprehensively assess the economic impact of computing facilities on host communities, I incorporate an array of localized economic variables capturing output, labor markets, household welfare, and housing costs. Table 2 outlines the institutional sources and specific metrics utilized in the empirical analysis. Due to varying federal publication lags, comprehensive data for certain economic outcomes are not yet available through 2025. Consequently, subsequent analyses of sectoral output, wage dynamics, and household welfare are constrained by the terminal years of their respective datasets. In addition to these economic indicators, I integrate county-level Heating and Cooling Degree Days (HDD/CDD) as time-varying controls, as local utility demand and retail prices are highly sensitive to weather fluctuations.

3.2 Sample Balance

Two critical features of data center development complicate the paper's causal identification strategy: (1) left-censoring of treatment timing, and (2) reverse causality between data center placement and outcome variables. As alluded to by Figures 1 and 2, the expansive development of cloud computing infrastructure partially predates the 2014 start of the study's observation window. Consequently, several counties in the treatment group of data center hosts are "early adopters" that already house significant demand loads at the beginning of the panel. Loudoun County serves as the most extreme example of this issue, hosting a little over 1.0 GW in demand load in 2014. Such left-censoring of treatment timing means that the 2014 baseline measurements for treatment counties may already capture early post-treatment effects. Second, the spatial distribution of data centers is highly endogenous. Technology companies likely select host communities based on the study's dependent variables. Specifically, data center operators are generally known to prefer regions with low industrial electricity rates, favorable climates for server cooling, and healthy local macroeconomic conditions.

To diagnose these identification challenges and assess cohort comparability, I examine the 2014 cross-sectional baseline balance between treatment and control counties (see Table 3). I define a "treatment county" as one recording at least 10 MW of operational data center demand at any point during the observation window. This threshold ensures counties hosting only negligible, small-scale computing facilities remain in the control group. The balance table highlights stark differences between treatment and control counties. The topmost row of Table 3 reveals the first challenge of left-censored treatment timing. At the start of observation, the average treatment county already houses 45 MW of data center demand load. While a normalized difference of 0.529 indicates a clear baseline imbalance between the two cohorts, I argue that this pre-existing capacity introduces minimal distortion to estimated treatment effects. First, the 45 MW average among the treatment group is heavily right-skewed by a small number of early adopters like Loudoun County. Critically, 140 of the 247 treatment counties (56.7%) possessed less than 10 MW of operational load at the 2014 baseline. Second, 45 MW roughly equates to a single legacy hyperscale facility. This demand load represents a fraction of the colossal, multi-100 MW load shocks that define data center development later in the observation window. Most importantly, the TWFE specification (introduced in Section 4) structurally absorbs the cross-sectional level effect of this 2014 baseline capacity through county-level fixed effects (FEs).

Subsequent rows in Table 3 highlight the second challenge, demonstrating that location selection of data centers is not a random process. Computing facilities are highly concentrated in wealthy, populous, metropolitan hubs. The lower treatment average in annual HDD also indicates a tendency to locate in areas with slightly milder winters. This finding aligns with the structural requirements of data center architecture, as operators prioritize proximity to major internet exchange points and dense fiber-optic networks over the ambient cooling benefits of colder, remote regions. This difference in local economic composition and climate poses a

Table 3: Balance Table of Treatment vs. Control Counties in 2014

Variable	Control Mean	Treatment Mean	Difference	Norm Diff
Demand (100 MW)	0.00	0.45	0.45	0.529
Population	64,792.95	516,260.65	451,467.70	0.674
GDP	2,964,718.16	36,701,753.08	33,737,034.92	0.633
Unemployment Rate	6.21	5.88	-0.34	-0.165
Median HHI	45,543.64	55,873.70	10,330.06	0.744
Annual HDD	5,406.19	4,812.80	-593.40	-0.268
Annual CDD	1,152.00	1,258.57	106.57	0.131
Service Intensity	0.15	0.16	0.01	0.196
Professional Intensity	0.05	0.13	0.08	1.089
Manufacturing Intensity	0.12	0.11	-0.01	-0.057
Agriculture Intensity	0.00	0.00	0.00	-0.268
Res. RC Approval	0.85	0.92	0.07	0.220
Res. Var Rate	9.81	10.22	0.41	0.169
Res. Fixed Charge	14.15	11.86	-2.30	-0.421
Com. RC Approval	0.86	0.93	0.07	0.240
Com. Var Rate	8.76	8.91	0.15	0.058
Com. Fixed Charge	54.36	43.93	-10.43	-0.108
Ind. RC Approval	0.76	0.84	0.09	0.216
Ind. Var Rate	5.48	5.20	-0.28	-0.124
Ind. Fixed Charge	1,093.04	695.76	-397.28	-0.145

Note: Norm Diff represents the Imbens and Rubin (2015) normalized difference. A normalized difference value $> |0.25|$ indicates a substantial imbalance between the treatment and control groups. Sector intensity is the percentage of employed individuals in the county working within the sector. The mean rate case (RC) approval represents the percentage of counties within each cohort that had a rate case approval in 2014. Variable rate is measured in cents per kWh, and the fixed charge is measured in dollars per month.

greater threat to causal identification than the first challenge, as county and year FEs in a traditional TWFE model do not absorb diverging macroeconomic trends that emerge from this stark baseline heterogeneity. If metropolitan hubs and rural counties were already on fundamentally different growth trajectories for utility pricing or economic outcomes prior to data center expansion in our observation window, the core parallel trends assumption would be violated.

Despite the evident imbalances in baseline macroeconomic scale and composition, the retail electricity market conditions exhibit reassuring comparability. Table 3 shows that rate case approval frequency, variable rates, and fixed charges are largely balanced across the three rate classes. The sole exception within the utility variables is the residential fixed charge, which is notably lower in treatment counties with a normalized difference of -0.421. This discrepancy can largely be attributed to a mechanical artifact of customer density. Utilities operating in highly concentrated urban hubs can socialize fixed grid capital expenditures and operating costs across a vastly larger ratepayer base per square mile, reducing the per-household fixed charge relative to sprawling control counties.

Because the retail electricity market conditions most relevant to data centers (C&I variables) are balanced at baseline, I opt to retain the full national sample of counties rather than artificially restricting the control group via an arbitrary method of sample trimming. I address the remaining challenge of spatial sorting by introducing time-varying FEs and climate controls in my TWFE specifications. These inclusions are designed to absorb unobserved, differential electricity market and economic growth trajectories correlated with initial economic composition and climate. I discuss both of these empirical strategies in Section 4.

4 Empirical Analysis

4.1 Identification Strategy

The objective of this section is to isolate the causal effect of data center expansion on retail electricity prices and local economic outcomes. The nature of computing infrastructure development presents two standard threats to causal identification that motivate the adoption of a TWFE approach. First, treatment and control counties possess fundamental structural differences (as proven in Section 3.2) that create varying baseline levels in macroeconomic outcomes. County FEs are thus required to absorb this time-invariant heterogeneity between counties. Second, electricity markets and local economies are subject to nationwide energy trends and macroeconomic cycles (e.g., fuel prices, inflation). These uniform time shocks must therefore be controlled for by implementing year FEs. While county FEs absorb a county's average climate, retail electricity prices (and thus local economic outcomes through the electricity channel) are highly sensitive to annual weather volatility. Therefore, I introduce county-level annual HDD and CDD as time-varying controls to ensure that weather shocks and spatially heterogeneous climate change do not confound the estimated effect of data center loads.[10]

4.1.1 Quintile-by-Year Fixed Effects

However, there remain two challenges to causal identification that are not addressed by the standard TWFE framework. As discussed in Section 3.2, the non-random spatial sorting of data centers into wealthy, populous, mild-winter hubs is one such challenge. While county FEs address level differences, they do not absorb trend differences. The divergence in regional economic trajectories between 2014 and 2025 underscores the plausibility of non-parallel growth paths between the two cohorts. Specifically, knowledge-based agglomeration externalities have fostered a disproportional concentration of high-tech, finance, and professional services within certain American metropolitan areas. This structural transformation likely accelerated local output, wage inflation, and C&I electricity demand loads independent of data center deployment. In contrast, counties with manufacturing- or agriculture-intensive economies did not benefit from these externalities and remained more susceptible to the secular pressures of globalized trade. Finally, these cohorts likely faced heterogeneous exposure to the COVID-19 pandemic. The viability of remote work in professional hubs likely insulated treatment counties against the acute labor market shocks that destabilized the in-person sectors prevalent in control localities.

Figure 3 illustrates where parallel trends in retail electricity market outcomes between the treatment and control groups hold and break down throughout the observation window. Across all three rate classes, the frequency of rate case approval moves in near-perfect lockstep—including the sharp, synchronized drop during the pandemic. The median commercial variable rate also tracks remarkably well throughout the study. However, clear divergences emerge among other electricity price variables that cannot be credibly attributed to treatment counties being exposed to greater data center loads. Industrial variable rates diverge visibly between 2018 and 2022, with the treatment group dropping significantly below the control group before sharply accelerating to converge by 2023. Fixed charges across all three rate classes feature sharp volatility and a lack of clear parallel movement, particularly for the industrial sector. These visual deviations are likely the manifestation of the aforementioned heterogeneous growth paths.

[10] Heating and Cooling Degree Days (HDD/CDD) capture the temperature of a location in the context of energy use. The mean temperature of a location in a single day $\bar{T}$ is defined as the average between the highest and lowest temperature recorded that day, measured in Fahrenheit. CDD is defined as $\max(\bar{T} - 65°\text{F}, 0)$ and indicates the air conditioning required to keep buildings cool. HDD is defined as $\max(65°\text{F} - \bar{T}, 0)$ and indicates the heating required to keep buildings warm. HDD/CDD are widely used in the energy sector to analyze and forecast energy demand.

Figure 3: Parallel Trends in Retail Electricity Pricing (2014–2025)

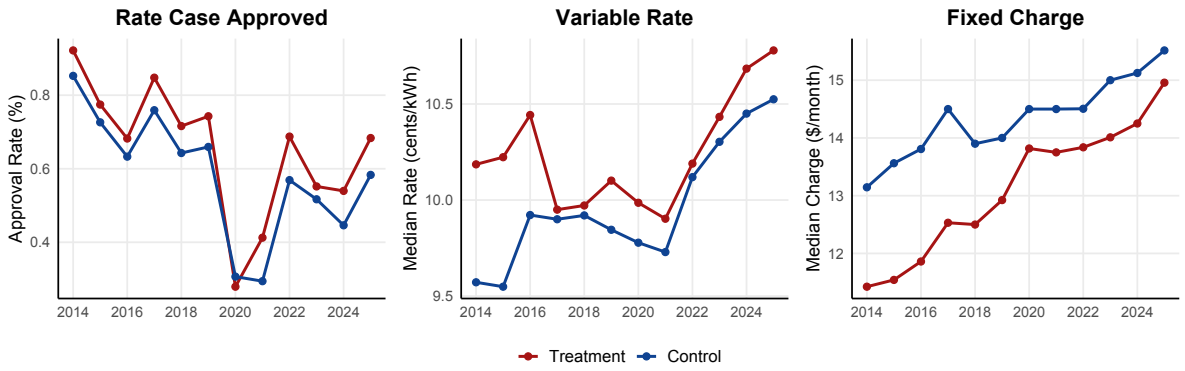

(a) Residential Sector Trends

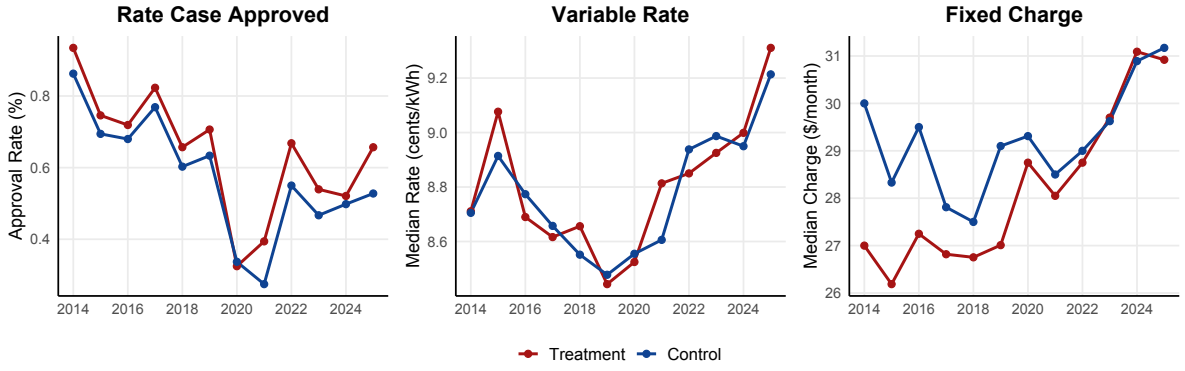

(b) Commercial Sector Trends

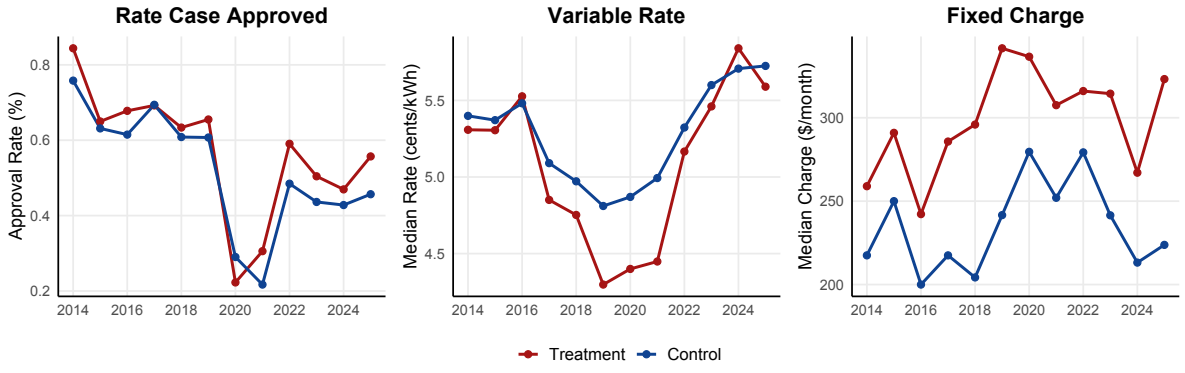

(c) Industrial Sector Trends

To address the challenge of spatial sorting, I introduce covariate-specific time trends by interacting year indicators with baseline county characteristic quintiles. For specifications estimating retail electricity market outcomes, I include 2014 population and GDP quintile-by-year FEs. This inclusion shifts the difference-in-differences requirement to a conditional parallel trends assumption. In the absence of data center load shocks, treatment and control counties must only observe parallel growth paths conditional on belonging to the same baseline demographic and economic output quintiles. This is a far more plausible proposition, as a treatment and

a control county that had similar population sizes and GDP in 2014 are highly likely to have faced the same macroeconomic trends over time. By absorbing time-varying shocks that are structurally correlated with a county's size and wealth at the beginning of the observation window, this augmented specification prevents divergent macroeconomic shifts from confounding the treatment effect. To confirm the empirical validity of this approach, Appendix A verifies the existence of sufficient common support between the treatment and control cohorts across all baseline covariate quintiles.

I impose an even stronger augmentation for specifications estimating local economic outcomes. Because I examine sector-specific causal effects of data centers, the specification must account for the fact that computing facilities systematically sort into regions with high baseline professional and technological intensity. Failure to control for these sector-specific growth paths would severely bias the estimators. Along with population and GDP quintile-by-year FEs, I also interact year indicators with four baseline sector intensity quintiles: professional and business services (includes high-tech), manufacturing, service, and agriculture. These interactions non-parametrically control for the distinct macroeconomic trends experienced by local economies with distinct sector compositions. These sector quintile-by-year FEs are omitted from the specifications for electricity market outcomes, as a county's baseline commercial and industrial intensity is intrinsically tied to how its local utility structures its retail products—a relationship absorbed by standard county FEs. Adding time-varying interactions for sector intensities in the electricity regressions risks absorbing the very utility pricing dynamics the paper aims to measure.

4.1.2 Leading Placebo Indicator

While the introduction of quintile-by-year FEs addresses the challenge of spatial sorting, a second challenge of reverse causality remains for specifications estimating local economic outcomes. Spatial sorting is driven by observable, pre-existing structural baseline traits that cause broad, secular trend divergence. Reverse causality is instead driven by unobserved, idiosyncratic shocks within the observation window that dictate the exact timing of data center development. For instance, facility site selection may be driven by municipal tax incentives and planned infrastructure upgrades that stimulate local economic outcomes in the absence of facility development. If treatment timing is positively correlated with such localized, unobserved economic catalysts, the TWFE estimator will suffer from upward bias. Conversely, facility site selection may be driven by de-industrializing localities that contain pre-existing industrial sites with no active owner.[11] The TWFE estimator will then falsely attribute declining manufacturing output and employment to the computing facility, thereby suffering from downward bias. Quintile-by-year FEs cannot absorb these unobserved shocks specific to a single county in a single year.

To explicitly test for this reverse causality in site selection and validate the timing of a causal mechanism, the macroeconomic specifications are further augmented with a 3-year leading placebo indicator. Recall from Section 3.2 that I define a treatment county as one that records at least 10 MW of operational demand load within the observation window. To create the placebo indicator, I identify the first year that each treatment county crosses the 10 MW threshold and assign the county-specific placebo indicator a value of 1 for the three years prior to this threshold year. The indicator is otherwise 0 for all control counties, treatment counties after crossing the 10 MW threshold, and treatment counties prior to the 3-year lead window. For "early adopter" counties like Loudoun that have crossed the 10 MW threshold before 2014, the placebo indicator takes the value of 0 across the entire panel.

If reverse causality is driving the results, the placebo indicator coefficient will capture anomalous pre-treatment growth or decline. Confirmation of a data center-induced macroeconomic shift therefore requires a precise null effect on this placebo indicator. I intentionally omit this

[11] xAI's first facility in Memphis is housed in a former Electrolux appliance manufacturing factory.

falsification test from the retail electricity specifications due to the rigid institutional frictions governing utility rate-setting across both regulated and deregulated regimes. Under the standard "used and useful" regulatory doctrine, the oversight of state PUCs prevents utilities from passing anticipatory capital expenditures onto ratepayers in advance of actual infrastructure deployment. Electricity price hikes are thus mechanically a reaction to realized grid strain, rendering pre-treatment trends in retail tariffs implausible by regulation.

4.2 TWFE Specification

I introduce two TWFE specification variants based on the discussion of the causal identification strategy in Section 4.1. The first type of specification estimates retail electricity market outcomes: (1) the probability of rate case approval, (2) the variable rate, and (3) the fixed charge across the residential, commercial, and industrial rate classes. While I estimate the probability of rate case approvals in levels, I log-transform the variable rate and fixed charge outcomes to normalize the risk-skewed distribution of price data. This log-level specification allows the estimated coefficients to be interpreted as semi-elasticities.

Let Y_{ist} be the class-specific retail electricity outcome for county i located in state s in year t. Let D_{it} be the operational data center demand load, expressed in units of 100 MW. Let Reg_s and Dereg_s be time-invariant state-level deregulation indicators. Let X_{it} be the vector of time-varying weather controls (annual HDD/CDD). To study spatial spillover effects of computing facilities at the Commuting Zone (CZ) and Balancing Authority (BA) levels, I define D_{it}^{CZ} and D_{it}^{BA} as the aggregate demand load of county i's CZ and BA, respectively, minus the demand load of county i. Given that CZ_i and BA_i are the sets of counties comprising the CZ and BA of county i, respectively, these aggregated "leave-one-out" demand loads are defined as:

$$D_{it}^{CZ} \equiv \sum_{j \in CZ_i} D_{jt} - D_{it},$$
$$D_{it}^{BA} \equiv \sum_{j \in BA_i} D_{jt} - D_{it},$$

Following the above notation, the retail electricity specification is as follows:

$$\begin{aligned} Y_{ist} = \beta_1(D_{it} \times \text{Reg}_s) + \beta_2(D_{it} \times \text{Dereg}_s) + \delta_1 D_{it}^{CZ} + \delta_2 D_{it}^{BA} \\ + \Gamma X_{it} + \alpha_i + \tau_t + \Phi_{q(i),t} + \varepsilon_{ist}, \end{aligned} \tag{1}$$

where α_i are county FEs, τ_t are year FEs, and $\Phi_{q(i),t}$ are the 2014 baseline population and GDP quintile-by-year FEs. By the logic of my identification strategy, the coefficients β_1 and β_2 capture the causal effect of each additional 100 MW of in-county demand load for regulated and deregulated counties, respectively. Similarly, δ_1 and δ_2 capture the causal effect of each incremental 100 MW of out-of-county demand load at the CZ and BA levels, respectively. I deliberately omit regulatory interactions with the spatial spillover terms, D_{it}^{CZ} and D_{it}^{BA}, as such a specification fundamentally mischaracterizes the institutional mechanism of regional cost transmission. Investments in regional transmission networks necessitated by localized data center load shocks are administered by the BA under federal regulation. Crucially, the capital expenditures for high-voltage grid maintenance and upgrades are socialized across the entire BA footprint. Consequently, a local utility will mechanically absorb its share of these regional costs, regardless of whether it operates in a regulated or deregulated state. Utilities then uniformly recover these exogenous, BA-imposed transmission costs by passing them down to their customers through a combination of fixed charges and TCAs. Imposing a retail regulatory interaction on the spillover terms would theoretically constrain a federally-regulated, wholesale cost-socialization mechanism.

To ensure valid inference, I cluster standard errors at the state level for two reasons. First,

counties within the same state are frequently served by the same utilities and are thus subject to uniform rate schedules. Second, the adjudicatory processes for utility regulations and rate case approvals fall under the unified jurisdiction of state-level PUCs. State-level clustering is therefore necessary to account for arbitrary spatial and serial correlation of residuals within states, which would otherwise lead to downward-biased standard errors and spurious findings of statistical significance.

The second type of TWFE specification estimates local economic outcomes: (1) aggregate and sectoral GDP, (2) aggregate and sectoral labor market outcomes, (3) household welfare metrics, and (4) housing value. I apply a natural log-transformation to all continuous economic dependent variables for three distinct purposes. First, monetary and demographic data are bounded at zero and severely right-skewed. Second, logging these outcomes stabilizes variance across a highly heterogeneous panel. Because absolute economic fluctuations scale with the size and wealth of counties, estimating in levels would allow the standard volatility of major metropolitan hubs to drown out meaningful causal shocks occurring in smaller counties. Finally, the log-level specification standardizes the interpretation of the treatment effect as a semi-elasticity across all outcome variables.

Let Y_{ijt} be the local economic outcome for county i located in CZ j at time t. Let P_{it} be the 3-year leading placebo indicator defined in Section 4.1. Following this notation, the local economic specification is as follows:

$$Y_{ijt} = \beta D_{it} + \phi P_{it} + \Gamma X_{it} + \alpha_i + \tau_t + \Phi_{q(i),t} + \Psi_{q(i),t} + \varepsilon_{ijt}. \tag{2}$$

The economic TWFE specification retains the county (α_i) and year (τ_t) FEs established in Equation 1, along with the population and GDP quintile-by-year FEs ($\Phi_{q(i),t}$). As discussed in Section 4.1, the second specification expands non-parametric time trends to include baseline sector intensities for professional and business services, manufacturing, service, and agriculture. These sector intensity quintile-by-year FEs are collectively denoted by $\Psi_{q(i),t}$. For the coefficient β to be interpreted as the causal effect of a 100 MW data center load shock on economic outcome Y_{ijt}, the specification must satisfy the falsification test: $\phi = 0$. A precise null effect on the 3-year leading placebo indicator indicates that economic shifts do not precede the treatment, thereby ruling out the introduction of bias from reverse causality.

Because labor markets, wage spillovers, and residential sorting operate fluidly across county lines, I cluster standard errors for all economic specifications at the CZ level.[12] CZs are geographical clusters of counties delineated by the U.S. Department of Agriculture's Economic Research Service to represent contiguous local economies and labor markets. They therefore serve as the ideal level of clustering to account for arbitrary spatial and serial correlation in the residuals.

4.3 Empirical Results: Retail Electricity Markets

Results from the retail electricity specifications reveal striking patterns of households bearing the burden of data center demand loads. Table 4 indicates a sharp contrast in retail cost pass-through to residential customers between regulated and deregulated states. For each 100 MW increment in host-county load, the probability of a residential rate case approval filed by utilities in regulated states increases by 0.65 percentage points. Given the unconditional sample mean of 0.588, this point estimate translates to approximately a 1.1% relative increase in approval probability.[13] Columns 2 and 3 indicate that regulated utilities are securing residential rate case approvals that translate into higher monthly fixed charges for households: fixed charges

[12] I use 2020 Commuting Zone delineations to most accurately represent the boundaries of local labor economies during the observation window.

[13] The calculation $0.65/0.588 = 0.011$ represents a conservative lower bound for the relative effect, as the pooled sample mean of 0.588 is mechanically inflated by treated observations.

Table 4: Effect of Data Center Load on Residential Electricity Rates

	Dependent Variable		
	(1) Rate Case App.	(2) log(Var Rate)	(3) log(Fixed Charge)
Demand (100MW) × Regulated	0.0065**	0.0026	0.0074***
	(0.0027)	(0.0017)	(0.0023)
Demand (100MW) × Deregulated	-0.0124	-0.0079	-0.0005
	(0.0082)	(0.0063)	(0.0109)
CZ Demand Spillover (100MW)	0.0036	-0.0020	0.0021
	(0.0035)	(0.0017)	(0.0021)
BA Demand Spillover (100MW)	0.0001	-0.0001	0.0006**
	(0.0006)	(0.0003)	(0.0003)
County & Year FEs	Yes	Yes	Yes
Quintile-by-Year FEs	Yes	Yes	Yes
Weather Controls	Yes	Yes	Yes
Observations	35,886	35,876	35,448
Dependent Var. Mean (Unlogged)	0.588	10.04	13.83
R^2	0.425	0.779	0.851
Within R^2	0.005	0.008	0.005

Notes: Standard errors clustered by state are reported in parentheses. Significance levels are denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The variable rate is expressed in cents per kWh, and the fixed charge in dollars per month. Quintile-by-Year FEs encompass GDP quintile-by-year and population quintile-by-year fixed effects. County quintiles are constructed using 2014 as the base year.

are raised by 0.74% for each 100 MW load shock, while the change in variable rates remains statistically indistinguishable from zero. By taking the ratio of these two point estimates, I can approximate the relative magnitude of this pass-through: for every 1 percentage point increase in the probability of a rate case approval, residential fixed charges in regulated states increase by approximately 1.13%. Given the unconditional sample fixed charge mean of \$13.83/month, this 1.13% increase approximately translates to an additional \$0.16/month per 100 MW.[14] As discussed in Section 3.1, the estimated effect of computing infrastructure demand loads on local residential electricity costs in regulated states is a lower bound. The Utility Rate Database (URDB) omits the PCAs and TCAs that are mechanically sensitive to the exact grid strains introduced by data centers. When accounting for these direct pass-through mechanisms that are missing from rate case data, it is highly plausible that facility load shocks impose an aggregate financial burden on residential ratepayers that substantially exceeds the base tariff adjustments captured in this specification.

