CITY OF MARSHALLTOWN
PLAN ZONING COMMISSION
NOTICE OF PUBLIC MEETING
CITY HALL COUNCIL CHAMBERS
10 WEST STATE STREET
AUGUST 6, 2026, 5:00 PM

AGENDA

CALL TO ORDER

ROLL CALL

Jon Boston, Matthew Brodin, Mahala Casady, Deirdre Gruendler, Benjamin Harris-Medina, Patrick Streit, Stephen Valbracht

BUSINESS

1. Public Hearing and Recomendation on Ordinance 15123 on the Amendment to Chapter 156 of the Code of Ordinances of the City of Marshalltown, Iowa, Establishing a Temporary Data Center Moratorium

2. Work Session on Data Center Regulations

ADJOURNMENT

MARSHALLTOWN
IOWA
HOUSING & COMMUNITY DEVELOPMENT

Deb Millizer, Director
Clayton Ender, Assistant Director
24 North Center Street
Marshalltown, IA 50158-4911
Tel - (641) 754-5756
Fax - (641) 754-5717

TO: Planning and Zoning Commission
FROM: Clayton Ender, AICP, Assistant Housing & Community Development Director
DATE: August 6th, 2026
RE: Public Hearing and Recommendation on Ordinance 15123 on the Amendment to Chapter 156 of the Code of Ordinances of the City of Marshalltown, Iowa, Establishing a Temporary Data Center Moratorium

City Staff Contact:	Clayton Ender, AICP Assistant Director of Housing and Community Development Phone: 641-754-5756 Email: cender@marshalltown-ia.gov
Applicant:	City of Marshalltown 24 N Center Street Marshalltown IA 50158
Recommendation:	Staff recommends that the Planning and Zoning Commission conduct the public hearing and forward a recommendation of approval of Ordinance 15123, establishing a temporary Data Center moratorium, to the City Council.
Background	At its July 13, 2026 meeting, the City Council approved, by a 5-2 vote, a motion establishing a temporary moratorium on Data Center development within the City of Marshalltown. Following that meeting, the City Attorney advised that because the City's Zoning Ordinance regulates the acceptance and processing of development applications, a Council motion alone is not sufficient to temporarily suspend administration of the adopted ordinance. Implementation of the moratorium requires adoption of an ordinance. Ordinance 15123 establishes a temporary moratorium on the acceptance, processing, and issuance of permits, applications, and development approvals relating to Data Centers. The purpose of the temporary moratorium is to allow the City time to evaluate potential amendments to the Comprehensive Plan, Zoning Ordinance, subdivision regulations, development standards, and other applicable policies relating to Data Centers.

	The proposed ordinance would apply to: 1. New Data Centers proposed within the City of Marshalltown; and 2. The substantial expansion, redevelopment, or conversion of existing structures or uses for Data Center purposes. The ordinance defines a Data Center as a building, structure, facility, or portion thereof primarily used to house computer servers, data-processing equipment, digital storage equipment, networking equipment, telecommunications equipment, or related systems for the remote storage, management, processing, transmission, or distribution of digital data. The moratorium would not prohibit ordinary repair, maintenance, or replacement work necessary to protect public health, safety, or welfare at lawfully existing facilities, provided such work does not constitute establishment, expansion, or conversion into a Data Center. The proposed expiration date of the moratorium is November 23, 2026, unless earlier repealed, extended, or modified by action of the City Council.
Review Criteria:	In determining whether to approve, approve with conditions, or deny a zoning text amendment, the review bodies shall consider the following review criteria: 1. **The request complies with the applicable standards of this Zoning Ordinance, the City Code of Ordinances, and any applicable county, state, or federal requirements.** The proposed ordinance is being considered in accordance with the amendment procedures