Unlike their regulated counterparts, utilities in deregulated states appear to be non-responsive to data center load shocks. In deregulated regimes, privately owned generators bear the operating expenditures for increasing supply to meet demand and the capital expenditures to construct new power plants. Meanwhile, retail energy providers (REPs) absorb the wholesale price volatility passed on by generators. Consequently, local utilities face no exposure to generation-related costs, removing a major catalyst for filing a rate case. However, utilities in deregulated markets still own the transmission infrastructure exposed to load shocks. Due to massive capacity needs,

[14]The approximation of relative magnitude in regulated pass-through scales the reduced-form effect on fixed charges by the first-stage effect on rate case approval: $(0.00740/0.00654) \times 0.01 = 0.0113$. As discussed in footnote 13, the calculation $13.83 \times 0.0113 \approx 0.16$

Table 5: Effect of Data Center Load on Commercial Electricity Rates

	Dependent Variable		
	(1) Rate Case App.	(2) log(Var Rate)	(3) log(Fixed Charge)
Demand (100MW) × Regulated	0.0118	-0.0025	-0.0073
	(0.0117)	(0.0067)	(0.0111)
Demand (100MW) × Deregulated	0.0041	0.0029**	0.0029
	(0.0027)	(0.0014)	(0.0031)
CZ Demand Spillover (100MW)	0.0016	-0.0027	-0.0036
	(0.0034)	(0.0029)	(0.0028)
BA Demand Spillover (100MW)	0.0003	-0.0005	0.0005
	(0.0005)	(0.0003)	(0.0006)
County & Year FEs	Yes	Yes	Yes
Quintile-by-Year FEs	Yes	Yes	Yes
Weather Controls	Yes	Yes	Yes
Observations	35,957	35,903	35,809
Dependent Var. Mean (Unlogged)	0.580	8.52	34.01
R^2	0.402	0.694	0.793
Within R^2	0.010	0.008	0.004

Notes: Standard errors clustered by state are reported in parentheses. Significance levels are denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The variable rate is expressed in cents per kWh, and the fixed charge in dollars per month. Quintile-by-Year FEs encompass GDP quintile-by-year and population quintile-by-year fixed effects. County quintiles are constructed using 2014 as the base year.

large-scale computing facilities often require interconnection at the high-voltage bulk transmission level, building their own substations for conversion to low-voltage. Because transmission costs are partially passed through TCAs, utilities in deregulated states have less incentive than those in regulated states (who are exposed to both generation and transmission risk) to open a prolonged rate case with state PUCs. This likely explains the null effect in rate case approval and fixed charges observed in the URDB data. Furthermore, the observed divergence implies that the immense capital requirements of generation capacity expansion serve as the primary driver for data center-driven formal rate cases in regulated states.

Because utility rate case data also omit the changes in competitive pricing among the generation-exposed REPs, the null coefficients estimated for deregulated localities cannot be interpreted as evidence that deregulated households are immune to the costs of data center development. Rather, the bifurcated structure of deregulated markets between competitive REPs and PUC-regulated local utilities fragments the cost-recovery process in a manner that evades formal rate cases.

Alongside in-county demand load effects, Table 4 also shows that a 100 MW shock in out-of-county demand load within the Balancing Authority (BA) footprint drives a 0.06% increase in local residential fixed charges. As discussed in Section 4.2, high-voltage transmission network upgrades necessitated by massive load shocks are facilitated by federally regulated BAs. This empirical estimate captures the capital expenditures of transmission investment that are mechanically socialized across all utilities within a BA, regardless of state regulatory regime. The statistical significance of the fixed charge hike and null result in rate case approval with respect to BA demand spillover indicate that these socialized transmission costs are substantial enough that direct pass-through TCAs cannot fully absorb the socialized costs, yet too diffuse to trigger

Table 6: Effect of Data Center Load on Industrial Electricity Rates

	Dependent Variable		
	(1) Rate Case App.	(2) log(Var Rate)	(3) log(Fixed Charge)
Demand (100MW) × Regulated	0.0151	0.0060	0.0074
	(0.0119)	(0.0067)	(0.0128)
Demand (100MW) × Deregulated	-0.0017	-0.0104*	0.0239
	(0.0026)	(0.0053)	(0.0183)
CZ Demand Spillover (100MW)	0.0036	-0.0016	0.0093
	(0.0041)	(0.0022)	(0.0059)
BA Demand Spillover (100MW)	-0.0008	0.0004	-0.0004
	(0.0005)	(0.0006)	(0.0008)
County & Year FEs	Yes	Yes	Yes
Quintile-by-Year FEs	Yes	Yes	Yes
Weather Controls	Yes	Yes	Yes
Observations	34,285	34,077	33,428
Dependent Var. Mean (Unlogged)	0.524	4.91	274.87
R^2	0.429	0.791	0.854
Within R^2	0.005	0.003	0.002

Notes: Standard errors clustered by state are reported in parentheses. Significance levels are denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The variable rate is expressed in cents per kWh, and the fixed charge in dollars per month. Quintile-by-Year FEs encompass GDP quintile-by-year and population quintile-by-year fixed effects. County quintiles are constructed using 2014 as the base year.

the legal and administrative friction of accelerated rate case filings.

Interestingly, the point estimate for Commuting Zone (CZ) spillovers (+0.21%) is substantially larger in magnitude than its BA counterpart (+0.06%) yet remains statistically indistinguishable from zero. This empirical divergence likely reflects the institutional perimeters of the power grid. Because CZs denote labor market geographies rather than administrative grid boundaries, the cost-sharing mechanisms of CZ-level transmission are highly arbitrary and idiosyncratic. For instance, a CZ may be covered by a single utility, may cover multiple utility service territories, or even cover multiple BA regions. Consequently, the CZ estimate for residential fixed charges suffers from extreme variance and yields a null result.

Tables 5 and 6 show the retail electricity specification results for the commercial and industrial (C&I) rate classes. Under regulated environments, data center demand loads yield null effects across all dependent variables for C&I. This result stands in stark contrast with the +0.74% fixed charge effect for residential ratepayers, indicating that utilities in regulated localities strategically insulate C&I consumers from massive generation and transmission capital expenditure shocks. This behavior may be driven by large C&I entities possessing lobbying power in rate case processes and the ability to relocate operations, resulting in a highly elastic long-term energy demand. Utilities therefore shift fixed cost-recovery burdens onto their captive, highly inelastic residential ratepayer base.

The cost burden in deregulated states appears to also include commercial entities (along with households through TCAs, as discussed previously). Table 5 reveals a statistically significant, albeit small, 0.29% increase in the commercial variable rate driven by each incremental 100 MW addition in data center load. Recall that URDB data in deregulated states strictly reflects the local utility's delivery tariff driven by transmission and distribution costs, not the retail

energy provider's (REP) generation-exposed price. This statistically significant result implies that massive data center loads inevitably create localized strain on utility networks. Rather than accelerating rate cases for any of their rate classes, utilities in deregulated localities appear to recover the costs of this localized strain volumetrically through variable delivery rates rather than through fixed charges.

Most remarkably, Table 6 appears to imply that a localized 100 MW load shock leads to a marginally significant 1.04% decrease ($p = 0.054$) in the deregulated industrial variable rate. This variable rate drop likely reflects volume-discount delivery tariffs negotiated directly between large-scale computing facilities and local utilities.[15] These negotiated utility discounts mechanically drag down the county's median industrial variable rate. The negative result may raise endogeneity concerns: rather than revealing discount agreements, the negative coefficient captures the dynamic of data centers sorting into counties where industrial power is becoming cheaper independent of facility presence. This critique can be addressed by two observations. First, the 2014 baseline population and GDP quintile-by-year FEs explicitly absorb the secular spatial sorting of facilities into structurally cheaper or de-industrializing regions. Second, the presence of reverse causality should result in the decline of industrial variable rates prior to treatment across both regulatory regimes. Table 6 shows no such negative relationship in regulated counties, thereby suggesting the presence of an institutional mechanism particular to deregulated jurisdictions (e.g., utility tariff negotiation) rather than a violation of parallel trends.

Similar to residential customers, C&I customers are subject to direct pass-through PCAs and TCAs that do not appear in the URDB. Therefore, no definitive claim can be asserted regarding the aggregate financial incidence of data center expansion across the three rate classes. However, the strategic direction of utility cost allocation, as revealed by base tariff adjustments, remains unambiguous. Utilities in regulated localities deliberately structure their base rate cases to shield commercial and industrial entities from the immense capital expenditures required to accommodate these new demand loads. In deregulated markets, utilities are actively creating specialized rate classes to offer bespoke discounts and secure the revenue generated by massive data center loads.

Unlike energy-intensive C&I entities—which exhibit highly elastic long-run energy demand due to credible flight risk—households are geographically sticky and exhibit highly inelastic energy demand. The empirical results demonstrate that this captive ratepayer base systematically faces the burden of utility cost-recovery in regulated communities. While the precise magnitude of total retail pass-through is partially obscured by direct cost adjustments, the deliberate asymmetry in base tariff effects confirms that the financial incidence of data center development on the grid falls disproportionately upon local households.

4.4 Empirical Results: Local Economic Outcomes

Results from the economic specifications indicate that data centers have localized effects on host economies that are more nuanced than the conventional promise of unconstrained economic revitalization. As described in Sections 4.1 and 4.2, I impose a falsification test through the 3-year leading placebo P_{it} with regression coefficient ϕ to address the challenge of reverse causality. This augmentation to the TWFE framework proved highly necessary, as the falsification test failed five out of the thirty-one economic variables used: labor force, median household income (HHI), manufacturing GDP, service employment, and service wages (see Appendix B for complete regression tables of all economic specifications). These variables are therefore excluded from causal interpretation.

[15] These negotiations are distinct from Power Purchase Agreements (PPAs), where large C&I entities like hyperscalers and AI data centers may directly purchase electricity wholesale from generators at a discounted rate. The URDB database only captures tariffs between retail customers and utilities, not wholesale transactions.

Figure 4: Local Aggregate Economic Outcomes

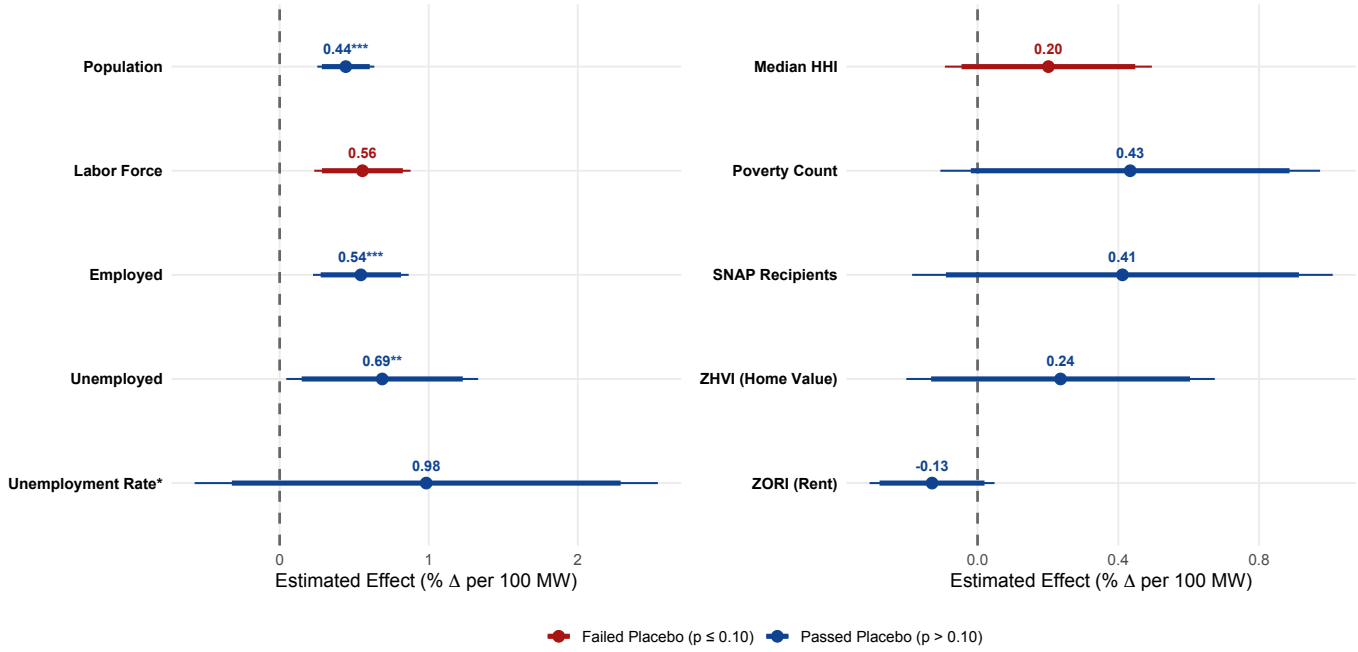

Figure 5: Local Sectoral Economic Outcomes

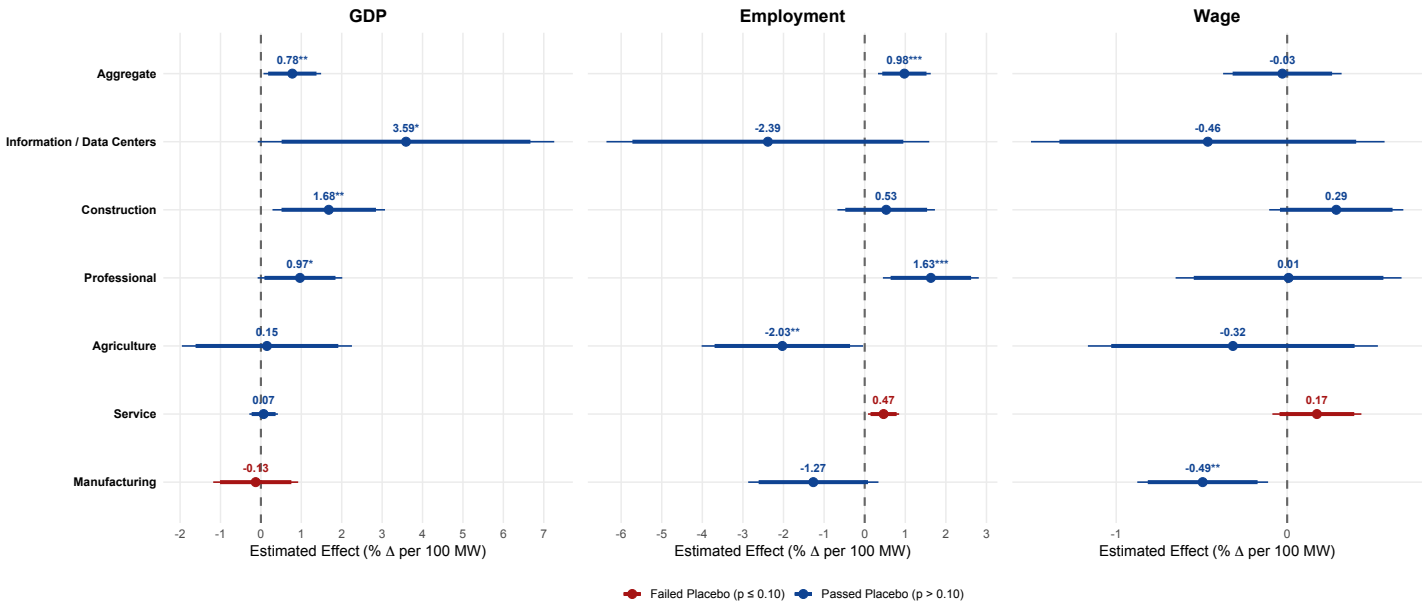

Note: *The estimated effect on the unemployment rate must be interpreted as percentage point change per 100 MW, not relative percentage. "Employed" in Figure 4 is sourced from the Bureau of Labor of Statistics' Local Area Unemployment Statistics, whereas "Aggregate Employment" in Figure 5 is calculated from the U.S. Census Bureau's County Business Patterns.

Figure 4 captures the labor market and household welfare outcome impact of data centers, whereas Figure 5 shows the impact on sector-level output, employment, and wages. Each additional 100 MW in data center load drives up aggregate county GDP by 0.78%. This output growth is highly concentrated in information (+3.59%), construction (+1.68%), and professional and business services (0.97%)—the three sectors with direct exposure to data center development

and operations. The arrival of large-scale computing facilities appears to deliver a permanent employment boom, as a 100 MW marginal increase in host-county load raises county population by 0.44% and the number of employed residents by 0.54%. These permanent jobs appear highly concentrated in professional and business services, the only sector to have a positive employment estimate (+1.63%). The effect on employment in construction and information yields null results, likely due to (1) the highly transient nature of data center construction, and (2) the small number of dedicated technicians required to operate a data center.

While the rise in county population and employment count lends credence to the prospect of job creation promised by data center operators, the null effects in the unemployment rate, aggregate wage, output in non-related sectors, poverty count, SNAP recipients, and housing value confirm that the capital and labor influx from data centers fails to alter broader community welfare. In fact, two economic outcomes indicate distinct patterns of crowding-out: the (1) rising unemployment count and (2) falling manufacturing wages. Although these coefficient estimates could accommodate a multitude of competing mechanistic interpretations, I propose a distinct, institutionally grounded channel for each.

First, the simultaneous rise in both employed (+0.54%) and unemployed (+0.69%) resident counts highlights a structural labor market friction. While the facility's arrival induces permanent population in-migration concentrated in professional and business services, this influx of job-seekers outpaces the long-term job absorption capacity of the host community. "Professional and business services" (NAICS 54–56) includes highly specialized roles in technology, but the sector also encompasses the extensive non-specialized support ecosystem of a data center: physical security, facilities maintenance, and administrative staffing. The initial development of computing infrastructure may draw in regional, non-specialized labor seeking such roles. However, the steady-state operation of a data center is exceptionally lean (see Section 2.2), and the local labor market cannot sustain this inflated support sector indefinitely. Unlike highly compensated, specialized talent with national mobility, these non-specialized migrants are often geographically sticky. Consequently, the saturation of the local support market results in surplus labor stranded in the host municipality.

Second, the decline in manufacturing wages (−0.49%) likely represents the labor-side incidence of the electricity market dynamics established in Section 4.3. Because manufacturing is uniquely energy-intensive among the studied sectors, electricity cost hikes (either from base tariffs or direct cost adjustments) induced by localized grid strain severely constrain manufacturing profit margins. The empirical evidence suggests that these elevated operational costs are ultimately passed down to the local workforce in the form of reduced wages. The precise theoretical mechanisms that drive this empirical result are explored in the structural model introduced in Section 5.

Taken all together, my reduced-form analysis rejects the narrative of a broad economic windfall. The colossal investment in computing infrastructure creates an isolated technology and construction boom with permanent employment gains in professional and business services. However, this expansion is shown to create surplus labor in the data center support ecosystem and actively penalize energy-intensive wages, leaving the net welfare of the host community statistically unchanged. Although this framework provides a cohesive explanation for the observed results, an empirical puzzle remains: why do manufacturing output and employment remain unaffected while the energy-intensive sector's wage erodes? To solve this labor market puzzle, I turn to a structural model in Section 5.

5 Structural Model

5.1 DSGE Set-Up

I introduce a closed-economy, two-region Dynamic Stochastic General Equilibrium (DSGE) framework to formalize the macroeconomic incidence of large-scale computing infrastructure. Region 1 is defined as the host locality subject to exogenous data center demand shocks, while region 2 represents the aggregate non-host remainder of Region 1's Balancing Authority (BA) footprint. The two regions are structurally integrated through a shared administrative grid perimeter and draw power from a unified regional supply of generators. Each region i is populated by three entities: (1) a shared electricity sector, (2) a continuum of identical households of measure N_i, which I model as a single representative household, and (3) a perfectly competitive manufacturing firm. As defined, region 1 also exclusively hosts an exogenous data center sector.

Although the empirical evidence demonstrates that large-scale computing facilities independently shock aggregate output and local labor markets, this framework intentionally abstracts from the channels of direct job creation and agglomeration. Instead, the model is strictly disciplined to isolate the retail electricity mechanism, providing a targeted theoretical resolution to the manufacturing sector puzzle established in Section 4.4.

To isolate the impact of the regulatory regime on macroeconomic incidence, I parametrize and solve two distinct electricity sector designs representing regulated and deregulated market structures. I formalize this institutional divergence by characterizing the electricity sector variants via a binary indicator $\Omega = \{0, 1\}$, where $\Omega = 0$ represents the regulated architecture and $\Omega = 1$ represents the deregulated architecture. I abstract from regulatory heterogeneity within BAs and assume that both regions have the same electricity market structure. This simplification preserves analytical tractability and strictly isolates the cost-recovery mechanisms of each respective regime. The modeling choice also ensures that the analysis is not confounded by asymmetric tariff structures or regulatory arbitrage within the unified grid.

5.1.1 Data Center Sector

Because this paper is primarily concerned with the macroeconomic incidence of localized data center load shocks and not the spatial sorting decisions of facility operators, I model the data center (DC) sector as an exogenous shock to region 1's total demand load. Borrowing notation from my reduced-form work, I define D_t as the aggregated maximum power demand (in 100 MW) of computing facilities in region 1 at year t. The annual energy demand corresponding to the maximum power D_t is as follows:

$$E_t^{DC} \equiv D_t \cdot \delta \cdot (8760 \text{ hours/year}),$$

where $\delta \in [0, 1]$ is the aggregate load factor (utilization) of the locality's data centers.

5.1.2 Regulated Electricity Sector: $\Omega = 0$

The regulated electricity sector faces variable and capital costs that it recovers through different pass-through mechanisms. Under the regulated market architecture, it is analytically convenient to think of the shared electricity sector as each region housing its own vertically integrated utility. Let E_t be the aggregate energy demand at year t across both regions:

$$E_t \equiv E_{1t}^H + E_{2t}^H + E_{1t}^M + E_{2t}^M + E_t^{DC},$$

where E_{it}^H represents the demand of the representative household in region i (aggregated across the continuum) and E_{it}^M represents the demand of the manufacturing firm in region i. Following Borenstein et al. (2000) and Borenstein and Holland (2003), the sector faces convex variable

costs:[16]

$$V(E_t) \equiv v_0 E_t + \frac{v_1}{2} E_t^2.$$

Note that the variable cost of generation is symmetric across regions due to the assumption of a unified generation supply. Because utilities in regulated regimes operate under frameworks that authorize the automatic reconciliation of fluctuating generation expenses outside of rate cases, I assume that the sector's variable costs are directly passed to households and manufacturing firms through the volumetric rate via the Power Cost Adjustment (PCA). The variable rate of electricity for households p_t^H and firms p_t^M across both regions is thus the average variable cost:

$$p_t^H = p_t^M \equiv \frac{V(E_t)}{E_t} = v_0 + \frac{v_1}{2} E_t.$$

The exogenous shock of data center power demand D_t requires the electricity sector to incur significant capital expenditure in generation, transmission, and distribution infrastructure. Let Γ_t be the net uncovered capital cost of data center support, scaling linearly with D_t:

$$\Gamma_t(D_t) \equiv \Psi D_t,$$

where Ψ is the net uncovered capital cost per 100 MW of peak aggregate data center power demand. As demonstrated empirically in Section 4.3, these uncovered capital expenditures are partially socialized throughout the entire BA region:

$$\Gamma_t = \omega \Gamma_t + (1 - \omega)\Gamma_t,$$

where $\omega \in [0, 1]$ is the share of cost localization, $\omega\Gamma_t$ is the capital cost that is recovered solely in region 1, and $(1 - \omega)\Gamma_t$ is the capital cost that is socialized across both regions. As Section 4.3 also shows that utilities actively shield highly elastic C&I entities from capacity-driven rate hikes, I make the modeling assumption that the full incidence of infrastructure capital costs falls on the households in the form of fixed charges. Therefore, manufacturing firms face a time-invariant fixed tariff $\bar{F}^M$ that strictly reflects the legacy costs established prior to the exogenous demand shock.

Following the structure of fixed charges levied by utilities, let the individual household fixed charge F_{it}^H in region i at year t consist of a base fixed tariff and a transmission cost adjustment (TCA):

$$F_{it}^H \equiv F_{it}^{\text{Base}} + F_{it}^{\text{TCA}},$$

where the TCA component represents the BA-facilitated investment into shared regional transmission infrastructure. Therefore, regulated utilities flexibly adjust F_{it}^{TCA} to cover the socialized segment of net uncovered capital cost:

$$F_{it}^{\text{TCA}} = \frac{(1 - \omega)\Psi D_t}{N_1 + N_2}, \quad i \in \{1, 2\}.$$

Unlike the TCA, the base fixed charge F_{it}^{Base} is regulated through the PUC-mediated rate case process. Each region's utility has a target base fixed rate it would like to charge to cover localized capital costs:

$$F_{1t}^{\text{Base}*} \equiv \frac{\omega \Psi D_t}{N_1}, \quad F_{2t}^{\text{Base}*} \equiv 0,$$

[16] This functional form is often used to represent the increasing marginal cost of dispatching less efficient peaker plants as energy demand rises.

where $F_{it}^{\text{Base}} = 0$ is assumed at $t = 0$ for both regions. The utility in the host region thus faces a local cost gap Δ_t that is dictated by how far the regulated base fixed charge in the previous year $F_{1,t-1}^{\text{Base}}$ is from the target $F_{1t}^{\text{Base}*}$:

$$\Delta_t \equiv \omega \Psi D_t - N_1 F_{1,t-1}^{\text{Base}}.$$

As the local cost gap grows, the utility in region 1 (1) has a greater incentive to file a rate case with the state PUC to raise its base fixed tariff, and (2) has a greater probability of succeeding in the regulatory process. I therefore model the regulatory approval of a rate case as a stochastic hazard process. Let ρ_t be the probability that the PUC approves the host utility's rate case in year t, modeled as a logistic function strictly increasing in local cost gap Δ_t:

$$\rho_t \equiv \frac{1}{1 + \exp[-(\gamma_0 + \gamma_1 \Delta_t)]},$$

where γ_0 captures the unconditional baseline probability of a rate case occurring, and $\gamma_1 > 0$ dictates the combined utility and regulatory sensitivity to the utility's unrecovered capital burden.

5.1.3 Deregulated Electricity Sector: $\Omega = 1$

Under the deregulated market structure, the generation, retail, and delivery (transmission and distribution) functions of the electricity sector are structurally unbundled. To formalize this institutional environment, it is analytically convenient to abstract the fragmented sector into three distinct functional entities within the model's two-region framework: perfectly competitive generators operating across the shared BA footprint to form a unified supply stack, a representative Retail Energy Provider (REP) in each region engaged in perfect competition exclusively within its region, and a local utility in each region that maintains the physical grid.

The unbundling of generation assets alters the recovery mechanism of infrastructure costs. Unlike the vertically integrated utility, perfectly competitive generators do not recover capital expenditures on power plant infrastructure through PUC-regulated fixed charges and TCAs. Rather, they recoup both their variable operating costs and infrastructure investments exclusively through the centralized wholesale energy market. Under perfect competition, generators bid their marginal cost, yielding a base system-wide energy price shared across both regions at year t:

$$p_t^{\text{energy}} \equiv V'(E_t) = v_0 + v_1 E_t.$$

In the absence of significant and rapid transmission investment, the exogenous data center shock D_t induces localized grid congestion in the host region.[17] To capture this spatial friction, I define the wholesale market price as the Locational Marginal Price (LMP_{it}), which adds a region-specific nodal congestion premium p_t^{nodal} to the base energy price. Let the nodal premium scale linearly with the localized computing load in region 1:

$$p_t^{\text{nodal}} \equiv \kappa E_t^{DC},$$

where $\kappa > 0$ is the sensitivity parameter capturing the severity of localized grid congestion. The

[17] Congestion occurs when the electricity grid lacks sufficient transmission capacity to meet power demand. The phenomenon occurs when the transmission infrastructure delivering electricity is the binding constraint to meeting power demand, not the lack of supply from generators. Upgrading transmission infrastructure often takes more than seven years, meaning that the implicit modeling assumption of data centers inducing localized grid congestion is a highly defensible one.

LMP of each region can thus be defined as:

$$LMP_{it} \equiv \begin{cases} v_0 + v_1 E_t + \kappa E_t^{DC}, & i = 1 \\ v_0 + v_1 E_t, & i = 2. \end{cases}$$

Serving as the intermediary between the wholesale market and retail customers, region-specific REPs purchase electricity at the prevailing LMP_{it} and retail it to households. Under perfect competition, the REP operates under a zero-expected profit condition. To shield households from the real-time volatility of the wholesale market, REPs offer fixed-rate contracts spanning k years.[18] To satisfy the zero-expected profit condition, the price of a newly issued contract p_{it}^* in region i at year t must be the unweighted average of expected future wholesale LMPs over the contract duration:

$$p_{it}^* = \frac{1}{k} \sum_{s=0}^{k-1} \mathbb{E}_t[LMP_{i,t+s}].$$

To preserve the representative agent framework, I assume perfect risk-pooling across the continuum of households. Equivalently, the representative household manages a continuously overlapping portfolio of retail contracts. Assuming a uniform distribution of contract renewals, a $1/k$ fraction of this portfolio is renegotiated at the new price p_{it}^* in any given year t. The aggregated effective retail price p_{it}^H faced by the representative household is the average of all active, overlapping contracts:

$$p_{it}^H = \frac{1}{k} \sum_{m=0}^{k-1} p_{i,t-m}^*.$$

Conversely, large-scale manufacturing firms possess sophisticated energy procurement capabilities and higher risk tolerances compared to households. Therefore, I assume that the firms are directly exposed as buyers to the volatile wholesale market:

$$p_{it}^M = LMP_{it}.$$

Finally, the deregulated market regime leaves the two region-specific utilities strictly responsible for maintaining their respective physical transmission and distribution networks. While FERC-jurisdictional transmission upgrades continue to be socialized across the BA, these wires-only expenditures likely represent a minor fraction of the total capital cost required to support data center expansion (as discussed in Section 4.3). Because this socialized transmission component is a second-order effect compared to the incidence of wholesale energy volatility, I abstract from TCAs in the deregulated model to preserve analytical focus.

Consequently, the local utility simply levies a time-invariant fixed delivery charge $\bar{F}$ on both households and manufacturing firms across both regions:

$$F_{it}^H = F_{it}^M = \bar{F}.$$

By neutralizing the fixed-charge mechanism, this framework isolates REP and manufacturing firm exposure to the wholesale market as the primary vector of macroeconomic incidence in deregulated markets.

[18] In the real world, many REPs also offer variable-rate contracts. However, fixed-rate plans are the most common retail product and offer greater model tractability.

5.1.4 Representative Household

Each region i is populated by a continuum of identical households of measure N_i, which I model as an infinitely-lived representative household that maximizes lifetime utility. I assume that the utility of the representative household in region i adheres to an additively separable functional form with King-Plosser-Rebelo (KPR) preferences (King et al., 1988):

$$\max_{C_{it},E_{it}^H,L_{it}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\ln(C_{it}) + \nu \cdot \frac{(E_{it}^H)^{1-\sigma_E}}{1-\sigma_E} - \chi \cdot \frac{L_{it}^{1+\varphi}}{1+\varphi} \right]$$

where β is the discount factor, C_{it} the consumption of the composite good, E_{it}^H the household energy consumption with weight ν and elasticity parameter σ_E, and L_{it} the labor supply with weight χ and inverse Frisch elasticity φ.

The household optimizes the above utility function subject to the flow budget constraint:

$$C_{it} + p_{it}^H E_{it}^H + F_{it}^H + \frac{B_{i,t+1}}{1+r_t} \leq w_{it} L_{it} + B_{it} + \Pi_{it},$$

where p_{it}^H and F_{it}^H are the variable rate and fixed charge faced by the household, respectively, w_{it} is the region-specific real wage, B_{it} denotes the household's holdings of a one-period, risk-free bond traded in a national credit market yielding an exogenous real interest rate r_t, and Π_{it} represents the lump-sum transfer of firm profit from the perfectly competitive manufacturing firm in region i.

The household is also subject to a standard No-Ponzi condition to preclude infinite borrowing:

$$\lim_{t\to\infty} \mathbb{E}_0 \left[\left(\prod_{s=0}^{t} \frac{1}{1+r_s} \right) B_{i,t+1} \right] \geq 0.$$

To solve the representative household's problem, I formulate the standard Lagrangian and derive the first-order conditions with respect to consumption, energy demand, and labor supply (full derivations are relegated to Appendix C). Imposing the transversality condition, optimality yields the following conditions.

First, the intertemporal Euler equation dictates the household's consumption smoothing behavior over time:

$$\frac{1}{C_{it}} = \beta(1+r_t)\mathbb{E}_t \left[\frac{1}{C_{i,t+1}} \right].$$

Second, the intratemporal labor supply condition equates the marginal disutility of working to the marginal utility of the real wage:

$$\chi L_{it}^{\varphi} = \frac{w_{it}}{C_{it}}.$$

Finally, the intratemporal energy demand condition equates the marginal utility of household energy consumption to its relative price in terms of foregone composite consumption:

$$\nu (E_{it}^H)^{-\sigma_E} = \frac{p_{it}^H}{C_{it}}.$$

Together with the budget constraint and the transversality condition, these equations fully characterize the optimal behavior of the representative household in region i given the prevailing wage and residential electricity price.

5.1.5 Manufacturing Firm

Each region i is populated by a representative manufacturing firm that produces a homogeneous composite consumption good Y_{it} and engages in perfect competition on the national market. I assume that the firm operates a Constant Elasticity of Substitution (CES) production technology over energy and labor:

$$Y_{it} = A \left[\alpha (E_{it}^M)^{\frac{\sigma-1}{\sigma}} + (1-\alpha) L_{it}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}},$$

where A represents a constant total factor productivity, $\alpha \in (0,1)$ is the distribution parameter governing the factor share of energy, and $\sigma > 0$ is the elasticity of substitution between manufacturing energy demand E_{it}^M and labor L_{it}. I choose to abstract from capital accumulation and effectively treat the manufacturing sector as operating with a fixed stock of specialized capital. This modeling choice ensures that the analysis strictly isolates the labor-energy substitution channel, which represents the primary challenge for local firms facing short-to-medium-run energy price volatility.

The firm maximizes static period-by-period profit Π_{it}, taking the price of the composite good (normalized to 1 as the numeraire), the region-specific real wage w_{it}, and the variable industrial electricity price p_{it}^M as given. The firm's optimization problem is therefore:

$$\max_{L_{it}, E_{it}^M} \Pi_{it} = Y_{it} - w_{it} L_{it} - p_{it}^M E_{it}^M - F_{it}^M.$$

To solve the representative firm's problem, I derive the first-order conditions with respect to labor and energy demand (full derivations are relegated to Appendix D). Optimality yields the standard marginal product conditions.

First, the intratemporal labor demand condition equates the marginal product of labor to the region-specific real wage:

$$(1-\alpha) A^{\frac{\sigma-1}{\sigma}} \left(\frac{Y_{it}}{L_{it}} \right)^{\frac{1}{\sigma}} = w_{it}.$$

Second, the intratemporal energy demand condition equates the marginal product of energy to the real variable electricity price faced by the firm:

$$\alpha A^{\frac{\sigma-1}{\sigma}} \left(\frac{Y_{it}}{E_{it}^M} \right)^{\frac{1}{\sigma}} = p_{it}^M.$$

Because the manufacturing sector is perfectly competitive and operates under constant returns to scale, the firm earns zero economic profit in equilibrium ($\Pi_{it} = 0$).[19] Consequently, the price of the composite good must equal the firm's marginal cost of production. This equality yields the region-specific factor price frontier:

$$1 = \frac{1}{A} \left[(1-\alpha)^\sigma w_{it}^{1-\sigma} + \alpha^\sigma (p_{it}^M)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}.$$

The curvature of the factor price frontier is determined by the elasticity of substitution σ. As $\sigma \to 0$ and approaches Leontief technology, the firm's ability to substitute energy vanishes, forcing the local real wage to bear the full burden of any wholesale price shocks to maintain the zero-profit condition under fixed output.

[19] While $\Pi_{it} = 0$ in equilibrium, it enters the household budget constraint as a potential lump-sum transfer to ensure accounting closure of the regional economy.

5.2 Model Mechanics

With the two-region DSGE model formally set up, this section explores the comparative statics of an exogenous data center demand shock D_t through the introduction and interpretation of six propositions. These propositions on the model's mechanics validate the empirical findings from Section 4 and offer a coherent solution to the puzzle posed at the end of the section: in response to data center demand shocks, why do manufacturing wages decline while the sector's output and employment remain unaffected? The propositions also formalize the macroeconomic incidence of computing infrastructure shocks, demonstrating a clear burden on households in host communities. Appendix E features the complete proof of each proposition.

Proposition 1: An exogenous increase in localized data center power demand D_t strictly increases the variable rate of electricity p_{1t} for region 1's representative household and manufacturing firm, regardless of the governing regulatory architecture.

I evaluate the partial derivatives of the region-specific variable electricity price equations with respect to exogenous data center energy demand. Across both the regulated ($\Omega = 0$) and deregulated ($\Omega = 1$) market architectures, the derivatives are strictly positive:

$$\frac{\partial p_{1t}}{\partial E_t^{DC}} = \frac{v_1}{2} > 0 \quad \text{and} \quad \frac{\partial LMP_{1t}}{\partial E_t^{DC}} = v_1 + \kappa > 0.$$

These positive derivatives formalize the merit-order dispatch mechanics of power grids. To serve the massive, inelastic demand of a new computing facility, the electricity sector must dispatch less efficient, more expensive power plants. In regulated markets, utilities mechanically pass this convex generation cost (v_1) through to retail users via the PCA. In deregulated markets, perfect competition among generators forces this same marginal cost into the base system-wide wholesale price, while the host region is uniquely penalized by the nodal congestion premium (κ) due to physical transmission constraints. Defining the log-linearized variable rate as $\hat{p}_{1t}$, this establishes the primary channel by which an exogenous data center demand shock friction affects equilibrium: $\hat{p}_{1t} > 0$.

Proposition 2: Under a perfectly competitive market and constant returns to scale, an exogenous increase in the localized firm electricity price p_{it}^M strictly depresses the local real wage w_{it}.

By taking the total differential of the unit cost function and applying Shephard's Lemma (see Appendix D for proof), I derive the log-linearized factor price frontier for the representative manufacturing firm:

$$\hat{w}_{1t} = -\frac{s_E}{s_L}\hat{p}_{1t}^M.$$

Because the steady-state factor shares of labor (s_L) and energy (s_E) are strictly positive, the localized energy price shock ($\hat{p}_{1t}^M > 0$) dictates that the region-specific real wage must strictly decrease ($\hat{w}_{1t} < 0$). This mathematical identity formalizes the theoretical mechanism driving the manufacturing wage penalty observed empirically in Section 4.4. Because the firm is a price taker in the national output market, it cannot pass localized electricity cost shocks onto its customers. To maintain its zero-profit condition in spatial equilibrium, the firm is forced to extract the entirety of the energy cost penalty from local labor. I label this firm-side dynamic the factor price displacement effect.

Critically, the magnitude of this wage erosion is strictly scaled by the ratio of the steady-state factor shares, s_E/s_L. Because the manufacturing sector exhibits a uniquely high energy

intensity relative to its labor share, it is structurally predisposed to a more severe wage penalty. The factor price displacement effect therefore disproportionately burdens manufacturing labor compared to the less energy-dependent, service-oriented sectors of the local economy.

Proposition 3: Following the data center demand shock, the steady-state consumption of the composite good by region 1's representative household strictly decreases.

I prove this consumption contraction by contradiction utilizing the log-linearized steady-state budget constraint:

$$\bar{C}_1\hat{C}_{1t} = \bar{w}_1\bar{L}_1(\hat{w}_{1t} + \hat{L}_{1t}) - \bar{p}_1^H\bar{E}_1^H(\hat{p}_{1t}^H + \hat{E}_{1t}^H) - \bar{F}_1^H\hat{F}_{1t}^H.$$

The household faces an unavoidable spatial squeeze. The right side of the constraint must be strictly negative, as the household is forced to pay elevated variable electricity rates ($\hat{p}_{1t}^H > 0$) and surging fixed grid capacity charges ($\hat{F}_{1t}^H \geq 0$) while simultaneously suffering from the eroded real wage income ($\hat{w}_{1t} < 0$) dictated by the factor price displacement effect. Because steady-state consumption $\bar{C}_1$ is positive, it must be that $\hat{C}_{1t} < 0$. This proposition confirms that households bear the incidence of data center expansion through reduced purchasing power.

Proposition 4: Region 1's labor supply response $\hat{L}_{1t}$ to a data center demand shock is theoretically ambiguous, governed by the counteraction between a household wealth effect and a firm factor price displacement effect.

By equating the firm's inverse labor demand curve with the household's inverse labor supply curve derived from KPR preferences, I isolate the equilibrium labor supply response:

$$\hat{L}_{1t} = \frac{1}{\varphi}\left(-\frac{s_E}{s_L}\hat{p}_{1t}^M - \hat{C}_{1t}\right).$$

The theoretical ambiguity of $\hat{L}_{1t}$ stems from the two competing terms inside the parentheses. The first term represents the firm factor price displacement effect ($-\frac{s_E}{s_L}\hat{p}_{1t}^M < 0$), which disincentivizes work as the real wage falls. The second term represents the household wealth effect ($-\hat{C}_{1t} > 0$). As surging electricity bills and falling wages compress the budget constraint, the household is compelled to increase its labor supply to smooth consumption. This tension provides the solution to the labor market puzzle from Section 4.4. The empirical observation that manufacturing employment remains statistically unchanged indicates that the wealth effect perfectly offsets the factor price displacement effect in equilibrium.

Proposition 5: The effect of a data center demand shock on region 1's manufacturing output $\hat{Y}_{1t}$ is theoretically ambiguous, though heavily biased downward by a high elasticity of substitution σ.

Log-linearizing the CES production technology and substituting the firm's optimal input ratio yields the aggregate output deviation:

$$\hat{Y}_{1t} = \hat{L}_{1t} - \sigma s_E\left(1 + \frac{s_E}{s_L}\right)\hat{p}_{1t}^M.$$

While the net change in output inherits the theoretical ambiguity of the labor response $\hat{L}_{1t}$, the second term isolates an unambiguous drag on output. A high elasticity of substitution σ

dictates that the firm aggressively slashes its use of the newly expensive energy input to protect profit margins. Because energy is a required factor of production, this divestment mechanically reduces total output scale.

The empirical finding that manufacturing output remains unaffected suggests that the true value of σ is relatively low. This implication aligns with economic reality, as additional labor cannot easily replace the energy requirements of modern manufacturing. Given Leontief-like inflexibility, local firms cannot substitute away from energy, forcing them to maintain input volumes and output levels by passing the energy shock entirely into wage erosion.

Proposition 6: The effect of a data center demand shock on region 1's representative household utility is theoretically ambiguous but heavily biased downward.

I translate the log-linearized utility deviation into a Consumption Equivalent Variation (CEV), denoted τ, which represents the exact percentage reduction in baseline steady-state consumption that yields an equivalent utility loss induced by the exogenous data center load shock:

$$\tau \approx \Delta U_{1t} \approx \hat{C}_{1t} + \nu(\bar{E}_1^H)^{1-\sigma_E}\hat{E}_{1t}^H - \chi\bar{L}_1^{1+\varphi}\hat{L}_{1t}.$$

While inheriting the theoretical ambiguity of the labor response $\hat{L}_{1t}$, the net welfare incidence is dominated by the strict decline in the consumption of the composite good ($\hat{C}_{1t} < 0$) and the forced reduction in household energy use ($\hat{E}_{1t}^H < 0$). This welfare metric provides a quantifiable parameter to evaluate the localized tax imposed on host residents through the electricity channel.

Collectively, these six propositions derived from the two-region DSGE model provide a structural resolution to the manufacturing puzzle established in Section 4.4. By isolating the retail electricity channel, the DSGE framework demonstrates that localized computing expansion imposes a structural tax on the host community. Because energy-intensive firms operate in competitive national markets with rigid production technologies (low σ), they are mechanically forced to pass energy price shocks onto captive local labor. I define this mechanism as the factor price displacement effect, which drives the observed erosion of real wages. This collapse in income, alongside surging utility costs, triggers a counteracting wealth effect. This effect compels households to maintain their labor supply to smooth consumption, thereby stabilizing aggregate employment and output while strictly deteriorating net welfare in the host locality.

While the theoretical mechanisms that arise from my structural model explain the empirical results in a cohesive manner, I acknowledge that this theoretical framework inherently captures a short-to-medium-run equilibrium. By intentionally abstracting from capital accumulation and firm mobility, the model isolates the immediate margins of adjustment available to the local economy. Over a sufficiently long horizon, sustained localized energy price premiums would likely induce structural capital investments or outright firm relocation to cheaper grid territories, thereby relaxing the localized grid constraints that drive the factor price displacement effect. Consequently, the derived mechanics represent the binding incidence placed on host communities before long-run capital and spatial adjustments can materialize.

To quantify the magnitude of macroeconomic incidence and evaluate the comparative frictions of regulated versus deregulated market architectures, I formally parameterize the model in Section 6.

6 Model Calibration

Following the temporal resolution of my data, I calibrate the two-region DSGE model to the U.S. economy at an annual frequency.[20] Due to the diversity of parameters introduced by my structural framework, I employ a mixed-methodology approach to parameterize the model. First, I draw standard macroeconomic parameters from established macroeconomic literature. Second, I calibrate baseline economic scales and steady-state factor shares directly from my constructed county-year panel that includes federal administrative data. Finally, I employ a direct analytical mapping of parameters to reduced-form estimates to extract the deep structural parameters governing grid market frictions. This approach ensures that the theoretical framework replicates the observed macroeconomic incidence of computing infrastructure without relying on simulation-based estimation.

The remainder of this section details this calibration strategy. Section 6.1 establishes the demographics of the two-region model and the exogenous data center shock. Section 6.2 details the parametrization of electricity market architecture across both regulatory regimes. Section 6.3 outlines the calibration of the household and firm sectors. A comprehensive summary of all calibrated parameters is provided in Table 7, while the full algebraic derivations for parameter extraction are relegated to Appendix F.

6.1 Demographic and Data Center Sector Parameters

The demographic scale of the baseline economy is anchored to the geometric means of the empirical dataset. The host region's population size (N_1) is calibrated to the average population of treated counties within the panel, while the broader grid region ($N_{grid} = N_1 + N_2$) is calibrated to the average population footprint of a Balancing Authority (BA). For the real-world scenarios evaluated in Section 7.2, I replace these generalized geometric means with the actual observed county-level and BA-level household populations corresponding to each specific host site.

The magnitude of the exogenous data center energy shock is governed by both the facility's capacity footprint (D_t) and the load factor δ. Shehabi et al. (2024) find significant variance in compute facility capacity utilization across tasks. Standard hyperscale activity yields a load factor of $\delta = 0.50$, while AI training and inference yield $\delta = 0.80$ and $\delta = 0.40$, respectively. For baseline analysis, I define the standard exogenous data center shock as the addition of a 100 MW facility operating at a 50% load factor ($\delta = 0.50$). This power demand and utilization rate best represents the industry standard for blended data center operations, capturing a mix of continuous AI training workloads, diurnal inference tasks, and traditional hyperscale operations. To capture the real-world scenarios discussed in Section 7.2, I adopt $\delta = 0.50$ for Loudoun County to reflect standard hyperscale expansion and $\delta = 0.80$ for xAI and OpenAI Stargate to reflect the facilities' AI training focus. Furthermore, I substitute the static 100 MW baseline shock with the empirical, multi-year capacity trajectories (D_t) to reflect historical data or projected forecasts for each specific development.

6.2 Electricity Market Parameters

The structural parameters governing the electricity sector dictate the magnitude of the localized energy price penalty and the regional socialization of infrastructure costs. I independently calibrate these parameters for both the regulated ($\Omega = 0$) and deregulated ($\Omega = 1$) market architectures.

Under the regulated market, I anchor the baseline variable cost of generation (v_0) to \$24/MWh, reflecting the Lazard (2025) lower-bound marginal cost estimate for combined-cycle gas turbine

[20] To ensure consistency with the terminal year of the empirical panel, all nominal monetary values are adjusted to real 2025 dollars using the Bureau of Labor Statistics Consumer Price Index for All Urban Consumers (CPI-U).

Table 7: Summary of Calibrated Model Parameters

Baseline Scales		
N_1	Host Region Households	68,414
N_{grid}	Grid Region Households	751,550
$\bar{p}$	Baseline Retail Price	\$133.47/MWh
$\bar{p}^{\text{energy}}$	Baseline Wholesale Price	\$42.80/MWh
$\bar{E}$	Total Baseline Load	21,119,717 MWh
$\bar{F}$	Baseline Fixed Tariff	\$200.47/year
$\bar{\rho}$	Baseline Hazard	0.5943
Electricity Sector		
v_0	Baseline Generation Cost	\$24.00
v_1	Marginal Cost Slope	8.902×10^{-7}
ω	Local Capital Share	0.6428
Ψ	Unrecovered Capital Cost Scalar	\$242,412.09
γ_0	Hazard Intercept	0.3817
γ_1	Hazard Sensitivity	1.744×10^{-7}
κ	Marginal Congestion Premium	5.756×10^{-6}
Household Preferences		
r	Annualized Real Interest Rate	0.04
β	Discount Factor	0.9615
σ_E	Energy Curvature	5.00
φ	Inverse Frisch Elasticity	5.00
ν	Energy Utility Weight	201.90
χ	Labor Disutility Weight	742.84
Firm Technology		
σ	Substitution Elasticity	0.60
α	Energy Distribution Weight (Representative)	0.2655
A	TFP (Representative)	169,148.41

Notes: All nominal monetary values are expressed in real 2025 dollars. The unrecovered capital gap Ψ is evaluated per standard 100 MW data center shock.

(CCGT) baseloads, which I adopt as the representative marginal technology for the stable, dispatchable capacity required to support continuous data center operations.[21] The convexity of the marginal cost curve (v_1) is then calibrated to match the observed steady-state wholesale energy price ($\bar{p}^{energy}$). To evaluate the three real-world scenarios in Section 7.2, I calibrate distinct marginal cost slopes (v_1) for each respective host Balancing Authority (PJM, TVA, and ERCOT). By anchoring these grid-specific parameters to localized steady-state wholesale prices and aggregate baseline loads, I force the structural generation supply stack to accurately replicate the heterogeneous marginal cost dynamics of each regional market.

To parameterize the capital recovery mechanism, I map the model's fixed charge and rate-case hazard functions directly to the empirical estimates. I extract the local capital cost share

[21] While renewable and nuclear generation exhibit near-zero marginal costs, they serve as inappropriate theoretical anchors for v_0. The inelastic, continuous power requirements of data centers necessitate dispatchable, stable capacity that intermittent renewables cannot reliably provide without massive auxiliary storage. Furthermore, defining the baseline generation floor near zero would mathematically distort the parameterization of the convex supply stack. CCGT therefore represents the true marginal baseload technology realistically dispatched to meet unyielding computing loads.

(ω) and the unrecovered capital gap per MW (Ψ) by equating the model's theoretical cost transmission to the empirical unconditional base rate effect and the BA-wide spillover effect. Similarly, I extract the hazard function intercept (γ_0) and sensitivity (γ_1) to perfectly match the baseline rate-case approval probability and the empirical 0.65 percentage point surge induced by a 100 MW load shock (see Appendix F.1 for derivations).

Under the deregulated market, local utilities are insulated from generation expenditures, and residential consumers are temporarily shielded from immediate wholesale volatility by fixed-rate Retail Electric Provider (REP) contracts. I set the duration of these household contracts to $k = 3$ years, preserving the lagged nature of retail prices in a model with annual frequency and reflecting the standard maturity of residential fixed-rate plans in competitive retail markets. However, wholesale Locational Marginal Prices (LMPs) remain highly sensitive to localized grid congestion. Because these volatile wholesale nodal premiums (κ) and specialized industrial discount tariffs are unobserved in the URDB panel data, I cannot calibrate κ through direct empirical matching. Instead, I utilize a direct analytical calibration of the model's equilibrium conditions. By taking the log-linearized Factor Price Frontier derived in Proposition 2, I back out the exact wholesale nodal LMP spike ($\hat{p}^M$) required to force the empirical -0.49% manufacturing wage penalty. Anchoring the nodal premium to this targeted empirical moment allows for the exact identification of the unobserved spatial congestion friction inherent to deregulated power grids (see Appendix F.2).

6.3 Household and Firm Parameters

I calibrate the representative household using standard macroeconomic conventions. The annualized real interest rate is set to $r = 0.04$, implying a steady-state discount factor of $\beta = 0.9615$. Following standard practice, the steady-state labor supply is fixed at $\bar{L} = 1/3$ of the available time endowment. I set the inverse Frisch elasticity to $\varphi = 5$, mapping to a labor supply elasticity of 0.20 as estimated by Elminejad et al. (2023). Symmetrically, the energy curvature parameter is set to $\sigma_E = 5$, matching a short-run residential electricity demand elasticity of -0.20, a value corroborated by recent estimates from Mamkhezri (2025). The KPR utility weights for energy (ν) and labor disutility (χ) are left as residual parameters, solved analytically via the household's intratemporal first-order conditions evaluated at the steady state (see Appendix F.3).

For the representative manufacturing firm, I set the CES elasticity of substitution to $\sigma = 0.60$, reflecting empirical estimates of capital-energy complementarity in heavy industry (Bretschger and Jo, 2024; Hassler et al., 2012). The steady-state factor cost shares are calibrated to $s_E = 0.025$ using a Department of Energy estimate for aggregate manufacturing energy intensity and $s_L = 0.30$ using BLS sector data. I utilize these manufacturing-specific cost shares exclusively to precisely identify the nodal premium via the direct analytical calibration detailed in Section 6.2.

To evaluate the broader macroeconomic incidence of data center shocks beyond the manufacturing sector, I also parameterize a "macro-average" representative firm. This generalized firm adopts a substantially lower energy intensity ($s_E = 0.010$) and an aggregate labor share ($s_L = 0.5985$) derived from the BLS average across all non-farm sectors between 2014 and 2024. For both firm variants, the CES distribution parameter (α) and total factor productivity (A) are solved analytically using the optimal input ratio and the zero-profit condition (see Appendix F.3). I report the representative calibration of these technology parameters in Table 7 and utilize this calibrated macro-average firm to conduct the baseline analysis of the standard exogenous data center shock. Subsequently, I calculate distinct, locality-specific representative firm cost shares and production technology parameters to accurately reflect the economic composition of the host communities evaluated in the three real-world scenarios (see Section 7.2).

7 Model Application

Equipped with the fully calibrated DSGE model, I move beyond the reduced-form empirical estimates of Section 4 to quantify the structural welfare incidence of large-scale computing infrastructure. To preserve the inherent nonlinearities of the electricity sector, I solve the model's transition path numerically using a period-by-period iterative root-finding algorithm. This recursive approach ensures that the system's market-clearing and optimality conditions are satisfied exactly at each step, avoiding the approximation errors associated with standard log-linearization techniques.

Section 7.1 establishes an exogenous 100 MW facility as a baseline scenario, mirroring the standard introduction of a major hyperscale data center to a local community. By applying this exact capacity shock to both the regulated and deregulated model variants, I conduct a direct comparative statics analysis. This controlled exercise isolates how divergent electricity market architectures mechanically alter the spatial distribution of grid costs, dictating whether infrastructure burdens are socialized across the broader balancing authority or heavily localized to the host county.

Section 7.2 then applies the structural framework to three high-fidelity, real-world case studies: the historic build-out of "Data Center Alley" in Loudoun County, Virginia, the ongoing xAI Colossus development in Memphis, Tennessee, and the OpenAI Stargate campus under construction in the Texas Panhandle. By grounding the model in site-specific demographics, localized grid cost slopes, and empirical multi-year capacity trajectories, these applications quantify the true welfare costs of both sustained historical hyperscale concentration and acute AI-driven expansion.

7.1 100 MW Benchmark Scenario

To rigorously isolate the institutional mechanisms of cost pass-through, I first establish a controlled comparative environment. I simulate the exogenous injection of a standard 100 MW capacity shock ($D_t = 1$) operating at a baseline 50% load factor ($\delta = 0.50$) into the host locality (Region 1). The remainder of the Balancing Authority acts as the non-host grid footprint (Region 2). This stylized benchmark mirrors the typical deployment of a major hyperscale facility, facilitating a direct evaluation of how regulated ($\Omega = 0$) versus deregulated ($\Omega = 1$) market architectures diverge in their spatial distribution of macroeconomic incidence.

Figures 6 and 7 illustrate the impulse response functions calculated by the nonlinear root-finding algorithm (see Appendix G for complete set of impulse response functions). The difference in impact between the two regulatory regimes is stark. Table 8 reveals that the deregulated host community suffers a catastrophic localized tax, bearing a Consumption Equivalent Variation (CEV) of -0.150%, which translates to an annual loss of \$126.40 per household. In contrast, both regions under the regulated regime fare significantly better, experiencing relatively symmetric CEVs of -0.019% (Region 1) and -0.016% (Region 2). Even the non-host deregulated region escapes largely unscathed with a CEV of -0.0125%.

This stark divergence in welfare loss is fundamentally driven by the structural differences in cost pass-through mechanisms. Under the regulated architecture, the institutional socialization of grid burdens effectively diffuses the shock across the entire ratepayer base. As illustrated in Figure 6, the unified grid variable price spikes symmetrically by 0.03 cents per kWh across both Region 1 and Region 2. Concurrently, the massive load shock generates an unrecovered capital gap (Δ), compelling the host utility to initiate a rate case. The stochastic probability of regulatory approval surges to a steady-state equilibrium of approximately 0.60. Once approved, this capital recovery manifests as an asymmetric fixed charge hike for Region 1 households. Driven by the localized component of infrastructure scaling, Region 1 households face an elevated monthly burden of \$0.20, whereas Region 2 ratepayers absorb a smaller \$0.01 monthly increase driven by socialized Transmission Cost Adjustments (TCAs). Ultimately, the regulated

Table 8: Welfare Cost of 100 MW Capacity Shock

	Region 1 (Treated Host)	Region 2 (Rest of BA)
Regulated Market		
Households	68,414	683,136
Welfare Loss	-0.0186	-0.0159
Welfare Cost per HH	15.67	13.39
Aggregate Welfare Cost	1,071,753	9,146,071
Deregulated Market		
Households	68,414	683,136
Welfare Loss	-0.1499	-0.0125
Welfare Cost per HH	126.40	10.58
Aggregate Welfare Cost	8,647,273	7,224,678

Note: Welfare impacts are calculated at the terminal steady state ($t = 10$). The baseline consumption ($\bar{C}$) is \$84,331.95 per household. Welfare cost per HH represents the monetary equivalent of the lost consumption per household, evaluated at baseline prices. Welfare loss is expressed as a percentage. Welfare cost per household and aggregate welfare cost are expressed in dollars per year.

framework socializes the bulk of both generational and transmission infrastructure investments required to accommodate the data center across the broader BA region.

Conversely, market deregulation structurally traps the economic incidence of data center expansion within the host municipality. Figure 7 demonstrates Region 1's Locational Marginal Price (LMP) surging by 0.28 cents per kWh, dwarfing the negligible 0.02 cents per kWh increase absorbed by Region 2. While this minor 0.02-cent baseline shift reflects the systemic, grid-wide rise in marginal generation costs shared by all participants in the competitive wholesale market, the severe spatial divergence is exclusively the manifestation of localized grid congestion.

In a deregulated environment governed by Locational Marginal Pricing, energy clearing prices are strictly bound by the physical realities of the transmission network. As the massive, inelastic computing load saturates local transmission capacity, the balancing authority is physically barred from delivering cheaper remote power to the host county. The deregulated architecture therefore forces the host locality's wholesale price to directly internalize this congestion. This unyielding nodal premium effectively traps the financial burden locally, imposing a severe penalty on the host economy while simultaneously insulating the broader balancing authority from meaningful spillover.[22]

The localized wholesale volatility in the deregulated host region is immediately transmitted to the local manufacturing firm, which operates directly as a buyer in the wholesale electricity market. For residential ratepayers, the deregulated impulse response functions exhibit a gradual, three-period phase-in. This dynamic captures the $k = 3$ year duration of fixed-rate Retail Electric Provider (REP) contracts, which temporarily shield households from immediate wholesale volatility before their tariffs irrevocably converge to the elevated LMP steady state.

The severe localized price hike in the deregulated regime fuels a violent factor price displacement effect. Because industrial operations are directly exposed to wholesale market clearing, the 0.28 cents per kWh LMP spike in Region 1 triggers an immediate, sharp collapse in the

[22]While restructured-regulated markets (RRMs) also operate competitive wholesale markets under the same locational marginal pricing system, they circumvent this severe spatial penalty through their retail architecture. Because utility monopolies in RRMs retain exclusive retail rights across expansive, multi-county service territories, they inherently function as financial aggregators. These utilities pool the extreme nodal congestion premiums incurred at the data center host site with the cheaper wholesale prices of surrounding counties, subsequently socializing that blended cost across their entire captive ratepayer base via the Power Cost Adjustment (PCA).

Figure 6: 100 MW Shock under Regulated Regime ($\Omega = 0$)

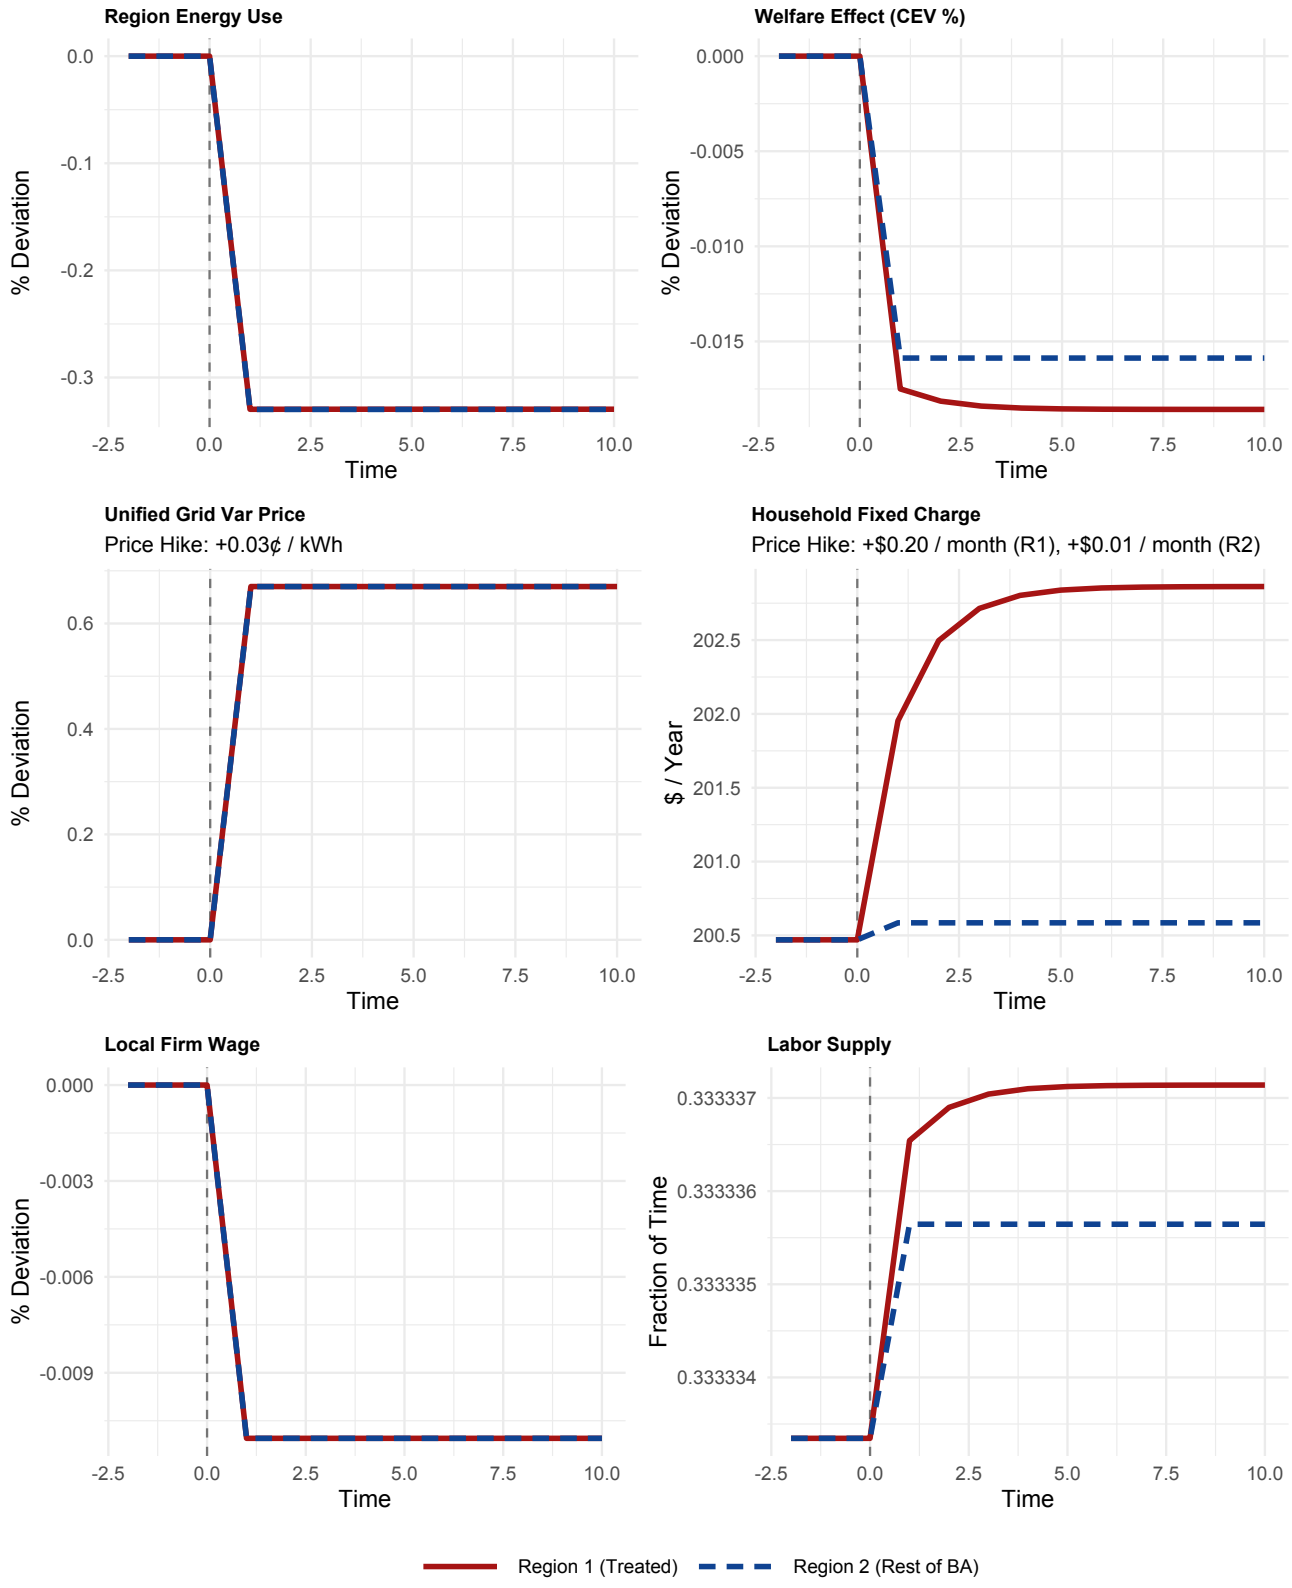

Note: For analytical tractability within the general equilibrium calculations, the impulse response function for the household fixed charge depicts the expected tariff trajectory. This is evaluated as the probability-weighted average of the baseline and target rates via the rate-case hazard function rather than a single stochastic realization of the regulatory approval process. Appendix G contains the full panel of impulse response functions.

Figure 7: 100 MW Shock under Deregulated Regime ($\Omega = 1$)

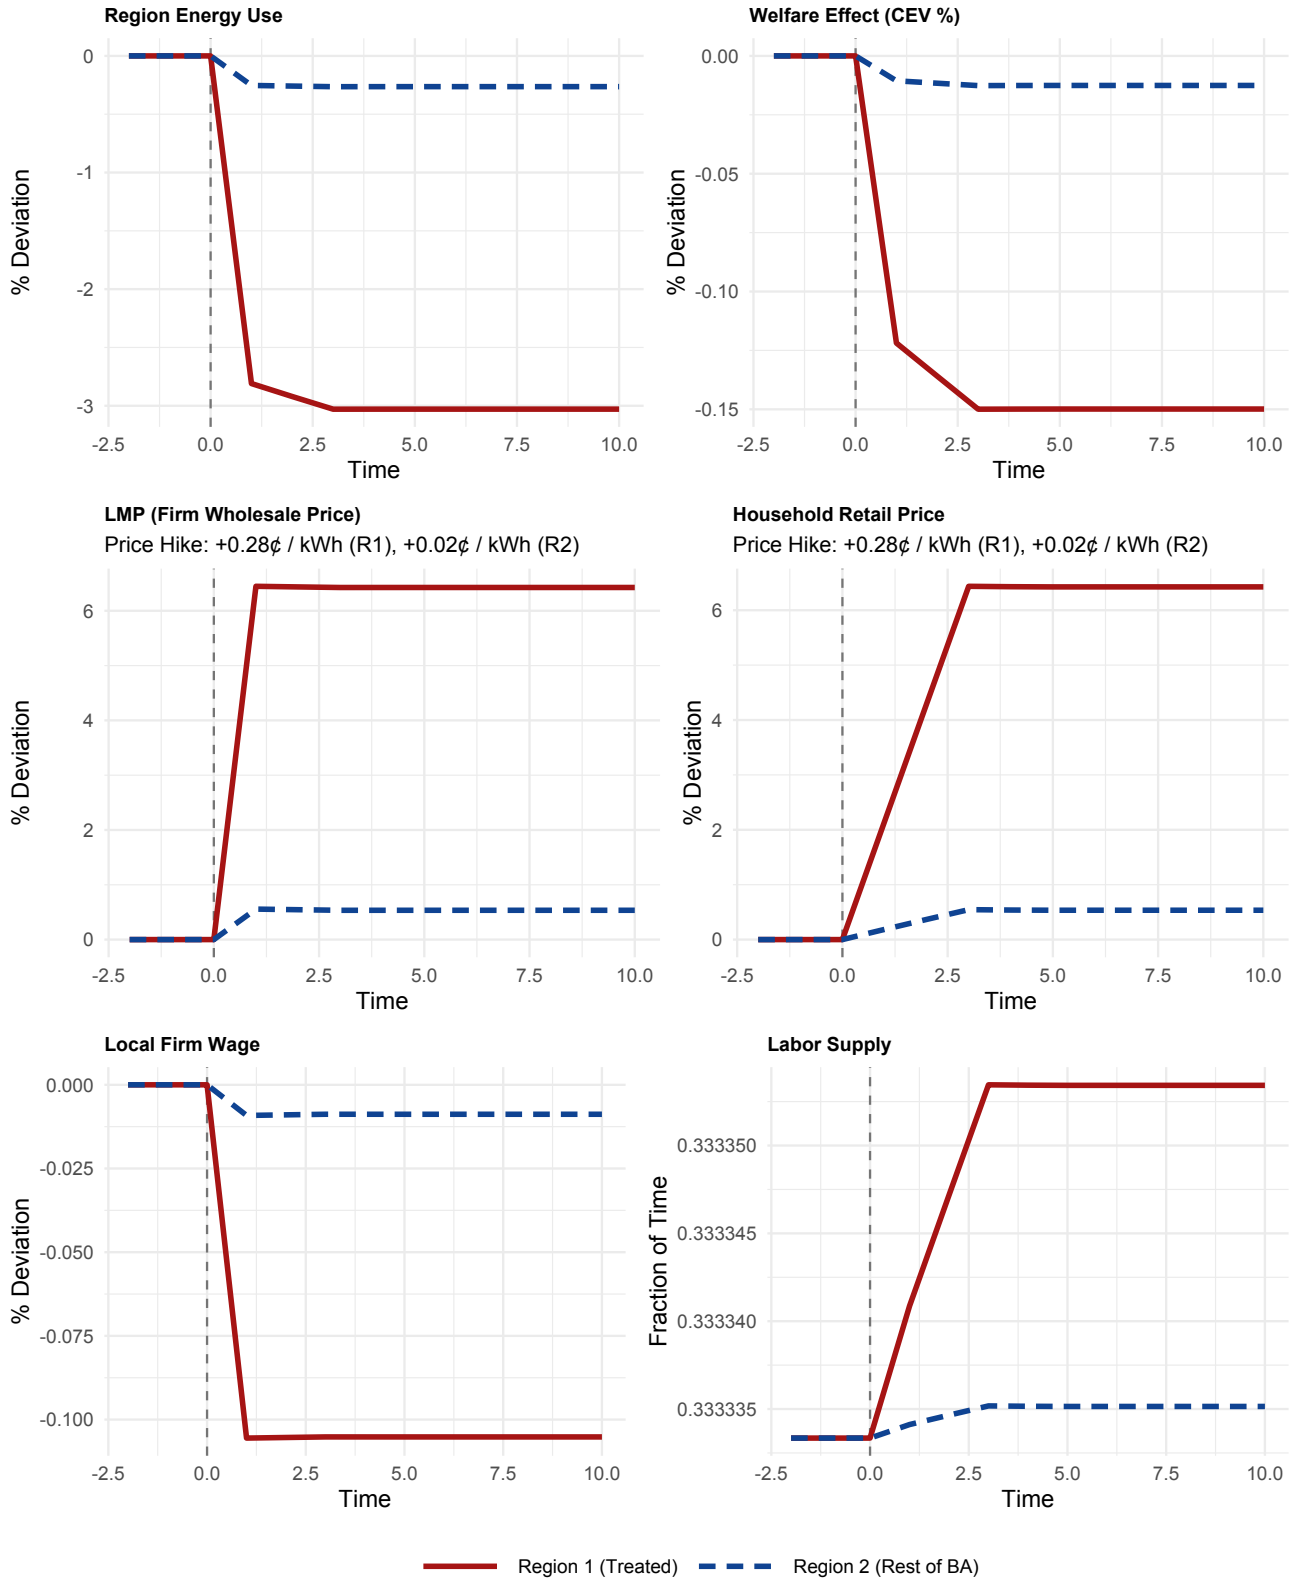

Note: Appendix G contains the full panel of impulse response functions.

local real wage of 0.105%. Perfectly competitive firms must offset surging energy input costs by extracting the penalty entirely from local labor.

This severe wage erosion, compounded by the eventual surge in household retail electricity prices, drastically compresses the household budget constraint. The result is a precipitous drop in steady-state consumption, which in turn triggers a powerful wealth effect. Driven by the necessity to smooth consumption, the representative household is compelled to increase its labor supply. The (1) factor price displacement effect suppressing labor demand and the (2) wealth effect expanding labor supply counteract each other under both regimes, resulting in a limited net change in the equilibrium labor supply. However, the absolute magnitude of both forces is vastly greater in the deregulated environment. The sheer scale of the localized wage collapse and subsequent consumption contraction mechanically dictates the catastrophic welfare loss (-0.150% CEV) observed in the deregulated host community.

The theoretical mechanisms formalized by the model's propositions actively manifest across both market architectures. Data center-induced price hikes structurally depress the local real wage via the factor price displacement effect, mechanically eroding the purchasing power of host households. To counteract this precipitous decline in consumption, residents are compelled by a wealth effect to increase their labor supply, ultimately yielding a depressed standard of living despite stabilized aggregate employment. Yet, the deregulated regime sustains a profoundly amplified welfare penalty because its unbundled framework structurally lacks the retail aggregation necessary to socialize volatile wholesale shocks. The extremity of this spatial divergence is overwhelmingly dictated by the severity of local nodal congestion.

7.2 Real-World Case Studies

While the stylized 100 MW benchmark isolates the foundational mechanics of cost pass-through, it inherently abstracts from the extreme scale and temporal dynamics that characterize modern computing infrastructure deployment. To bridge the gap between theoretical comparative statics and realized macroeconomic incidence, I apply the structural framework to three high-fidelity, real-world case studies. The primary merit of this exercise lies in its ability to quantify the cumulative welfare tax imposed on host communities by substituting static load assumptions with dynamic, multi-gigawatt capacity trajectories. By evaluating these specific geographic developments, the model captures the complex interaction between unprecedented infrastructure concentration, local industrial composition, and regional grid constraints.

To ensure each simulation accurately reflects the unique macroeconomic friction of its respective locality, I recalibrate the DSGE framework using the empirical strategy detailed in Section 6 and Appendix F. I first substitute the generalized geometric means of the baseline economy with the exact household populations of the host county (N_1) and the surrounding balancing authority (N_{grid}). To capture the distinct composition of the local labor market, I compute a bespoke composite firm ratio (s_E/s_L) using 2014 county-level sector intensities, thereby dictating the specific factor price displacement effect experienced by the host economy. Finally, I calibrate a marginal cost slope (v_1) for each scenario's BA region anchored to their respective steady-state wholesale prices and aggregate baseline loads. This tailored parametrization forces the structural generation supply stack to replicate the heterogeneous cost dynamics inherent to each grid architecture.

Loudoun County, Virginia, widely recognized as "Data Center Alley," presents a unique historical case study of sustained hyperscale concentration within a regulated market. Between 2014 and 2025, the county's data center footprint expanded aggressively, absorbing an unprecedented 7.0 GW of power demand. Because most development during the observation window was standard hyperscale facilities, I assume a load factor of 50%. As illustrated by Figure 8, the gradual trajectory of this capacity shock triggers a monotonic decline in regional welfare, resulting in a terminal Consumption Equivalent Variation (CEV) of -0.117% for the host county and -0.048% for the rest of the PJM grid. Translating these figures to monetary equivalents,

Table 9: Welfare Cost of Simulated Shock Scenarios

	Region 1 (Treated Host)	Region 2 (Rest of BA)
Loudoun County (Regulated 7.0 GW Shock)		
Households	173,199	25,865,345
BA Population Share	0.67	99.33
Composite Firm Ratio	0.0382	–
Welfare Loss	-0.1170	-0.0483
Welfare Cost per HH	98.68	40.77
Aggregate Welfare Cost	17,090,844	1,054,448,884
xAI Colossus (Regulated 1.9 GW Shock)		
Households	367,669	3,164,096
BA Population Share	10.41	89.59
Composite Firm Ratio	0.0326	–
Welfare Loss	-0.1534	-0.1437
Welfare Cost per HH	129.34	121.18
Aggregate Welfare Cost	47,553,814	383,426,641
OpenAI Stargate (Deregulated 4.1 GW Shock)		
Households	59,615	7,389,972
BA Population Share	0.80	99.20
Composite Firm Ratio	0.0384	–
Welfare Loss	-10.5009	-0.1373
Welfare Cost per HH	8,855.63	115.75
Aggregate Welfare Cost	527,930,328	855,359,378

Note: Welfare impacts are calculated at the terminal observation/forecast year for each respective case study. The baseline consumption ($\bar{C}$) is \$84,331.95 per household. Welfare cost per HH represents the monetary equivalent of the lost consumption per household, evaluated at baseline prices. BA population share and welfare loss are expressed as percentages. Welfare cost per household and aggregate welfare cost are expressed in dollars per year. The composite firm ratio represents the local structural energy-to-labor cost share ratio calibrated from 2014 county-level employment data.

Loudoun households face an annual welfare penalty of \$98.68, while the broader PJM ratepayer base absorbs a \$40.77 annual cost.

The mechanism driving this widespread economic burden is the regulated market's institutional cost socialization. The massive 7.0 GW capacity expansion necessitates continuous baseload generation dispatch, driving up the unified grid variable price by a maximum of 0.05 cents per kWh across both regions. Simultaneously, the staggering unrecovered capital costs of transmission and distribution upgrades force the host utility to execute multiple rate cases in succession, elevating the Loudoun household fixed charge by \$4.85 per month, while pushing \$0.02 onto the rest of PJM via transmission cost adjustments. This combined electricity price shock structurally depresses the local real wage by 0.04% through the factor price displacement effect. As energy becomes prohibitively expensive, both the local manufacturing firm and the representative household are forced to aggressively substitute away from electricity, leading to a 0.6% contraction in endogenous regional energy use.

The second scenario transitions to the nascent artificial intelligence arms race, evaluating the aggressive deployment of xAI's Colossus facilities in Shelby County, Tennessee. Engineered to rapidly train the Grok foundational model, this development is forecasted to inject a staggering 1.9 GW of power demand into the federally regulated Tennessee Valley Authority (TVA) between

Figure 8: Impact of Data Center Alley Build-Out (Loudoun County, VA)

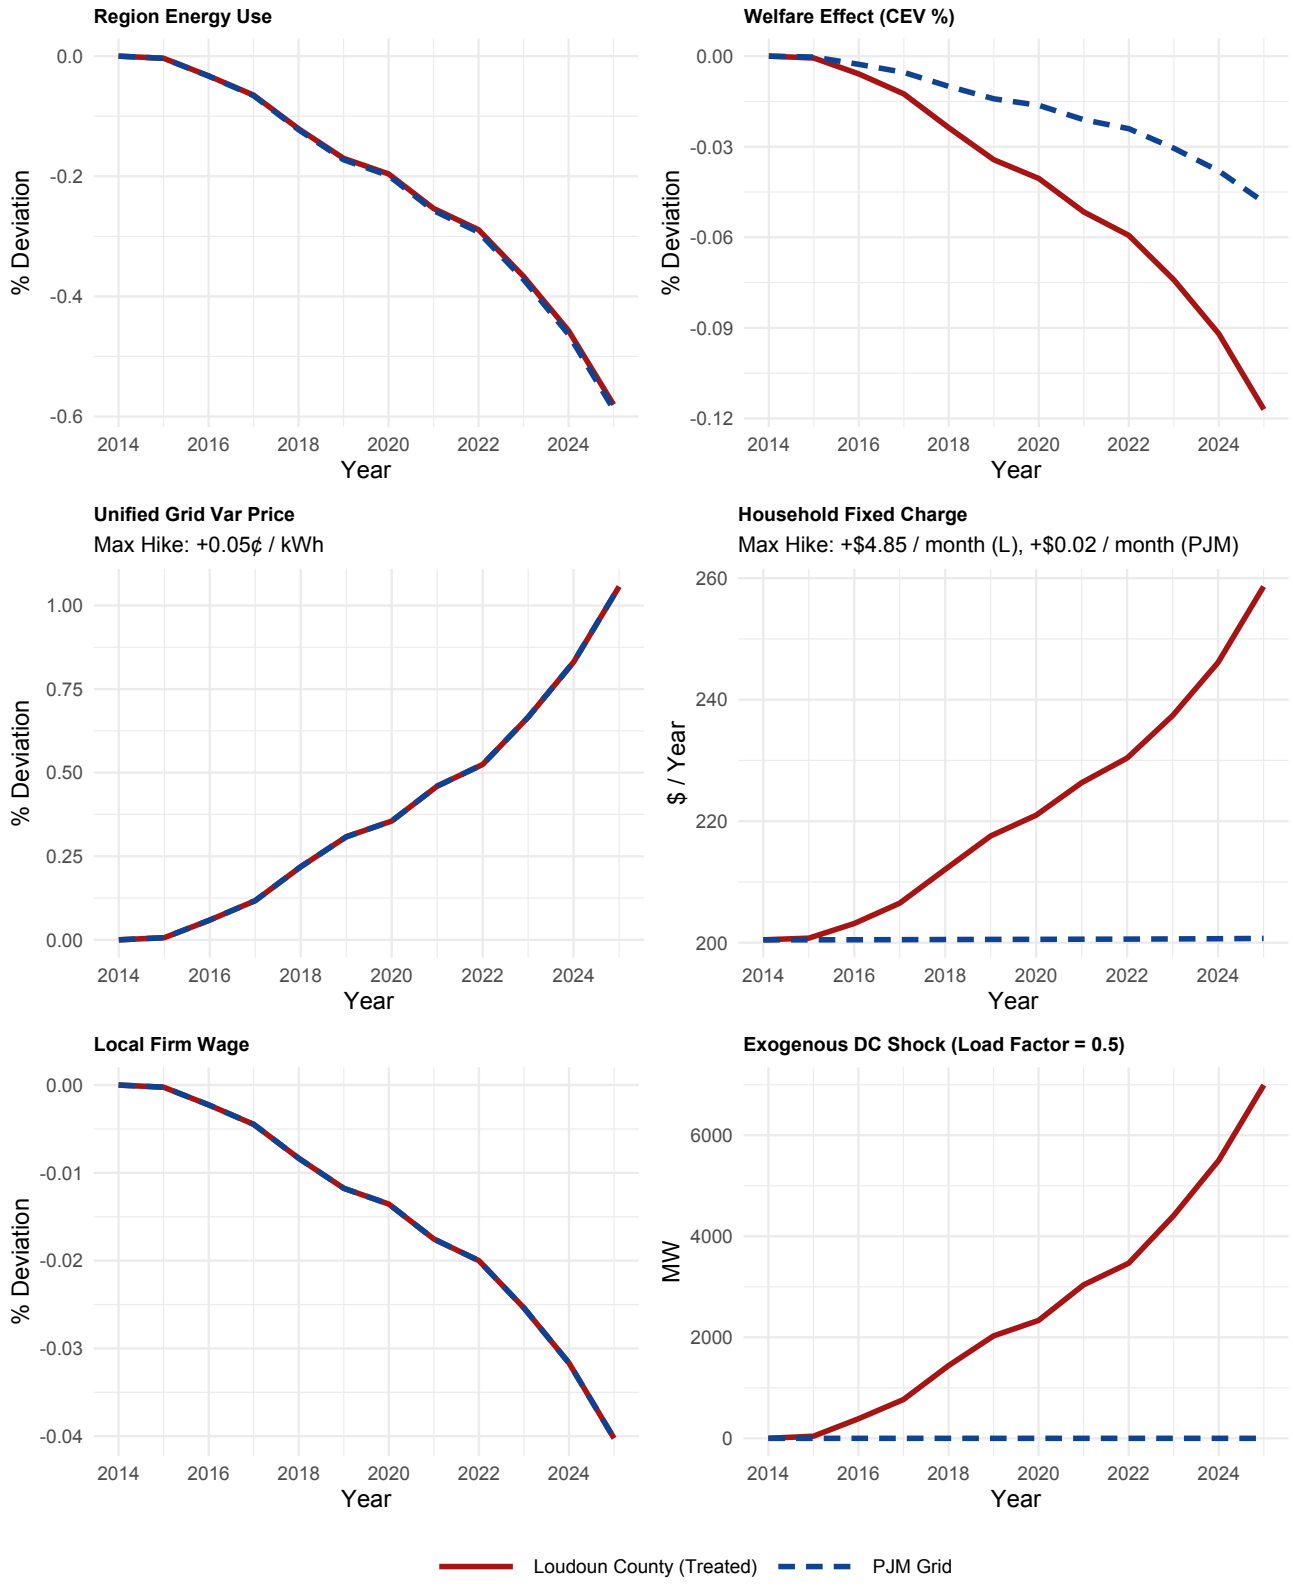

Note: For analytical tractability within the general equilibrium calculations, the impulse response function for the household fixed charge depicts the expected tariff trajectory. This is evaluated as the probability-weighted average of the baseline and target rates via the rate-case hazard function rather than a single stochastic realization of the regulatory approval process. Appendix H contains the full panel of impulse response functions.

Figure 9: Impact of xAI Colossus Development (Shelby County, TN)

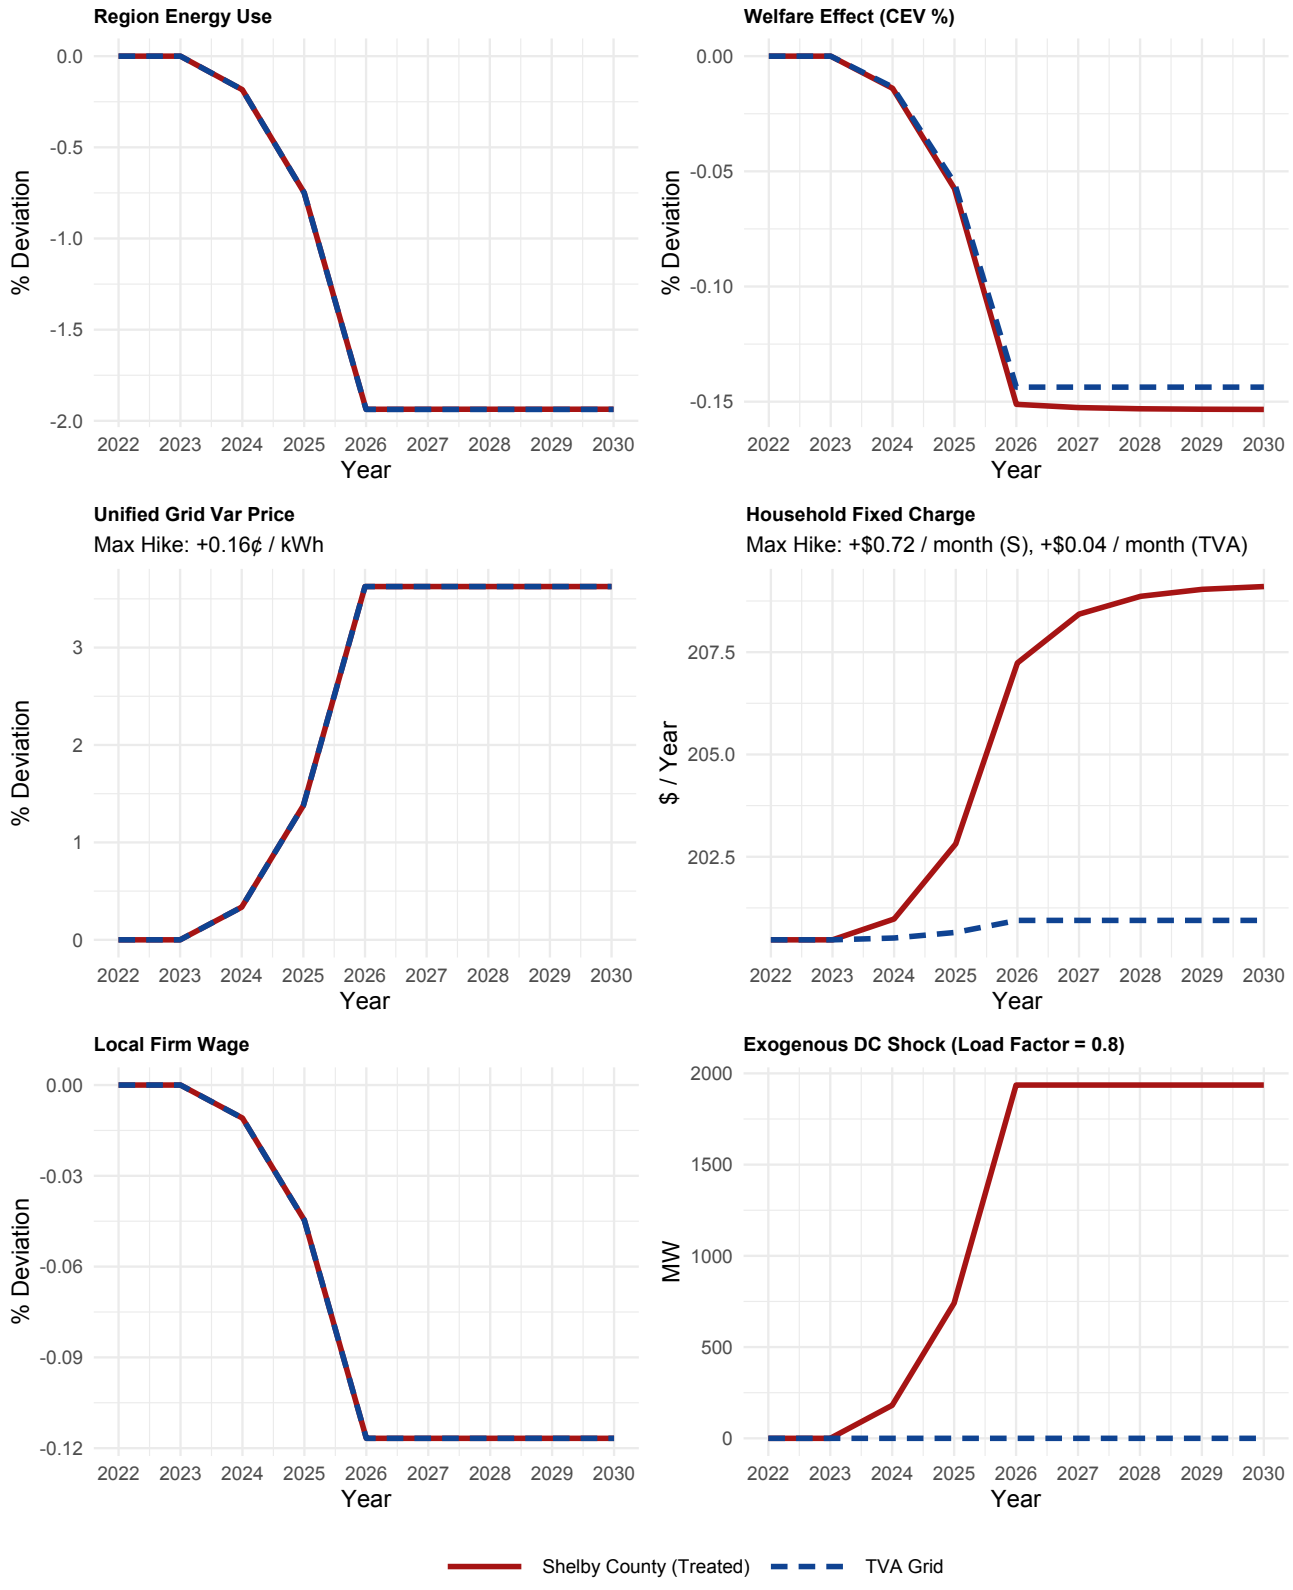

Note: For analytical tractability within the general equilibrium calculations, the impulse response function for the household fixed charge depicts the expected tariff trajectory. This is evaluated as the probability-weighted average of the baseline and target rates via the rate-case hazard function rather than a single stochastic realization of the regulatory approval process. Appendix H contains the full panel of impulse response functions.

Figure 10: Impact of OpenAI Development (Taylor & Shackelford Counties, TX)

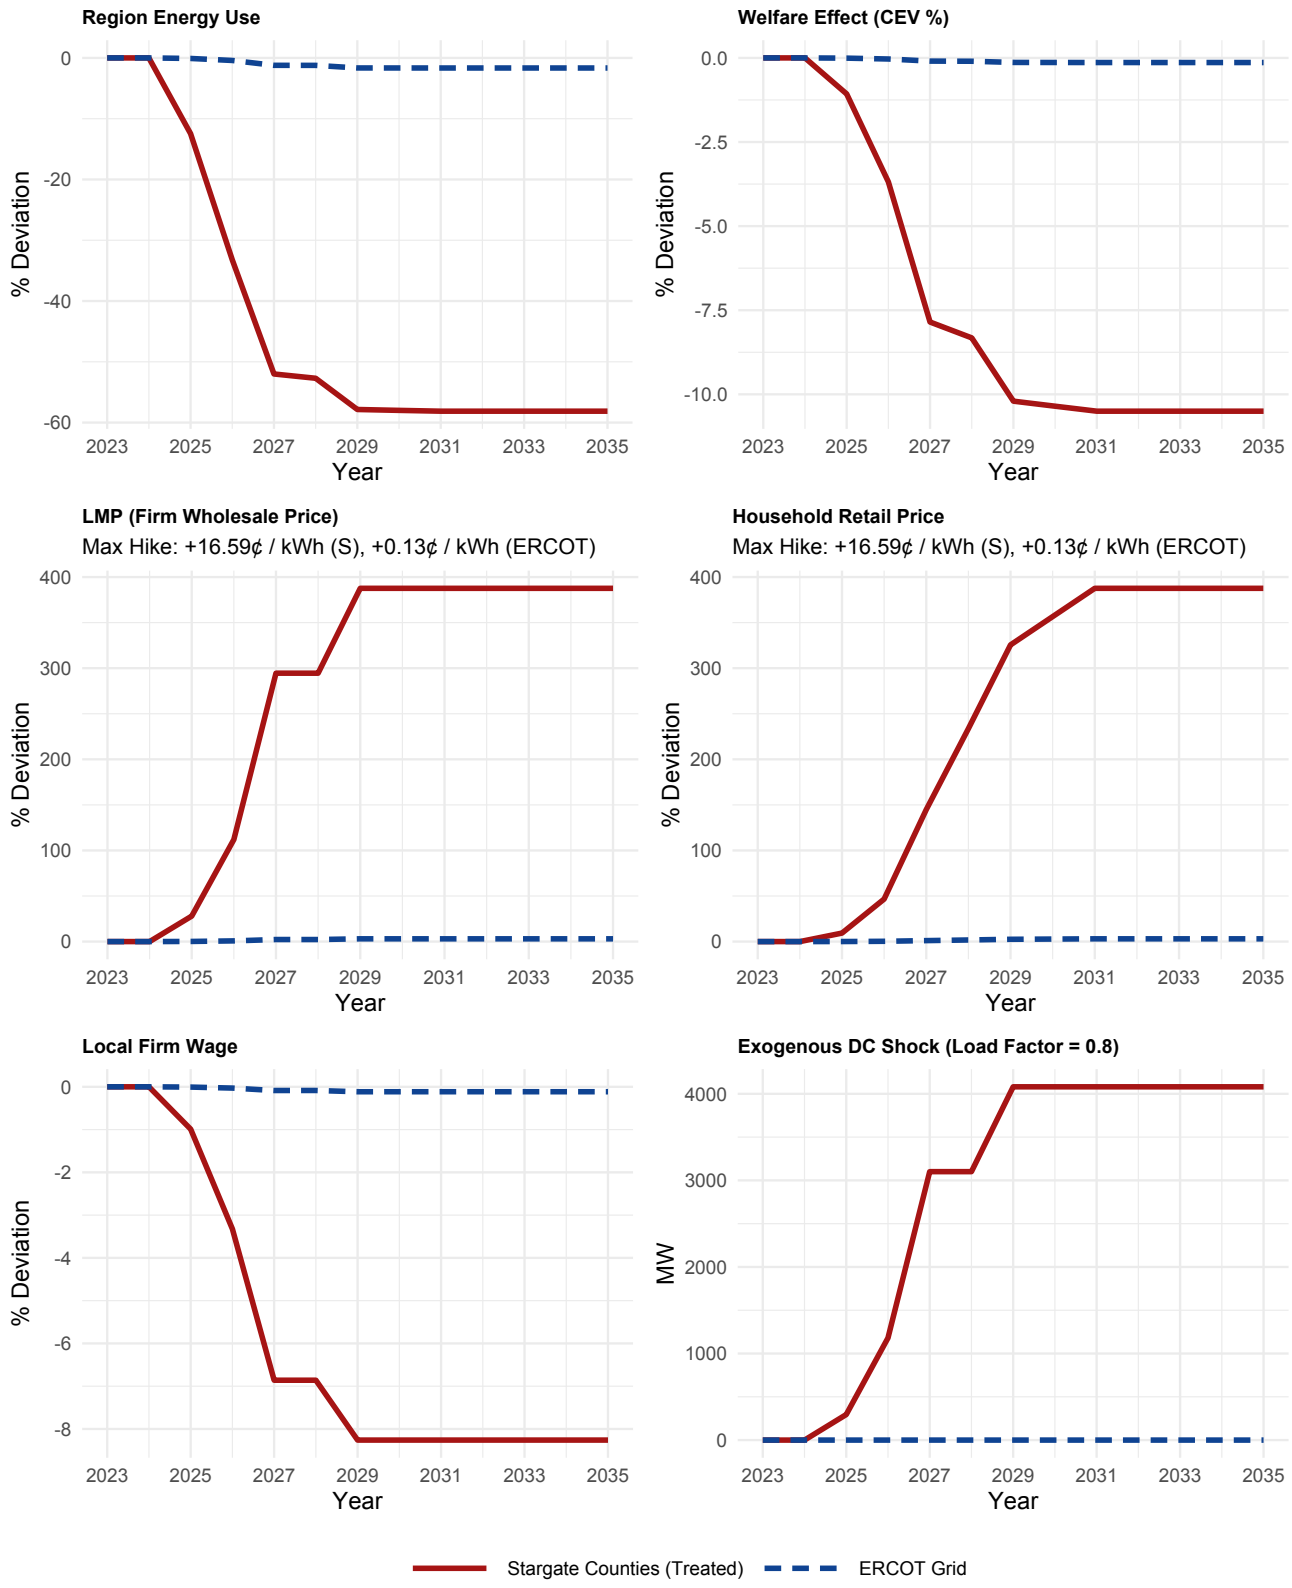

Note: Appendix H contains the full panel of impulse response functions.

2022 and 2030. Because Colossus is explicitly designed for AI training, I assume a load factor of 80%. As depicted in Figure 9, this acute capacity shock produces a remarkably symmetric economic burden across the grid, resulting in a terminal CEV of -0.153% for the host county and -0.144% for the broader TVA ratepayer base. In monetary terms, this translates to an annual welfare penalty of \$129.34 for Shelby households and \$121.18 for remaining grid constituents. This tight alignment in per-household welfare loss is an artifact of the TVA's unique population distribution. Because Shelby County constitutes a substantial 10.41% of the entire BA ratepayer base, the localized capital burden is distributed across a uniquely large municipal population, significantly diluting the host's per-household fixed charge hike and aligning its overall welfare penalty closely with that of the rest of the grid.

The relentless, high-utilization power demand of the AI cluster forces continuous baseload dispatch, driving the unified grid variable price up by 0.16 cents per kWh. Compounding this variable rate spike, the massive capital expenditures required to support the facility trigger successive rate cases, elevating the Shelby household fixed charge by \$0.72 per month while distributing a \$0.04 monthly increase to the rest of the TVA through transmission cost adjustments. This synthesized electricity price shock structurally depresses the local real wage by nearly 0.12% via the factor price displacement effect. As energy input costs surge, both the local manufacturing sector and the representative household must aggressively curtail their consumption, precipitating a steep 2.0% collapse in endogenous regional energy use.

The final scenario simulates the prospective macroeconomic fallout of OpenAI's monumental Stargate campus slated for Taylor and Shackelford Counties in Texas. Situated within ERCOT's fully deregulated wholesale market, this forward-looking application projects a colossal 4.1 GW demand shock onto the Texas Panhandle between 2023 and 2035. Because exact operational details remain undisclosed, I assume an 80% load factor consistent with the facility's explicit mandate for intensive AI training. In stark contrast to the regional socialization observed in the TVA, the Stargate simulation reveals a catastrophic and highly localized welfare penalty. The terminal Consumption Equivalent Variation (CEV) violently diverges, plummeting to -10.5% for the host counties compared to a marginal -0.137% for the broader ERCOT ratepayer base, translating to an overwhelming annual burden of \$8,855.63 per local household.

This extreme spatial isolation is the direct consequence of nodal pricing with a deregulated retail environment. Lacking a utility aggregation apparatus to pool and diffuse costs, the computing load completely saturates local transmission capacity. Consequently, the host region's LMP surges by 16.59 cents per kWh, while the rest of the grid absorbs only a negligible 0.13 cent systemic increase. This wholesale volatility directly transmits to local industry, triggering a severe factor price displacement effect that erodes the real wage by 8.3%. Faced with astronomical energy input costs and collapsing household incomes, the local economy is forced into a drastic contraction.

Collectively, these case studies demonstrate that the macroeconomic incidence of large-scale computing infrastructure is inextricably linked to the underlying electricity market architecture. In regulated environments, the capital burdens of generation and transmission expansion are systematically socialized, resulting in a diffuse but inescapable erosion of regional welfare. Conversely, the unbundled nature of deregulated wholesale markets traps the physical frictions of massive load shocks within host municipalities, threatening severe factor price displacement and household consumption loss.

However, the magnitude of these projected impacts underscores an inherent limitation of the structural framework. The simulated localized penalty is highly sensitive to the nodal congestion parameter, κ. Because bespoke industrial power purchase agreements and granular locational marginal prices elude the URDB data, the model identifies this transmission friction by calibrating κ to directly match the targeted empirical moment of the manufacturing wage penalty. Therefore, catastrophic welfare losses forecasted in isolated deregulated environments may represent an analytical upper bound rather than a definitive trajectory.

8 Conclusion

As the race toward artificial general intelligence accelerates, technology companies are investing hundreds of billions of dollars into unprecedented computing infrastructure. These gigawatt-scale facilities place intense strain on the electrical grid, necessitating massive investments to accommodate sudden influxes of highly inelastic energy demand. Consequently, as policymakers debate zoning reforms and federal moratoriums, there is an urgent need to rigorously understand the localized macroeconomic externalities these data centers impose on host communities.

Reduced-form empirical analysis reveals that a 100 MW capacity shock increases host-county residential fixed charges by 0.74% in regulated states, while balancing authority spillovers raise regional fixed charges by an additional 0.06%. Because the Utility Rate Database omits direct pass-through mechanisms like power and transmission cost adjustments, these base tariff effects likely represent a lower bound of the true financial incidence. A two-region DSGE model resolves the empirical puzzle of eroding manufacturing wages (-0.49%) amidst stable sector output and employment. The framework demonstrates that elevated local electricity prices trigger a severe factor price displacement effect, forcing energy-intensive firms to pass input cost shocks onto local labor to maintain zero-profit conditions. This wage erosion, combined with surging utility bills, induces a counteracting wealth effect that compels households to increase labor supply to smooth consumption, stabilizing aggregate employment while precipitating a sharp decline in welfare. A baseline 100 MW shock yields a Consumption Equivalent Variation (CEV) of -0.019% in regulated host counties and a profoundly amplified -0.150% in deregulated environments.

While this framework effectively maps channels of macroeconomic incidence, it is subject to limitations. By abstracting from capital accumulation and firm mobility, the model restricts itself to a short-to-medium-run equilibrium. Over a sufficiently long horizon, sustained localized energy premiums would likely induce capital investments or prompt energy-intensive firms to relocate, relaxing the binding constraints of the factor price displacement effect. Furthermore, the extreme spatial divergence observed in deregulated market simulations is highly sensitive to the nodal congestion parameter, κ. Because granular locational marginal prices and bespoke industrial power purchase agreements elude the panel data, κ is analytically identified by directly matching the empirical manufacturing wage penalty. Lacking a mechanism in the deregulated variant for endogenous transmission expansion to alleviate physical congestion, the catastrophic welfare losses forecasted in isolated deregulated case studies—such as the -10.5% CEV projected for the OpenAI Stargate development—may represent an analytical upper bound rather than a definitive long-run trajectory.

Despite these constraints, the combined analyses establish an unambiguous theme: the macroeconomic incidence of large-scale computing infrastructure is systematically shifted onto local consumers. Whether through rate-based socialization under vertically integrated monopolies or extreme nodal congestion premiums in unbundled markets, the immense expenditures required to sustain AI facilities are ultimately borne by captive households through degraded real wages and compressed purchasing power.

Future research must address current constraints by integrating node-level wholesale pricing data and confidential industrial tariff structures to precisely calibrate spatial transmission frictions. Incorporating endogenous firm mobility and long-run capital accumulation into the general equilibrium framework will also be critical for forecasting the permanent reorganization of host economies. As policymakers and public utility commissions navigate the accelerating proliferation of AI infrastructure, they must actively consider the distributional consequences of grid expansion. Reevaluating cost-allocation frameworks will be essential to ensure that the pursuit of computing supremacy does not inadvertently impoverish the local communities that power it—ensuring that residents, like those in Memphis, are not left to single-handedly bear the unseen costs of the artificial intelligence revolution.

References

[1] Acemoglu, Daron. 2024. "The Simple Macroeconomics of AI." April 5.

[2] Acemoglu, Daron, and Pascual Restrepo. 2018a. "Artificial Intelligence, Automation and Work." *NBER Working Paper Series*.

[3] Acemoglu, Daron, and Pascual Restrepo. 2018b. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108, no. 6: 1488–1542. https://doi.org/10.1257/aer.20160696.

[4] Acemoglu, Daron, and Pascual Restrepo. 2022. "Tasks, Automation, and the Rise in US Wage Inequality." February.

[5] Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones. 2017. "Artificial Intelligence and Economic Growth." *NBER Working Paper Series*.

[6] Arenas Díaz, Guillermo, Mariacristina Piva, and Marco Vivarelli. 2025. *Artificial Intelligence as a Complement to Other Innovation Activities and as a Method of Invention*. UNU-MERIT. https://doi.org/10.53330/AEJJ9309.

[7] Athey, Susan, and Fiona Scott Morton. n.d. "Artificial Intelligence, Competition, and Welfare."

[8] Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. "Artificial Intelligence, Firm Growth, and Product Innovation." *Journal of Financial Economics* 151 (January): 103745. https://doi.org/10.1016/j.jfineco.2023.103745.

[9] Benetton, Matteo, Giovanni Compiani, and Adair Morse. n.d. "When Technology Processors Come to Town: High Electricity Use Spillovers to the Local Economy."

[10] Borenstein, Severin, James Bushnell, and Steven Stoft. 2000. "The Competitive Effects of Transmission Capacity in a Deregulated Electricity Industry." *The RAND Journal of Economics* 31, no. 2: 294–325.

[11] Borenstein, Severin, and Stephen Holland. 2003. *On the Efficiency of Competitive Electricity Markets With Time-Invariant Retail Prices*. Working Paper 9922. National Bureau of Economic Research. https://doi.org/10.3386/w9922.

[12] Bretschger, Lucas, and Ara Jo. 2024. "Complementarity between Labor and Energy: A Firm-Level Analysis." *Journal of Environmental Economics and Management* 124 (March): 102934. https://doi.org/10.1016/j.jeem.2024.102934.

[13] Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2025. "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence."

[14] Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2023. "Generative AI at Work." Working Paper 31161. *NBER Working Paper Series*. National Bureau of Economic Research, April. https://doi.org/10.3386/w31161.

[15] Cockburn, Iain, Rebecca Henderson, and Scott Stern. 2018. "The Impact of Artificial Intelligence on Innovation." *NBER Working Paper Series*, March.

[16] De Vere, Hugo, Srini Ramaswamy, and Seth Searls. 2026. "AI Financing." *Dallas Fed Economics*, February 10. https://www.dallasfed.org/research/economics/2026/0210-searls-aifinancing.

[17] Dell'Acqua, Fabrizio, Charles Ayoubi, Hila Lifshitz-Assaf, et al. 2025. "The Cybernetic Teammate: A Field Experiment on Generative AI Reshaping Teamwork and Expertise." Preprint, *SSRN*. https://doi.org/10.2139/ssrn.5188231.

[18] Dell'Acqua, Fabrizio, Edward McFowland, Ethan R. Mollick, et al. 2023. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." *SSRN Electronic Journal*, ahead of print. https://doi.org/10.2139/ssrn.4573321.

[19] Demirer, Mert, John J. Horton, Nicole Immorlica, Brendan Lucier, and Peyman Shahidi. n.d. *NBER Working Paper Series*.

[20] U.S. Energy Information Administration (EIA). n.d. "Electricity Explained - Data and Statistics." Accessed March 24, 2026. https://www.eia.gov/energyexplained/electricity/data-and-statistics.php.

[21] U.S. Energy Information Administration (EIA). n.d. "Frequently Asked Questions (FAQs)." Accessed March 24, 2026. https://www.eia.gov/tools/faqs/faq.php?id=97&t=3.

[22] U.S. Energy Information Administration (EIA). n.d. "Prices and Factors Affecting Prices." Accessed March 24, 2026. https://www.eia.gov/energyexplained/electricity/prices-and-factors-affecting-prices.php.

[23] Elminejad, Ali, Tomas Havranek, Roman Horvath, and Zuzana Irsova. 2023. "Intertemporal Substitution in Labor Supply: A Meta-Analysis." *Review of Economic Dynamics* 51 (December): 1095–1113. https://doi.org/10.1016/j.red.2023.10.001.

[24] Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs Are GPTs: Labor Market Impact Potential of LLMs." *Science* 384, no. 6702: 1306–1308. https://doi.org/10.1126/science.adj0998.

[25] Gans, Joshua S. n.d. *NBER Working Paper Series*.

[26] Hassler, John, Per Krusell, and Conny Olovsson. 2012. *Energy-Saving Technical Change*. Working Paper 18456. National Bureau of Economic Research. https://doi.org/10.3386/w18456.

[27] Hui, Xiang, Oren Reshef, and Luofeng Zhou. 2024. "The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market." October 12.

[28] Humlum, Anders, and Emilie Vestergaard. 2025. *Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*. Working Paper 33777. National Bureau of Economic Research. https://doi.org/10.3386/w33777.

[29] Imbens, Guido W., and Donald B. Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.

[30] Jones, Charles I. n.d. "A.I. and Our Economic Future."

[31] Kay, Owen, Robert Reaser, and Reid Taylor. 2026. "Processing Power: The Effect of Data Centers on Wholesale Electricity Markets." *Federal Reserve Bank of Dallas, Working Papers* 2026, no. 2606. https://doi.org/10.24149/wp2606.

[32] King, Robert G., Charles I. Plosser, and Sergio T. Rebelo. 1988. "Production, Growth and Business Cycles." *Journal of Monetary Economics* 21, no. 2–3: 195–232. https://doi.org/10.1016/0304-3932(88)90030-X.

[33] Knittel, Christopher, Juan Ramon Senga, and Shen Wang. 2025. *Flexible Data Centers and the Grid: Lower Costs, Higher Emissions?* Working Paper 34065. National Bureau of Economic Research. https://doi.org/10.3386/w34065.

[34] Lazard. 2025. *Levelized Cost of Energy+ (LCOE+)*. Lazard.

[35] Mamkhezri, Jamal. 2025. "Assessing Price Elasticity in US Residential Electricity Consumption: A Comparison of Monthly and Annual Data with Recession Implications." *Energy Policy* 200 (May): 114537. https://doi.org/10.1016/j.enpol.2025.114537.

[36] Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence."

[37] Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." Preprint, *arXiv*, February 13. https://doi.org/10.48550/arXiv.2302.06590.

[38] Restrepo, Pascual. 2025. "We Won't Be Missed: Work and Growth in the AGI World." In *The Economics of Transformative AI*. University of Chicago Press. https://www.nber.org/books-and-chapters/economics-transformative-ai/we-wont-be-missed-work-and-growth-agi-world.

[39] Shehabi, Arman, Sarah Smith, Dale Sartor, et al. 2016. *United States Data Center Energy Usage Report*. LBNL–1005775, 1372902. https://doi.org/10.2172/1372902.

A Common Support Validation

Figure 11: Distribution of Treated Counties across Baseline (2014) Covariate Quintiles

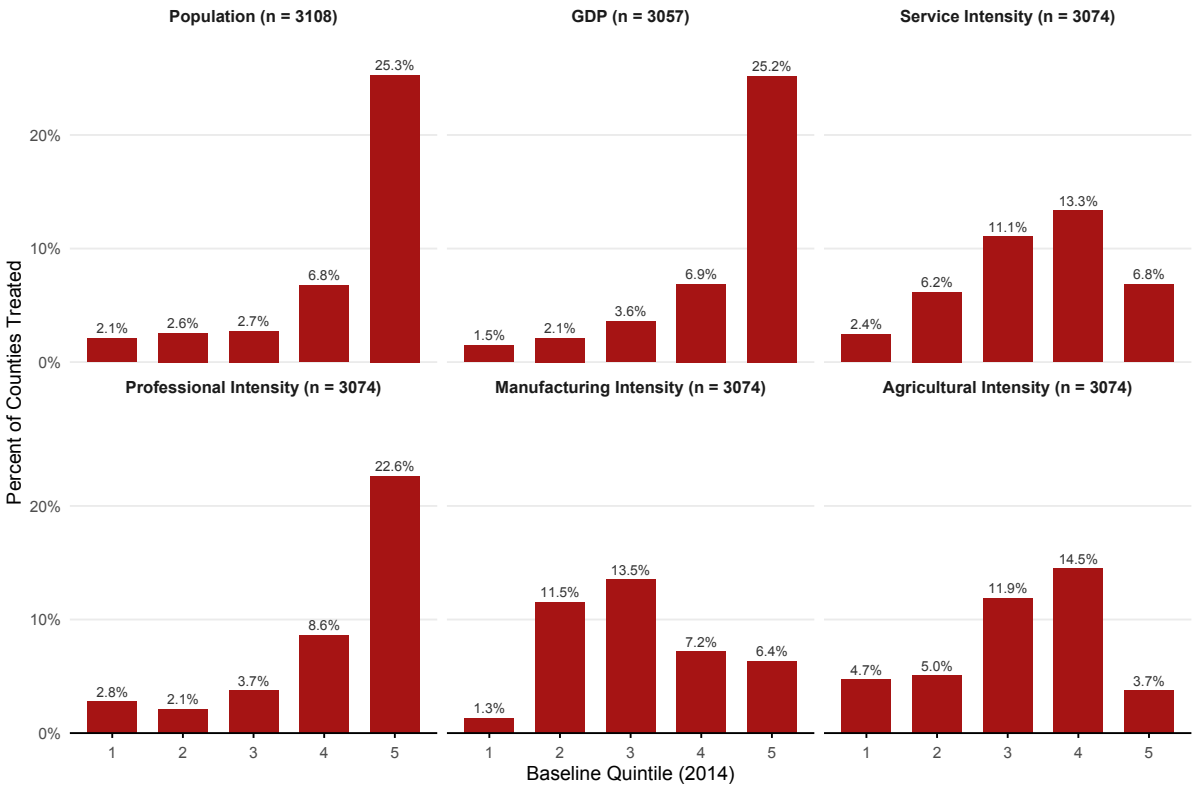

The introduction of covariate-specific time trends via quintile-by-year fixed effects rests on the critical assumption of common support. To validly identify the causal effect of data center expansion conditional on baseline characteristics, the panel must contain a sufficient number of both treated and untreated counties within every specified bin. If a particular quintile (e.g., the top 20% of populous counties) consists entirely of treated observations, the model lacks an appropriate control group to construct the counterfactual trend for that specific sub-population. In such a scenario, the quintile-by-year fixed effect would perfectly collinearize with the treatment, either absorbing the treatment effect entirely or rendering the conditional parallel trends assumption fundamentally untestable.

Figure 11 validates the empirical design by demonstrating robust common support across all interacted baseline covariates. As expected, the distribution of treatment is heavily right-skewed with respect to absolute scale; 25.3% of counties in the top population quintile and 25.2% in the top GDP quintile are treated, compared to less than 3% in the bottom quintiles. A similar skew is evident for baseline professional intensity, consistent with the spatial sorting of computing infrastructure into wealthy, high-tech metropolitan hubs. Conversely, data centers exhibit a higher relative concentration in counties with lower baseline service, manufacturing, and agricultural intensities.

Despite this stark non-random assignment, the critical condition for identification holds: no quintile across any of the six baseline metrics is exclusively composed of treated or untreated observations. Even in the most concentrated demographic bins (e.g., the 5th population quintile), nearly 75% of the counties remain untreated, providing a massive, structurally comparable control cohort to accurately estimate the local, time-varying macroeconomic trend.

B Economic Specification Regression Tables

Table 10: Effect of Data Center Load on Local Aggregate Economic Outcomes

	Dependent Variable					
	(1) log(GDP)	(2) log(Pop)	(3) log(Labor F.)	(4) log(Emp)	(5) log(Unemp)	(6) Unemp Rate
Demand	0.0078**	0.0044***	0.0056***	0.0054***	0.0069**	0.0098
	(0.0036)	(0.0010)	(0.0016)	(0.0016)	(0.0033)	(0.0079)
Placebo	0.0050	0.0024	0.0072*	0.0067	0.0153	0.0476
	(0.0114)	(0.0033)	(0.0041)	(0.0041)	(0.0157)	(0.0935)
Observations	33,132	33,132	33,132	33,132	33,132	33,132
$\bar{X}$ (no log)	1,326,428	29,911	13,673	13,019	605	4.761
R^2	0.996	0.999	0.999	0.999	0.994	0.870
Within R^2	0.004	0.008	0.005	0.004	0.008	0.004

Notes: Standard errors clustered by 2020 Commuting Zone are reported in parentheses. Significance levels are denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All specifications include county fixed effects (FEs), year FEs, quintile-by-year FEs, and weather controls via annual CDD/HDD. Demand is expressed in units of 100 MW. Placebo is an indicator variable that takes the value of 1 for the three years prior to a treated county crossing the 10 MW threshold; a precise null effect serves as a falsification test to rule out reverse causality. Quintile-by-Year FEs encompass baseline (2014) characteristics for population, GDP, and sector intensities (professional, manufacturing, agriculture, and service) interacted with year indicators.

Table 11: Effect of Data Center Load on Local Sectoral Output (GDP)

	Dependent Variable: log(GDP)						
	(1) Aggregate	(2) Mfg	(3) Prof	(4) Agri	(5) Constr	(6) Services	(7) Info
Demand	0.0078**	-0.0013	0.0097*	0.0015	0.0168**	0.0007	0.0359*
	(0.0036)	(0.0054)	(0.0053)	(0.0107)	(0.0071)	(0.0018)	(0.0187)
Placebo	0.0050	-0.0380*	0.0230	0.0615	0.0150	0.0075	-0.0022
	(0.0114)	(0.0210)	(0.0179)	(0.0517)	(0.0200)	(0.0081)	(0.0414)
Observations	33,132	29,815	28,644	28,553	29,970	30,412	26,266
$\bar{X}$ (no log)	1,326,428	140,928	64,602	24,343	62,318	39,144	21,363
R^2	0.996	0.986	0.994	0.881	0.984	0.996	0.977
Within R^2	0.004	0.002	0.002	0.004	0.003	0.005	0.004

Notes: Standard errors clustered by 2020 Commuting Zone are reported in parentheses. Significance levels are denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All specifications include county fixed effects (FEs), year FEs, quintile-by-year FEs, and weather controls via annual CDD/HDD. Demand is expressed in units of 100 MW. Placebo is an indicator variable that takes the value of 1 for the three years prior to a treated county crossing the 10 MW threshold; a precise null effect serves as a falsification test to rule out reverse causality. Quintile-by-Year FEs encompass baseline (2014) characteristics for population, GDP, and sector intensities (professional, manufacturing, agriculture, and service) interacted with year indicators.

Table 12: Effect of Data Center Load on Local Sectoral Employment

	Dependent Variable: log(Employment)						
	(1) Aggregate	(2) DCs	(3) Mfg	(4) Prof	(5) Agri	(6) Constr	(7) Services
Demand	0.0098***	-0.0238	-0.0127	0.0163***	-0.0203**	0.0053	0.0047**
	(0.0033)	(0.0203)	(0.0082)	(0.0060)	(0.0101)	(0.0061)	(0.0020)
Placebo	0.0087	0.0872	-0.0243	0.0267	-0.0073	0.0227	0.0222*
	(0.0111)	(0.1106)	(0.0395)	(0.0251)	(0.0602)	(0.0204)	(0.0133)
Observations	30,120	30,120	30,120	30,120	30,120	30,120	30,120
$\bar{X}$ (no log)	7,407	3	531	381	7	328	1,125
R^2	0.997	0.899	0.933	0.975	0.874	0.942	0.984
Within R^2	0.002	0.000	0.001	0.000	0.001	0.000	0.001

Notes: Standard errors clustered by 2020 Commuting Zone are reported in parentheses. Significance levels are denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All specifications include county fixed effects (FEs), year FEs, quintile-by-year FEs, and weather controls via annual CDD/HDD. Demand is expressed in units of 100 MW. Placebo is an indicator variable that takes the value of 1 for the three years prior to a treated county crossing the 10 MW threshold; a precise null effect serves as a falsification test to rule out reverse causality. Quintile-by-Year FEs encompass baseline (2014) characteristics for population, GDP, and sector intensities (professional, manufacturing, agriculture, and service) interacted with year indicators.

Table 13: Effect of Data Center Load on Local Sectoral Derived Wages

	Dependent Variable: log(Wage)						
	(1) Aggregate	(2) DC	(3) Mfg	(4) Prof	(5) Agri	(6) Constr	(7) Services
Demand	-0.0003	-0.0046	-0.0049**	0.0001	-0.0032	0.0029	0.0017
	(0.0018)	(0.0053)	(0.0020)	(0.0034)	(0.0043)	(0.0020)	(0.0013)
Placebo	-0.0062	0.0440	-0.0015	-0.0083	-0.0046	0.0060	-0.0129*
	(0.0051)	(0.0442)	(0.0089)	(0.0110)	(0.0189)	(0.0090)	(0.0066)
Observations	30,118	4,575	26,842	28,532	15,869	28,990	29,866
$\bar{X}$ (no log)	40,014	70,615	49,217	42,066	40,823	49,168	18,886
R^2	0.945	0.760	0.866	0.791	0.745	0.854	0.877
Within R^2	0.000	0.001	0.001	0.000	0.000	0.001	0.001

Notes: Standard errors clustered by 2020 Commuting Zone are reported in parentheses. Significance levels are denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All specifications include county fixed effects (FEs), year FEs, quintile-by-year FEs, and weather controls via annual CDD/HDD. Demand is expressed in units of 100 MW. Placebo is an indicator variable that takes the value of 1 for the three years prior to a treated county crossing the 10 MW threshold; a precise null effect serves as a falsification test to rule out reverse causality. Quintile-by-Year FEs encompass baseline (2014) characteristics for population, GDP, and sector intensities (professional, manufacturing, agriculture, and service) interacted with year indicators.

Table 14: Effect of Data Center Load on Local Welfare and Housing

	Dependent Variable				
	(1) log(Med HHI)	(2) log(Poverty)	(3) log(SNAP)	(4) log(ZHVI)	(5) log(ZORI)
Demand	0.0020	0.0043	0.0041	0.0024	-0.0013
	(0.0015)	(0.0028)	(0.0030)	(0.0022)	(0.0009)
Placebo	-0.0091*	-0.0020	-0.0106	0.0060	0.0028
	(0.0052)	(0.0100)	(0.0134)	(0.0096)	(0.0061)
Observations	33,128	33,131	33,132	34,161	6,900
$\bar{X}$ (no log)	52,680	4,083	1,337	169,058	1,296
R^2	0.965	0.992	0.987	0.975	0.985
Within R^2	0.003	0.001	0.001	0.017	0.015

Notes: Standard errors clustered by 2020 Commuting Zone are reported in parentheses. Significance levels are denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All specifications include county fixed effects (FEs), year FEs, quintile-by-year FEs, and weather controls via annual CDD/HDD. Demand is expressed in units of 100 MW. Placebo is an indicator variable that takes the value of 1 for the three years prior to a treated county crossing the 10 MW threshold; a precise null effect serves as a falsification test to rule out reverse causality. Quintile-by-Year FEs encompass baseline (2014) characteristics for population, GDP, and sector intensities (professional, manufacturing, agriculture, and service) interacted with year indicators.

C Household Problem Derivation

To fully characterize the optimal behavior of the representative household in region i, I set up the infinite-horizon Lagrangian. Let λ_{it} denote the sequence of Lagrange multipliers associated with the period-by-period flow budget constraint. The household's Lagrangian is formulated as:

$$\mathcal{L} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \Bigg\{ \ln(C_{it}) + \nu \cdot \frac{(E_{it}^H)^{1-\sigma_E}}{1-\sigma_E} - \chi \cdot \frac{L_{it}^{1+\varphi}}{1+\varphi} + \lambda_{it} \left(w_{it} L_{it} + B_{it} + \Pi_{it} - C_{it} - p_{it}^H E_{it}^H - F_{it}^H - \frac{B_{i,t+1}}{1+r_t} \right) \Bigg\}.$$

I derive the first-order conditions (FOCs) by differentiating the Lagrangian with respect to consumption C_{it}, household energy demand E_{it}^H, labor supply L_{it}, and bond holdings $B_{i,t+1}$.

First, the FOC with respect to consumption yields the shadow value of wealth:

$$\frac{\partial \mathcal{L}}{\partial C_{it}} = \beta^t \left(\frac{1}{C_{it}} - \lambda_{it} \right) = 0 \implies \lambda_{it} = \frac{1}{C_{it}}.$$

Second, the FOC with respect to labor supply equates the marginal disutility of labor to the marginal wage benefit:

$$\frac{\partial \mathcal{L}}{\partial L_{it}} = \beta^t \left(-\chi L_{it}^{\varphi} + \lambda_{it} w_{it} \right) = 0 \implies \chi L_{it}^{\varphi} = \lambda_{it} w_{it}.$$

Third, the FOC with respect to household energy demand equates the marginal utility of energy to its shadow cost:

$$\frac{\partial \mathcal{L}}{\partial E_{it}^H} = \beta^t \left(\nu (E_{it}^H)^{-\sigma_E} - \lambda_{it} p_{it}^H \right) = 0 \implies \nu (E_{it}^H)^{-\sigma_E} = \lambda_{it} p_{it}^H.$$

Fourth, the FOC with respect to one-period bond holdings yields the intertemporal condition linking sequential multipliers:

$$\frac{\partial \mathcal{L}}{\partial B_{i,t+1}} = -\beta^t \lambda_{it} \frac{1}{1+r_t} + \mathbb{E}_t \left[\beta^{t+1} \lambda_{i,t+1} \right] = 0 \implies \lambda_{it} = \beta (1+r_t) \mathbb{E}_t [\lambda_{i,t+1}].$$

Substituting the marginal utility of consumption ($\lambda_{it} = 1/C_{it}$) into the bond FOC yields the standard intertemporal Euler equation:

$$\frac{1}{C_{it}} = \beta (1+r_t) \mathbb{E}_t \left[\frac{1}{C_{i,t+1}} \right].$$

Substituting the marginal utility of consumption into the labor and energy FOCs yields the intratemporal optimality conditions for labor supply and energy demand:

$$\chi L_{it}^{\varphi} = \frac{w_{it}}{C_{it}},$$
$$\nu (E_{it}^H)^{-\sigma_E} = \frac{p_{it}^H}{C_{it}}.$$

Finally, I impose the transversality condition (TVC) to ensure that the household does not indefinitely accumulate wealth that is never consumed. The TVC dictates that the discounted

present value of the household's asymptotic bond holdings must equal zero:

$$\lim_{t\to\infty} \mathbb{E}_0 \left[\beta^t \lambda_{it} \frac{B_{i,t+1}}{1+r_t} \right] = 0.$$

Substituting the shadow price of consumption into this limit provides the rigorous boundary condition for the household's infinite-horizon problem:

$$\lim_{t\to\infty} \mathbb{E}_0 \left[\beta^t \left(\frac{1}{C_{it}} \right) \frac{B_{i,t+1}}{1+r_t} \right] = 0.$$

D Firm Problem Derivation

To solve the representative manufacturing firm's problem, I first derive the intratemporal optimality conditions by maximizing the static period-by-period profit function. The firm chooses labor L_{it} and energy E_{it}^M to maximize:

$$\max_{L_{it},E_{it}^M} \Pi_{it} = A\left[\alpha(E_{it}^M)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)L_{it}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} - w_{it}L_{it} - p_{it}^M E_{it}^M - F_{it}^M.$$

Taking the partial derivative of the profit function with respect to labor yields:

$$\frac{\partial \Pi_{it}}{\partial L_{it}} = \left(\frac{\sigma}{\sigma-1}\right) A\left[\alpha(E_{it}^M)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)L_{it}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{1}{\sigma-1}} (1-\alpha)\left(\frac{\sigma-1}{\sigma}\right) L_{it}^{-\frac{1}{\sigma}} - w_{it} = 0.$$

To simplify this expression, I isolate the core CES bracket from the original production function, noting that $Y_{it}^{\frac{1}{\sigma}} = A^{\frac{1}{\sigma}}[\cdot]^{\frac{1}{\sigma-1}}$. Substituting this relationship back into the derivative yields the intratemporal labor demand condition:

$$(1-\alpha)A^{\frac{\sigma-1}{\sigma}}\left(\frac{Y_{it}}{L_{it}}\right)^{\frac{1}{\sigma}} = w_{it}.$$

By symmetry, taking the partial derivative with respect to energy demand yields the intratemporal energy demand condition:

$$\frac{\partial \Pi_{it}}{\partial E_{it}^M} = 0 \implies \alpha A^{\frac{\sigma-1}{\sigma}}\left(\frac{Y_{it}}{E_{it}^M}\right)^{\frac{1}{\sigma}} = p_{it}^M.$$

To derive the factor price frontier, I formulate the firm's dual cost-minimization problem. The firm minimizes variable costs $VC_{it} = w_{it}L_{it} + p_{it}^M E_{it}^M$ subject to producing a given level of output Y_{it}. Dividing the energy demand condition by the labor demand condition yields the firm's optimal input ratio:

$$\frac{p_{it}^M}{w_{it}} = \frac{\alpha}{1-\alpha}\left(\frac{L_{it}}{E_{it}^M}\right)^{\frac{1}{\sigma}} \implies E_{it}^M = L_{it}\left(\frac{\alpha}{1-\alpha}\frac{w_{it}}{p_{it}^M}\right)^{\sigma}.$$

Substituting this optimal ratio back into the CES production function allows me to express output entirely as a function of labor demand and input prices:

$$Y_{it} = A\left[\alpha\left(L_{it}\left(\frac{\alpha}{1-\alpha}\frac{w_{it}}{p_{it}^M}\right)^{\sigma}\right)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)L_{it}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}$$

$$Y_{it} = AL_{it}\left[\alpha\left(\frac{\alpha}{1-\alpha}\frac{w_{it}}{p_{it}^M}\right)^{\sigma-1} + (1-\alpha)\right]^{\frac{\sigma}{\sigma-1}}.$$

To algebraically simplify the bracketed term, I factor out the wage and denominator components:

$$Y_{it} = AL_{it}\left[\frac{\alpha^{\sigma}(p_{it}^M)^{1-\sigma}w_{it}^{\sigma-1} + (1-\alpha)^{\sigma}}{(1-\alpha)^{\sigma-1}}\right]^{\frac{\sigma}{\sigma-1}}$$

$$Y_{it} = AL_{it}\frac{w_{it}^{\sigma}}{(1-\alpha)^{\sigma}}\left[\alpha^{\sigma}(p_{it}^M)^{1-\sigma} + (1-\alpha)^{\sigma}w_{it}^{1-\sigma}\right]^{\frac{\sigma}{\sigma-1}}.$$

Rearranging to solve for the optimized labor demand L_{it}^* gives:

$$L_{it}^* = \frac{Y_{it}}{A}\left(\frac{1-\alpha}{w_{it}}\right)^{\sigma}\left[(1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right]^{\frac{\sigma}{1-\sigma}}.$$

By symmetry, the optimized energy demand E_{it}^{M*} is:

$$E_{it}^{M*} = \frac{Y_{it}}{A}\left(\frac{\alpha}{p_{it}^M}\right)^{\sigma}\left[(1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right]^{\frac{\sigma}{1-\sigma}}.$$

Substituting both the optimized labor and energy demands into the variable cost equation yields the total variable cost as a function of output and input prices:

$$\begin{aligned}
VC_{it} &= w_{it}L_{it}^* + p_{it}^M E_{it}^{M*} \\
VC_{it} &= \frac{Y_{it}}{A}\left[(1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right]^{\frac{\sigma}{1-\sigma}}\left(w_{it}\left(\frac{1-\alpha}{w_{it}}\right)^{\sigma}+p_{it}^M\left(\frac{\alpha}{p_{it}^M}\right)^{\sigma}\right) \\
VC_{it} &= \frac{Y_{it}}{A}\left[(1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right]^{\frac{\sigma}{1-\sigma}}\left((1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right).
\end{aligned}$$

Combining the terms with the same base simplifies the total variable cost to:

$$VC_{it} = Y_{it}\cdot\frac{1}{A}\left[(1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}.$$

Dividing total variable cost by Y_{it} isolates the unit cost function $c(w_{it}, p_{it}^M)$:

$$c(w_{it}, p_{it}^M) = \frac{1}{A}\left[(1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}.$$

Because the manufacturing firm operates in a perfectly competitive market under constant returns to scale, economic profit is driven to zero in spatial equilibrium. Consequently, the price of the composite good (normalized to 1 as the numeraire) must equal the marginal cost of production, which is exactly the unit cost ($MC_{it} = c(w_{it}, p_{it}^M) = 1$). This equality formally establishes the region-specific factor price frontier:

$$1 = \frac{1}{A}\left[(1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}.$$

D.1 Proof of Shephard's Lemma

Finally, to rigorously support the total differential log-linearization utilized in Proposition 2, I explicitly derive Shephard's Lemma for this unit cost function. Taking the partial derivative of $c(w_{it}, p_{it}^M)$ with respect to the real wage w_{it} yields:

$$\begin{aligned}
\frac{\partial c(w_{it}, p_{it}^M)}{\partial w_{it}} &= \frac{1}{A}\left(\frac{1}{1-\sigma}\right)\left[(1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right]^{\frac{1}{1-\sigma}-1}(1-\sigma)(1-\alpha)^{\sigma}w_{it}^{-\sigma} \\
&= \frac{1}{A}\left[(1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right]^{\frac{\sigma}{1-\sigma}}(1-\alpha)^{\sigma}w_{it}^{-\sigma}.
\end{aligned}$$

Rearranging this derivative, I can express it precisely as the ratio of the optimized labor demand L_{it}^* to output Y_{it}:

$$\frac{\partial c(w_{it}, p_{it}^M)}{\partial w_{it}} = \frac{1}{A}\left(\frac{1-\alpha}{w_{it}}\right)^{\sigma}\left[(1-\alpha)^{\sigma}w_{it}^{1-\sigma}+\alpha^{\sigma}(p_{it}^M)^{1-\sigma}\right]^{\frac{\sigma}{1-\sigma}} = \frac{L_{it}}{Y_{it}}.$$

A symmetric differentiation with respect to the industrial electricity price p_{it}^M yields the conditional energy demand per unit of output:

$$\frac{\partial c(w_{it}, p_{it}^M)}{\partial p_{it}^M} = \frac{E_{it}^M}{Y_{it}}.$$

This confirms that the CES unit cost function fully encapsulates the firm's optimal factor demands, validating the structural identity necessary to establish the local wage penalty mechanism.

E Complete Proofs of DSGE Model Propositions

E.1 Proposition 1

P1: An exogenous increase in localized data center power demand D_t strictly increases the variable rate of electricity p_{1t} for region 1's representative household and manufacturing firm, regardless of the governing regulatory architecture.

Proof. The proof follows directly from the region-specific variable electricity price equations established in the model setup. Because data center energy demand E_t^{DC} scales linearly with peak power demand D_t, I evaluate the comparative statics with respect to E_t^{DC} under each regulatory regime.

Under $\Omega = 0$, the variable rate of electricity is determined by the average variable cost of the unified generation supply:

$$p_{1t} = v_0 + \frac{v_1}{2}(E_{1t}^H + E_{2t}^H + E_{1t}^M + E_{2t}^M + E_t^{DC}).$$

Taking the partial derivative with respect to the exogenous data center load gives:

$$\frac{\partial p_{1t}}{\partial E_t^{DC}} = \frac{v_1}{2} > 0,$$

which is strictly positive given the convex variable cost of generation, $v_1 > 0$.

Under $\Omega = 1$, variable rates for the household and firm are anchored to the LMP in region 1:

$$LMP_{1t} = v_0 + v_1(E_{1t}^H + E_{2t}^H + E_{1t}^M + E_{2t}^M + E_t^{DC}) + \kappa E_t^{DC}.$$

Taking the partial derivative with respect to the exogenous load gives:

$$\frac{\partial LMP_{1t}}{\partial E_t^{DC}} = v_1 + \kappa > 0,$$

which is strictly positive given $v_1 > 0$ and $\kappa > 0$.

By defining the log-linearized variable rate of region 1 as:

$$\hat{p}_{1t} \equiv \ln\left(\frac{p_{1t}}{\bar{p}_1}\right) \approx \frac{dp_{1t}}{\bar{p}_1},$$

where $\bar{p}_1$ is the steady-state variable rate prior to the load shock, the positive partial derivatives confirm that the localized price shock is strictly positive across both regulatory regimes:

$$\hat{p}_{1t} > 0.$$

□

E.2 Proposition 2

P2: Under a perfectly competitive market and constant returns to scale, an exogenous increase in the localized firm electricity price p_{it}^M strictly depresses the local real wage w_{it}.

Proof. Because the manufacturing sector operates under perfect competition and constant returns to scale (CRS), the firm earns zero economic profit. The price of the output composite good, normalized to 1 as the numeraire, must exactly equal the marginal cost of production. Equivalently, the unit cost function $C(w_{it}, p_{it}^M)$ is fixed at 1. Taking the total differential of the unit cost function yields:

$$dC(w_{it}, p_{it}^M) = \frac{\partial C}{\partial w_{it}} dw_{it} + \frac{\partial C}{\partial p_{it}^M} dp_{it}^M = 0.$$

By Shephard's Lemma (see Appendix D for proof), the partial derivative of the unit cost function with respect to an input price equals the conditional demand for that input required to produce one unit of output. Thus,

$$\frac{\partial C}{\partial w_{it}} = \frac{L_{it}}{Y_{it}}, \quad \frac{\partial C}{\partial p_{it}^M} = \frac{E_{it}^M}{Y_{it}}.$$

Substituting these conditions into the differential equation yields:

$$\left(\frac{L_{it}}{Y_{it}}\right) dw_{it} + \left(\frac{E_{it}^M}{Y_{it}}\right) dp_{it}^M = 0.$$

To express this relationship in terms of percentage deviations from the steady state, I multiply each term by its respective price-to-price ratio (e.g., w_{it}/w_{it}) to log-linearize the equation:

$$\left(\frac{w_{it} L_{it}}{Y_{it}}\right) \hat{w}_{it} + \left(\frac{p_{it}^M E_{it}^M}{Y_{it}}\right) \hat{p}_{it}^M = 0.$$

Let the steady-state factor shares of labor and energy be represented by:

$$s_L \equiv \frac{w_{it} L_{it}}{Y_{it}}, \quad s_E \equiv \frac{p_{it}^M E_{it}^M}{Y_{it}},$$

respectively. The relationship simplifies to:

$$\hat{w}_{it} = -\frac{s_E}{s_L} \hat{p}_{it}^M.$$

By Proposition 1, the exogenous data center shock strictly increases the variable price of electricity for the host region ($\hat{p}_{1t}^M > 0$). Because the steady-state factor shares of labor and energy are strictly positive, the region-specific real wage must strictly decrease ($\hat{w}_{1t} < 0$).

□

E.3 Proposition 3

P3: Following the data center demand shock, the steady-state consumption of the composite good by region 1's representative household strictly decreases.

Proof. By way of contradiction, I assume the opposite is true: following the exogenous data center shock, household consumption remains constant or increases ($\hat{C}_{1t} \geq 0$).

From the household's intratemporal labor supply condition, the optimal labor supply is given by $\chi L_{1t}^{\varphi} = w_{1t}/C_{1t}$. Log-linearizing this condition yields:

$$\varphi \hat{L}_{1t} = \hat{w}_{1t} - \hat{C}_{1t}.$$

By Proposition 2, the energy price shock strictly depresses the local real wage ($\hat{w}_{1t} < 0$). Under my assumption that $\hat{C}_{1t} \geq 0$, the right-hand side of the labor supply equation is strictly negative. Because the inverse Frisch elasticity is positive ($\varphi > 0$), this implies that labor supply must strictly decrease: $\hat{L}_{1t} < 0$.

Now, I evaluate the representative household's steady-state budget constraint. In the steady state, bond holdings are constant ($B_t = B_{t+1}$), and the perfectly competitive firm earns zero profit ($\Pi_t = 0$). The budget constraint simplifies to:

$$\bar{C}_1 + \bar{p}_1^H \bar{E}_1^H + \bar{F}_1^H = \bar{w}_1 \bar{L}_1.$$

Taking the total differential and log-linearizing around the steady state yields:

$$\bar{C}_1 \hat{C}_{1t} + \bar{p}_1^H \bar{E}_1^H (\hat{p}_{1t}^H + \hat{E}_{1t}^H) + \bar{F}_1^H \hat{F}_{1t}^H = \bar{w}_1 \bar{L}_1 (\hat{w}_{1t} + \hat{L}_{1t}).$$

Rearranging to isolate the change in consumption expenditure gives:

$$\bar{C}_1 \hat{C}_{1t} = \bar{w}_1 \bar{L}_1 (\hat{w}_{1t} + \hat{L}_{1t}) - \bar{p}_1^H \bar{E}_1^H (\hat{p}_{1t}^H + \hat{E}_{1t}^H) - \bar{F}_1^H \hat{F}_{1t}^H.$$

I now analyze the signs of the components on the right-hand side. The income term ($\hat{w}_{1t} + \hat{L}_{1t}$) is strictly negative, as both the real wage (by P2) and labor supply (by my initial assumption) have declined. The variable energy cost term ($\hat{p}_{1t}^H + \hat{E}_{1t}^H$) represents expenditure. The net effect on expenditure is dominated by the positive price shock $\hat{p}_{1t}^H > 0$ (by P1). Finally, the fixed charge shock is strictly non-negative ($\hat{F}_{1t}^H \geq 0$), representing either the socialized capacity costs under $\Omega = 0$ or a constant delivery charge under $\Omega = 1$.

Therefore, the entire right-hand side of the log-linearized budget constraint is strictly negative. This requires $\bar{C}_1 \hat{C}_{1t} < 0$. Because steady-state consumption $\bar{C}_1$ is strictly positive, it must be that $\hat{C}_{1t} < 0$. This contradicts my initial assumption that $\hat{C}_{1t} \geq 0$. Thus, steady-state consumption must strictly decrease.

□

E.4 Proposition 4

P4: Region 1's labor supply response $\hat{L}_{1t}$ to a data center demand shock is theoretically ambiguous, governed by the counteraction between a household wealth effect and a firm factor price displacement effect.

Proof. I begin with the representative household's log-linearized intratemporal labor supply condition, derived from the King-Plosser-Rebelo (KPR) preferences:

$$\varphi \hat{L}_{1t} = \hat{w}_{1t} - \hat{C}_{1t},$$

where $\varphi > 0$ is the inverse Frisch elasticity of labor supply. Rearranging this condition yields the region's inverse labor supply curve:

$$\hat{w}^S_{1t} = \varphi \hat{L}_{1t} + \hat{C}_{1t}.$$

From Proposition 2, the region-specific real wage mandated by the manufacturing firm's zero-profit condition acts as the inverse labor demand curve:

$$\hat{w}^D_{1t} = -\frac{s_E}{s_L}\hat{p}^M_{1t}.$$

In spatial equilibrium, the local labor market must clear, such that $\hat{w}^S_{1t} = \hat{w}^D_{1t}$. Equating the supply and demand conditions yields:

$$\varphi \hat{L}_{1t} + \hat{C}_{1t} = -\frac{s_E}{s_L}\hat{p}^M_{1t}.$$

Solving for the equilibrium labor supply response $\hat{L}_{1t}$ gives:

$$\hat{L}_{1t} = \frac{1}{\varphi}\left(-\frac{s_E}{s_L}\hat{p}^M_{1t} - \hat{C}_{1t}\right).$$

The theoretical ambiguity of $\hat{L}_{1t}$ is established by analyzing the signs of the two terms on the right-hand side. The first term inside the parentheses represents the firm factor price displacement effect. By Proposition 1, the exogenous data center shock strictly increases the localized industrial electricity price ($\hat{p}^M_{1t} > 0$). Because the steady-state factor shares s_E and s_L are strictly positive, this entire term is strictly negative:

$$-\frac{s_E}{s_L}\hat{p}^M_{1t} < 0.$$

The second term inside the parentheses represents the household wealth effect. By Proposition 3, the steady-state consumption of the composite good strictly decreases following the shock ($\hat{C}_{1t} < 0$). Therefore, the negative of this deviation is strictly positive:

$$-\hat{C}_{1t} > 0.$$

Because the equilibrium labor response $\hat{L}_{1t}$ is proportional to the sum of a strictly negative displacement effect and a strictly positive wealth effect, its net sign is theoretically ambiguous and depends entirely on the relative magnitudes of the local energy price penalty and the resulting household consumption contraction.

□

E.5 Proposition 5

P5: The effect of a data center demand shock on region 1's manufacturing output $\hat{Y}_{1t}$ is theoretically ambiguous, though heavily biased downward by a high elasticity of substitution σ.

Proof. I begin by taking the total differential of the CES production technology, which yields:

$$dY_{1t} = \frac{\partial Y}{\partial L_{1t}} dL_{1t} + \frac{\partial Y}{\partial E^M_{1t}} dE^M_{1t}.$$

Under perfect competition, marginal products equal real factor prices:

$$\frac{\partial Y}{\partial L_{1t}} = w_{1t}, \quad \frac{\partial Y}{\partial E^M_{1t}} = p^M_{1t}$$

Dividing by Y_{1t} and log-linearizing around the steady state yields:

$$\hat{Y}_{1t} = s_L \hat{L}_{1t} + s_E \hat{E}^M_{1t},$$

where s_L and s_E are the steady-state factor shares of labor and energy. From the firm's first-order conditions, the optimal input ratio is:

$$\frac{E^M_{1t}}{L_{1t}} = \left(\frac{\alpha}{1-\alpha} \frac{w_{1t}}{p^M_{1t}} \right)^{\sigma}.$$

Log-linearizing this relationship yields:

$$\hat{E}^M_{1t} - \hat{L}_{1t} = -\sigma(\hat{p}^M_{1t} - \hat{w}_{1t}).$$

Substituting the local wage penalty from Proposition 2 ($\hat{w}_{1t} = -\frac{s_E}{s_L}\hat{p}^M_{1t}$) gives:

$$\hat{E}^M_{1t} = \hat{L}_{1t} - \sigma \left(1 + \frac{s_E}{s_L} \right) \hat{p}^M_{1t}.$$

Substituting this expression for $\hat{E}^M_{1t}$ back into the output equation yields:

$$\hat{Y}_{1t} = s_L \hat{L}_{1t} + s_E \left[\hat{L}_{1t} - \sigma \left(1 + \frac{s_E}{s_L} \right) \hat{p}^M_{1t} \right].$$

Because the firm operates under constant returns to scale ($s_L + s_E = 1$), this simplifies to:

$$\hat{Y}_{1t} = \hat{L}_{1t} - \sigma s_E \left(1 + \frac{s_E}{s_L} \right) \hat{p}^M_{1t}.$$

The ambiguity of local manufacturing output $\hat{Y}_{1t}$ stems from the theoretically ambiguous labor supply response $\hat{L}_{1t}$ established in Proposition 4. However, because σ, s_E, s_L, and $\hat{p}^M_{1t}$ are all strictly positive, the second term is strictly negative. This isolates the unambiguous drag on output caused by the firm substituting away from the expensive energy input, demonstrating that the aggregate output response is heavily biased downward by a high elasticity of substitution σ.

□

E.6 Proposition 6

P6: The effect of a data center demand shock on region 1's representative household utility is theoretically ambiguous but heavily biased downward.

Proof. I begin by evaluating the total differential of the representative household's period utility function $U(C_{1t}, E_{1t}^H, L_{1t})$ evaluated at the steady state:

$$dU_{1t} = \frac{\partial U}{\partial C_{1t}} dC_{1t} + \frac{\partial U}{\partial E_{1t}^H} dE_{1t}^H + \frac{\partial U}{\partial L_{1t}} dL_{1t}.$$

Substituting the marginal utilities from the King-Plosser-Rebelo (KPR) functional form yields:

$$dU_{1t} = \left(\frac{1}{\bar{C}_1}\right) dC_{1t} + \left(\nu(\bar{E}_1^H)^{-\sigma_E}\right) dE_{1t}^H - \left(\chi \bar{L}_1^{\varphi}\right) dL_{1t}.$$

To express the utility change in terms of percentage deviations from the steady state, I multiply and divide each term by its respective steady-state quantity (e.g., $dC_{1t} \approx \bar{C}_1 \hat{C}_{1t}$):

$$\Delta U_{1t} \approx \left(\frac{1}{\bar{C}_1}\right) \bar{C}_1 \hat{C}_{1t} + \left(\nu(\bar{E}_1^H)^{-\sigma_E}\right) \bar{E}_1^H \hat{E}_{1t}^H - \left(\chi \bar{L}_1^{\varphi}\right) \bar{L}_1 \hat{L}_{1t}.$$

Simplifying the exponents gives the log-linearized deviation in period utility:

$$\Delta U_{1t} \approx \hat{C}_{1t} + \nu(\bar{E}_1^H)^{1-\sigma_E} \hat{E}_{1t}^H - \chi \bar{L}_1^{1+\varphi} \hat{L}_{1t}.$$

The theoretical ambiguity of the total utility deviation stems from the labor supply response $\hat{L}_{1t}$, which is ambiguous as established in Proposition 4. However, the first two terms represent strict utility losses. First, steady-state consumption strictly decreases ($\hat{C}_{1t} < 0$) by Proposition 3. Second, the household energy demand deviation ($\hat{E}_{1t}^H < 0$) is strictly negative due to the localized variable electricity price hike ($\hat{p}_{1t}^H > 0$ from Proposition 1) and the falling real wage ($\hat{w}_{1t} < 0$ from Proposition 2). Because two of the three components strictly decrease utility while the third remains ambiguous, the net utility deviation is theoretically ambiguous but heavily biased downward.

To formalize the magnitude of this welfare incidence, I define the Consumption Equivalent Variation (CEV), denoted τ. The CEV represents the uniform percentage change in steady-state consumption that would yield an equivalent change in utility, holding energy consumption and labor supply constant at their pre-shock levels.

I evaluate the CEV by setting the utility difference equal to the total utility deviation ΔU_{1t}:

$$U((1+\tau)\bar{C}_1, \bar{E}_1^H, \bar{L}_1) - U(\bar{C}_1, \bar{E}_1^H, \bar{L}_1) = \Delta U_{1t}.$$

Given the additively separable log-utility over consumption, this becomes:

$$\ln((1+\tau)\bar{C}_1) - \ln(\bar{C}_1) = \Delta U_{1t}.$$

Using logarithmic properties, the steady-state consumption terms cancel out:

$$\ln(1+\tau) = \Delta U_{1t}.$$

For small deviations around the steady state, $\ln(1+\tau) \approx \tau$. Therefore, the Consumption Equivalent Variation is approximately equal to the log-linearized utility deviation:

$$\tau \approx \Delta U_{1t}.$$

□

F Model Parameter Calibration

This appendix details the explicit algebraic derivations used to extract deep structural parameters from the steady-state equilibrium conditions and empirical estimates.

F.1 Regulated Grid Cost Parameters

To extract the local capital share (ω) and the unrecovered capital gap per MW (Ψ), I map the model's structural parameters directly to the empirical fixed charge estimates. I define the baseline hazard $\bar{\rho}$ as the average residential rate case approval frequency among control counties across the 2014–2025 observation window, yielding a value of 0.5943. Because the observed local base rate hike is mathematically conditional on regulatory approval, I first isolate the "true" local capital shock by dividing the empirical unconditional fixed charge effect (0.00683) by the post-shock probability of a rate case. This post-shock probability is calculated as the baseline hazard $\bar{\rho}$ plus the marginal rate case probability increase (0.00654):

$$\text{True Local Shock} = \frac{0.00683}{0.5943 + 0.00654} = 0.01136.$$

I then define R as the ratio of this true local shock (0.01136) to the socialized Transmission Cost Recovery (TCR) spillover (0.00057), scaled by the respective household populations of the host region ($N_1 = 68,414$) and the broader grid ($N_{grid} = 751,550$):

$$R = \left(\frac{0.01136}{0.00057}\right) \times \frac{68,414}{751,550}.$$

This ratio effectively compares the localized cost burden to the socialized cost burden. To analytically isolate the share of capital costs uniquely localized to the host region ($\omega = 0.6428$), I substitute the structural per-capita definitions of the target base fixed charge ($\frac{\omega \Psi D_t}{N_1}$) and the Transmission Cost Adjustment ($\frac{(1-\omega)\Psi D_t}{N_{grid}}$) into the empirical ratio:

$$R = \left(\frac{\frac{\omega \Psi D_t}{N_1}}{\frac{(1-\omega)\Psi D_t}{N_{grid}}} \right) \times \frac{N_1}{N_{grid}}.$$

The scale of the unrecovered capital gap (ΨD_t) and the regional population denominators perfectly cancel out, leaving the ratio strictly as a function of the localization parameter:

$$R = \left(\frac{\omega}{1-\omega} \times \frac{N_{grid}}{N_1} \right) \times \frac{N_1}{N_{grid}}$$
$$R = \frac{\omega}{1-\omega}.$$

Algebraically rearranging this expression to solve for ω yields the final parameter extraction:

$$\begin{aligned} R(1-\omega) &= \omega \\ R - R\omega &= \omega \\ R &= \omega(1+R) \\ \omega &= \frac{R}{1+R}. \end{aligned}$$

Given ω, I back out the total capital gap parameter Ψ by scaling the true local shock by the host population and the steady-state annualized fixed tariff ($\bar{F}$). I derive this baseline tariff $\bar{F}$ by computing the geometric mean of monthly real residential fixed charges across all

control observations and annualizing the result, yielding \$200.47. Evaluating the expression for a standard 100 MW shock ($D_t = 1$):

$$\Psi = \frac{0.01136 \cdot 68,414 \cdot 200.47}{0.6428 \cdot 1} = \$242,412.09 \text{ per 100 MW}.$$

To calibrate the logistic hazard function governing PUC rate-case approvals, I set the intercept γ_0 to match the unconditional baseline approval probability $\bar{\rho}$ (0.5943) observed in the control group:

$$\gamma_0 = \ln\left(\frac{0.5943}{1 - 0.5943}\right) = 0.3817.$$

The sensitivity parameter γ_1 is then derived to exactly match the empirical post-shock probability $\rho_{shock} = 0.5943 + 0.00654$:

$$\gamma_1 = \frac{1}{0.6428 \cdot 242,412.09 \cdot 1}\left[\ln\left(\frac{0.60084}{1 - 0.60084}\right) - 0.3817\right] = 1.744 \times 10^{-7}.$$

F.2 Exact Identification of the Nodal Premium

In the deregulated market ($\Omega = 1$), the nodal premium κ is unobserved in retail utility data. I analytically identify it by forcing the model to exactly match the empirical manufacturing wage penalty ($\hat{w} = -0.0049$). By utilizing the calibrated manufacturing factor cost shares for energy ($s_E = 0.025$) and labor ($s_L = 0.300$), I isolate the energy price shock within the log-linearized Factor Price Frontier derived in Proposition 2. This backs out the precise percentage jump in the industrial electricity price needed to force this specific wage collapse under a zero-profit condition:

$$\hat{p}^M = \frac{\hat{w}}{-(s_E/s_L)} = \frac{-0.0049}{-(0.025/0.300)} = 0.0593.$$

I must convert this 5.93% price deviation into an absolute required Locational Marginal Price (LMP) spike for the host region. Because the Utility Rate Database strictly records retail tariffs and omits wholesale market transactions, I construct a proxy for the baseline steady-state wholesale energy cost ($\bar{p}^{energy}$) using retail data. I begin with the 2014 mean real retail variable rate across all control counties $\bar{p}$, which equates to \$133.40 per MWh. To isolate the strictly variable generation component of this retail price, I apply two scaling factors. First, I multiply the retail rate by the Energy Information Administration's (EIA) estimate for the aggregate generation share of retail electricity costs (61.6%). Second, I multiply the resulting total generation cost by the variable cost share of a combined-cycle gas turbine (CCGT) baseload. Using Lazard's (2025) Levelized Cost of Energy estimates, I calculate this variable share as the ratio of variable operating costs (\$25/MWh) to total LCOE (\$48/MWh), yielding 52.08%. Applying these scalars to the baseline retail rate isolates the steady-state wholesale energy cost:

$$\bar{p}^{energy} = \$133.40 \times 0.616 \times 0.5208 = \$42.80 \text{ per MWh.}$$

By evaluating the structural price deviation against this baseline wholesale cost, I find the total absolute price increase ΔLMP_{1t}:

$$\Delta LMP_{1t} = 0.0593 \times 42.80 = \$2.54 \text{ per MWh.}$$

As established in the DSGE setup, the wholesale price in the host region (LMP_{1t}) comprises a system-wide base energy price (p_t^{energy}) and a localized nodal premium (p_t^{nodal}). Therefore, the total absolute price shock can be decomposed into the marginal increase in grid-wide generation costs and the spatial congestion penalty: $\Delta LMP_{1t} = \Delta p_t^{energy} + p_t^{nodal}$. To isolate the unobserved nodal premium, I first calculate the annual energy demand of a standard 100 MW data center shock ($D_t = 100$) operating at a 50% load factor ($\delta = 0.50$), yielding $E_t^{DC} = 438,000$ MWh.

I then net out the systemic component of the price increase ($\Delta p_t^{energy} = v_1 E_t^{DC} = \$0.02$ / MWh), which reflects the grid-wide marginal cost slope (v_1) applied to the facility's inelastic demand. Because the generation supply stack is relatively flat over this specific demand increment, the systemic cost increase is negligible. Consequently, nearly the entirety of the wholesale price jump is driven by physical transmission constraints rather than generation expenses. This leaves a residual nodal premium of $p_t^{nodal} = \$2.52$ / MWh. Because the nodal premium is defined to scale linearly with the localized load shock, I analytically solve for the marginal congestion parameter κ:

$$\begin{aligned} \kappa E_t^{DC} &= \Delta LMP_{1t} - v_1 E_t^{DC} \\ \kappa &= \frac{2.54 - 0.02}{438,000} = 5.756 \times 10^{-6}. \end{aligned}$$

F.3 Extraction of Deep Preference and Technology Parameters

To calibrate the household utility weights, I first anchor steady-state labor income to the 2024 U.S. median household income, adjusted to real 2025 dollars ($\bar{w}\bar{L} = \$85,933.21$). Because standard macroeconomic convention fixes the steady-state labor supply at one-third of the available time endowment ($\bar{L} = 1/3$), this income level analytically implies a baseline real wage rate of $\bar{w} = \$257,799.62$. I then evaluate the steady-state budget constraint to find the residual consumption of the composite good $\bar{C}$. By subtracting the annualized steady-state fixed grid tariff ($\bar{F} = \$200.47$) and baseline household energy expenditures ($\bar{p}\bar{E}_H = \$1,400.79$) from labor income, I isolate $\bar{C}$:

$$\bar{C} = \bar{w}\bar{L} - \bar{p}\bar{E}_H - \bar{F} = 85,933.21 - 1,400.79 - 200.47 = \$84,331.95.$$

With $\bar{C}$ determined, and drawing upon the previously calibrated energy curvature ($\sigma_E = 5.00$) and inverse Frisch elasticity ($\varphi = 5.00$), the energy utility weight ν and labor disutility weight χ are solved directly from the household's intratemporal first-order conditions:

$$\begin{aligned}
\nu &= \frac{\bar{p}(\bar{E}_H)^{\sigma_E}}{\bar{C}} = 201.90, \\
\chi &= \frac{\bar{w}}{\bar{L}^{\varphi}\bar{C}} = \frac{257,799.62}{(1/3)^5 \times 84,331.95} = 742.84.
\end{aligned}$$

For the "macro-average" representative firm, the CES distribution parameter α is analytically derived from the optimal input ratio evaluated at the steady state. I utilize the generalized macro-average factor cost shares for energy ($s_E = 0.010$) and labor ($s_L = 0.5985$), alongside the CES elasticity of substitution ($\sigma = 0.60$). By establishing the steady-state input price ratio between the baseline real wage ($\bar{w} = \$257,799.62$) and the wholesale energy cost ($\bar{p}^{\text{energy}} = \$42.80$), I formulate the structural identity:

$$\begin{aligned}
\frac{s_E}{s_L} &= \left(\frac{\alpha}{1-\alpha}\right)^{\sigma} \left(\frac{\bar{w}}{\bar{p}^{\text{energy}}}\right)^{\sigma-1} \\
\frac{0.010}{0.5985} &= \left(\frac{\alpha}{1-\alpha}\right)^{0.60} \left(\frac{257,799.62}{42.80}\right)^{-0.40}.
\end{aligned}$$

Rearranging this expression to isolate α yields the required energy distribution weight:

$$\begin{aligned}
K &= \left(\frac{0.010}{0.5985}\right)^{\frac{1}{0.60}} \left(\frac{42.80}{257,799.62}\right)^{\frac{-0.40}{0.60}}, \\
\alpha &= \frac{K}{1+K} = 0.265491.
\end{aligned}$$

Finally, Total Factor Productivity A is extracted using the perfectly competitive zero-profit condition. This restriction dictates that the unit cost of production must equal the numeraire price of the composite good ($MC = 1$):

$$\begin{aligned}
1 &= \frac{1}{A}\left[(1-\alpha)^{\sigma}\bar{w}^{1-\sigma} + \alpha^{\sigma}(\bar{p}^{\text{energy}})^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \\
1 &= \frac{1}{A}\left[(1-0.265491)^{0.60}(257,799.62)^{0.40} + (0.265491)^{0.60}(42.80)^{0.40}\right]^{\frac{1}{0.40}}.
\end{aligned}$$

Solving for the scaling parameter A perfectly grounds the firm's production technology to the empirical baseline:

$$A = \left[(1-0.265491)^{0.60}(257,799.62)^{0.40} + (0.265491)^{0.60}(42.80)^{0.40}\right]^{\frac{1}{0.40}} = 169,148.4095.$$

G Impulse Response Panel: 100 MW Benchmark Scenario

Figure 12: 100 MW Shock under Regulated Regime ($\Omega = 0$)

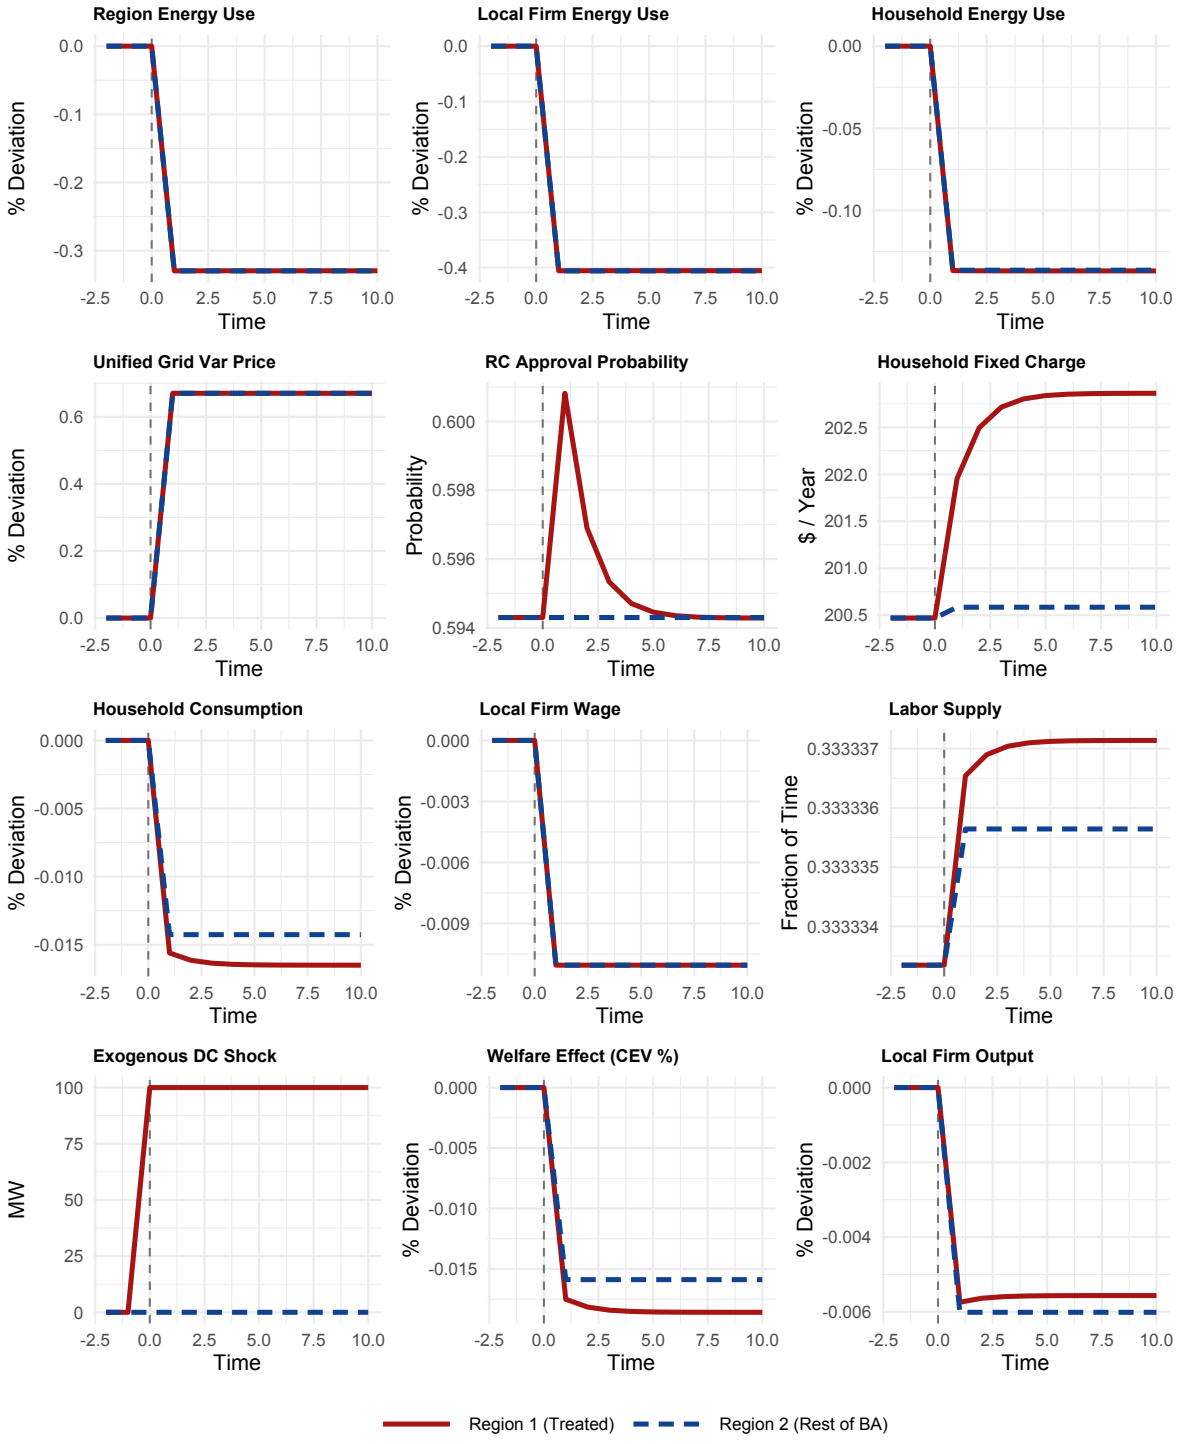

Note: For analytical tractability within the general equilibrium calculations, the impulse response function for the household fixed charge depicts the expected tariff trajectory. This is evaluated as the probability-weighted average of the baseline and target rates via the rate-case hazard function rather than a single stochastic realization of the regulatory approval process.

Figure 13: 100 MW Shock under Deregulated Regime ($\Omega = 1$)

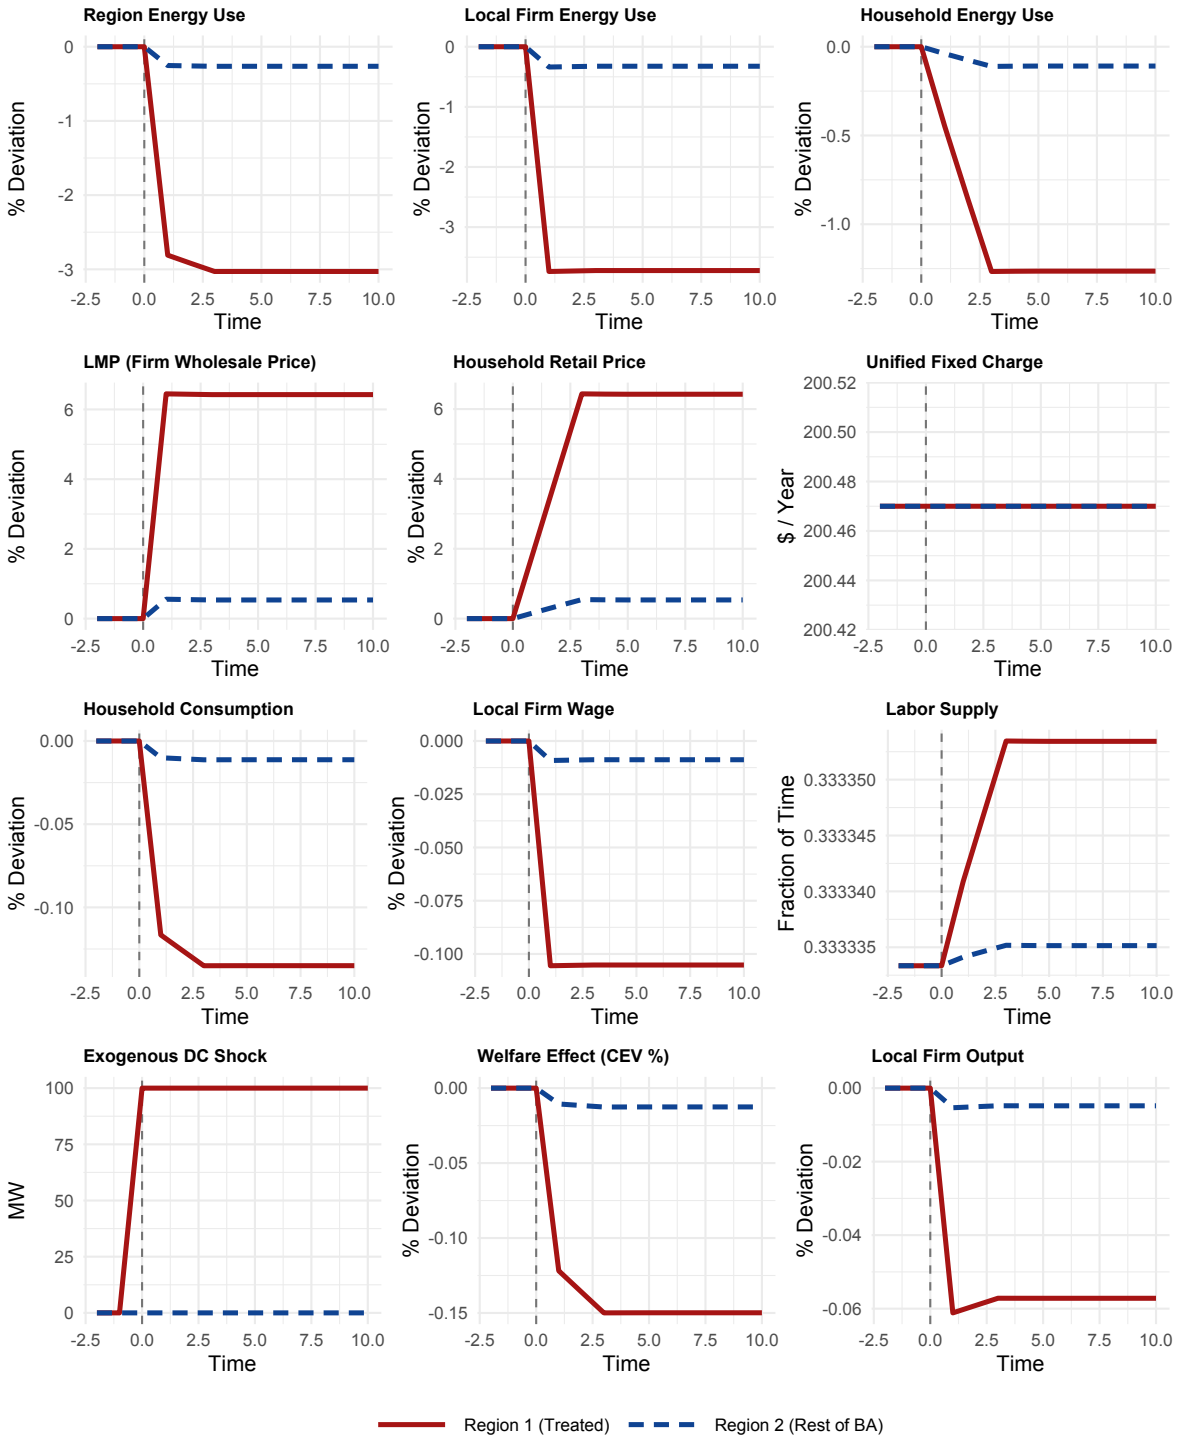

H Impulse Response Panel: Real-World Case Studies

Figure 14: Impact of Data Center Alley Build-Out (Loudoun County, VA)

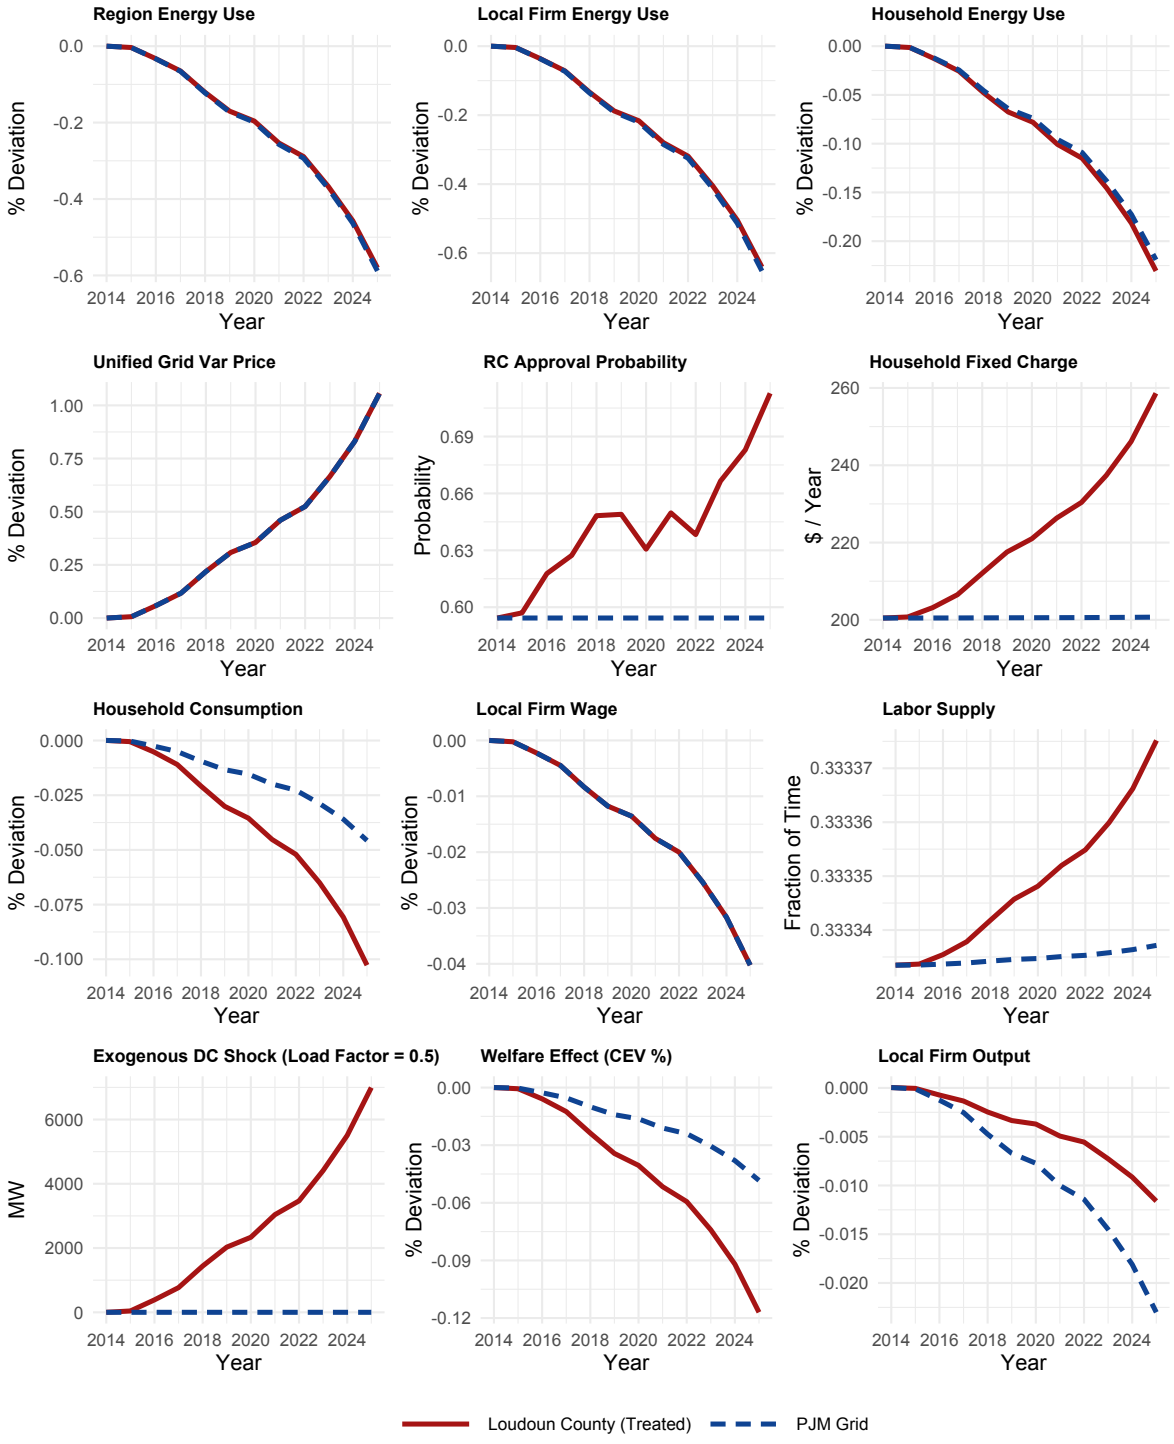

Note: For analytical tractability within the general equilibrium calculations, the impulse response function for the household fixed charge depicts the expected tariff trajectory. This is evaluated as the probability-weighted average of the baseline and target rates via the rate-case hazard function rather than a single stochastic realization of the regulatory approval process.

Figure 15: Impact of xAI Colossus Development (Shelby County, TN)

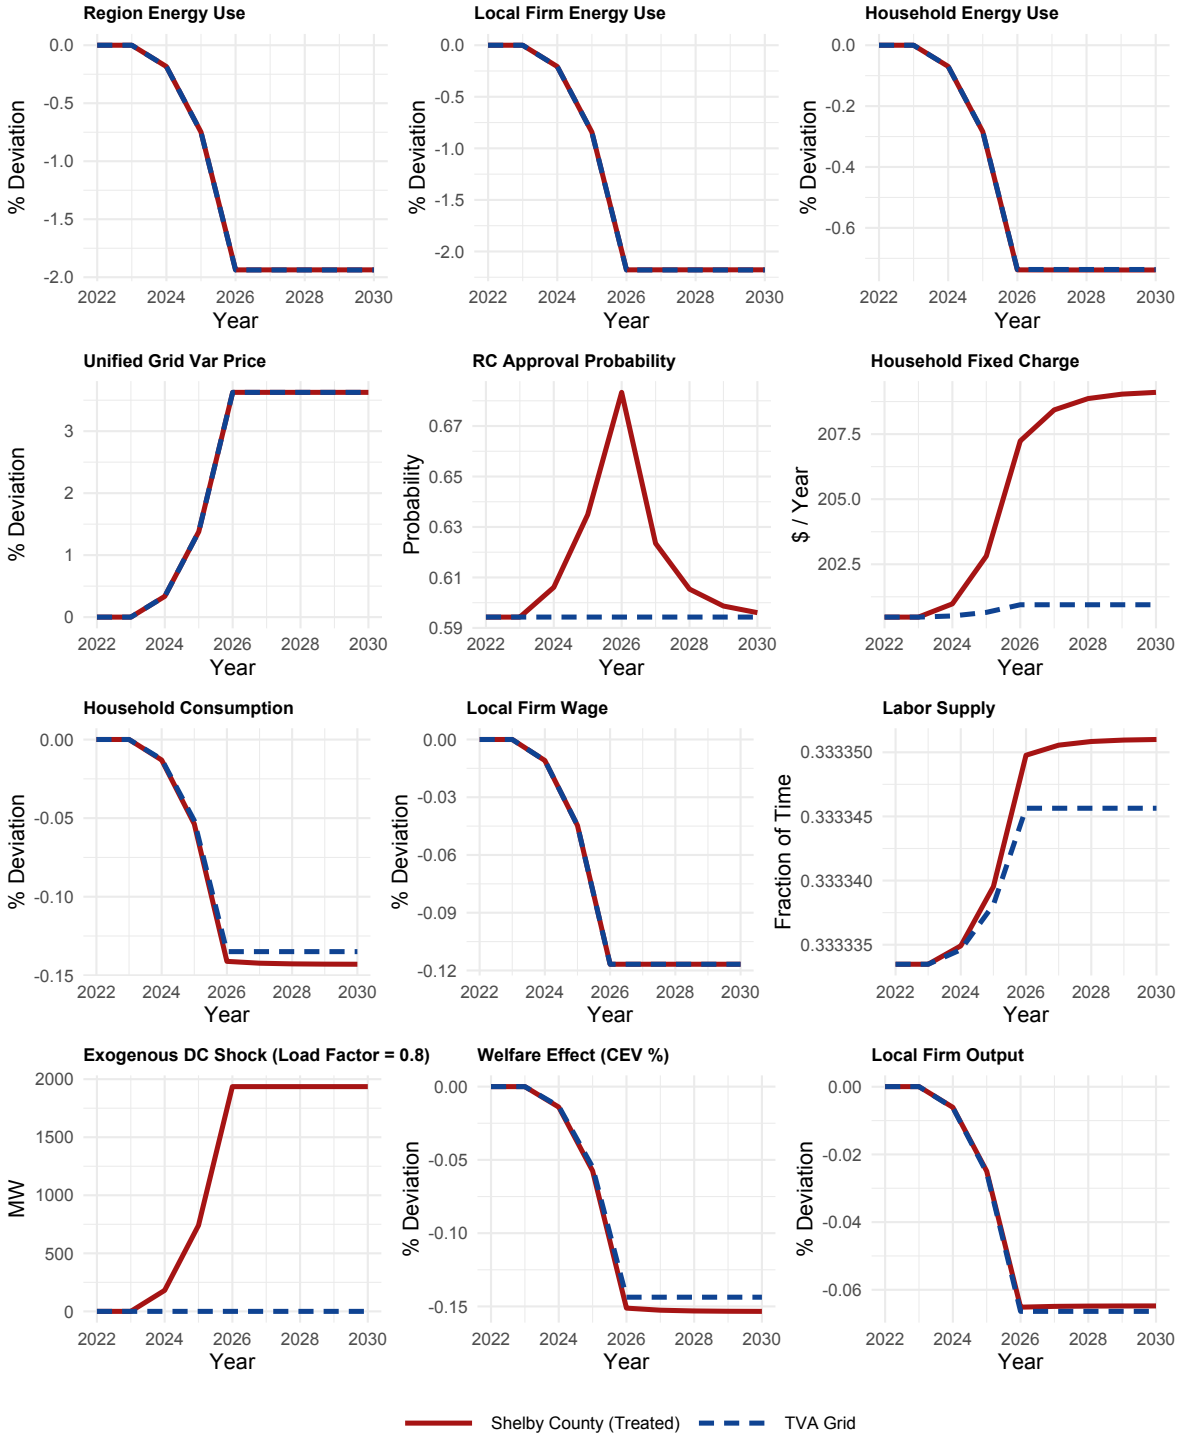

Note: For analytical tractability within the general equilibrium calculations, the impulse response function for the household fixed charge depicts the expected tariff trajectory. This is evaluated as the probability-weighted average of the baseline and target rates via the rate-case hazard function rather than a single stochastic realization of the regulatory approval process.
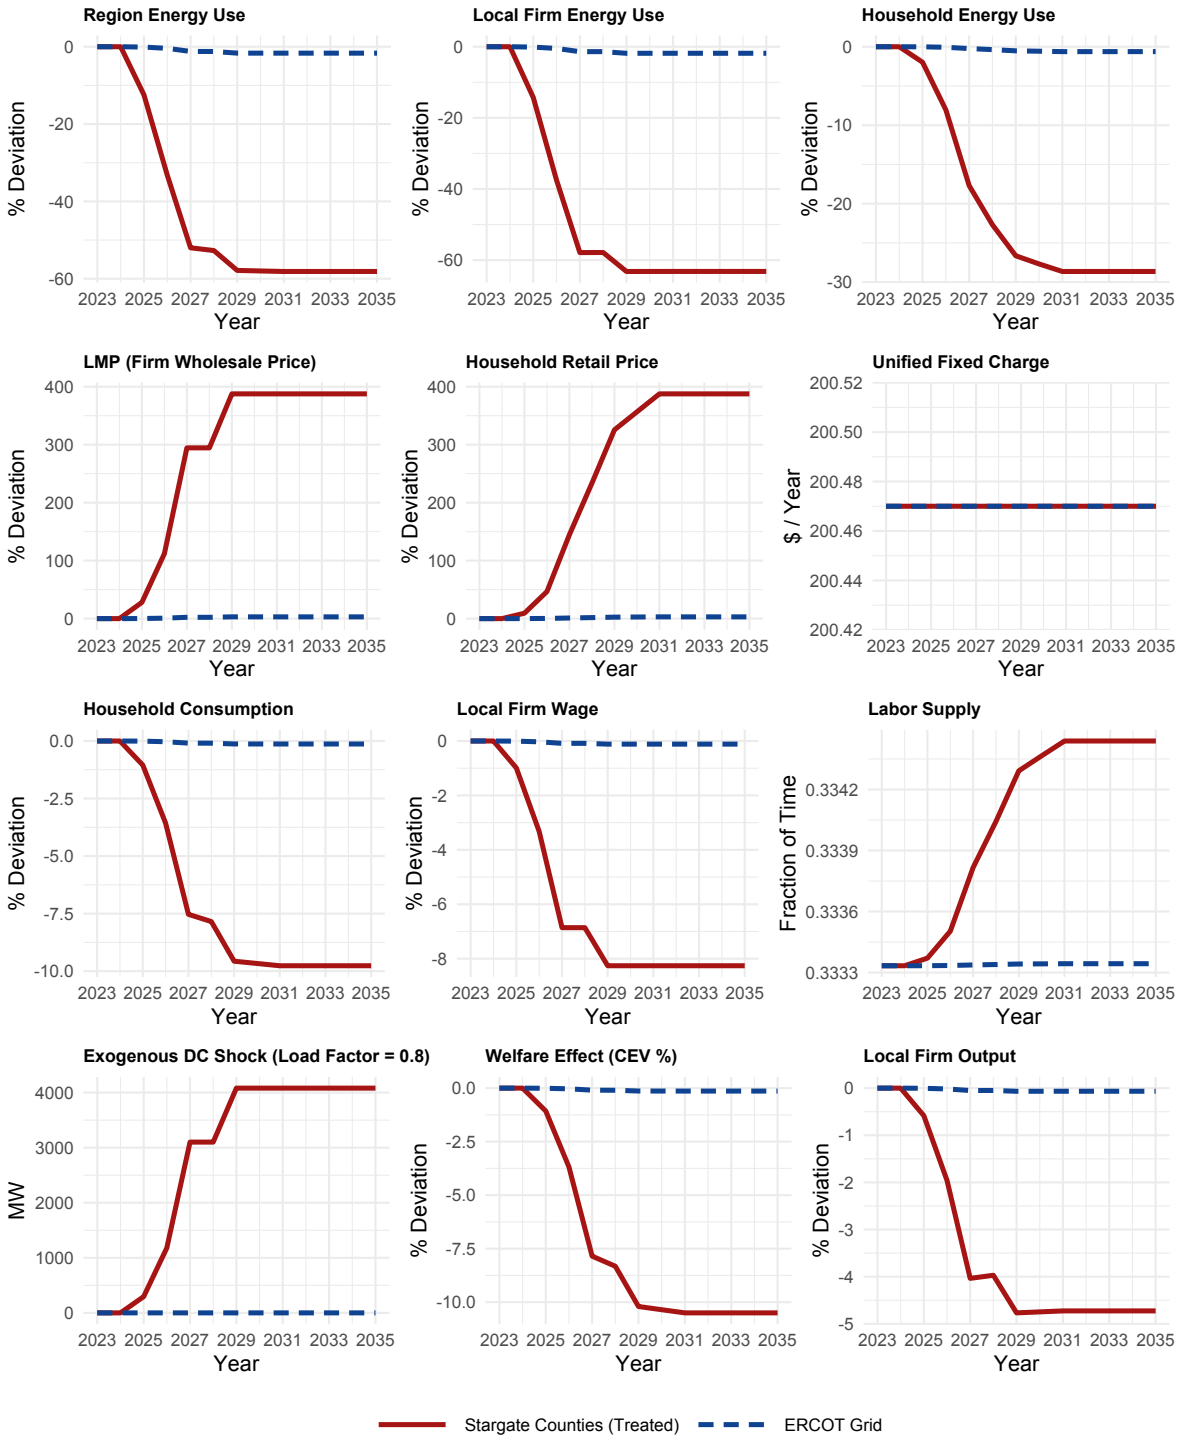
Figure 16: Impact of OpenAI Development (Taylor & Shackelford Counties, TX)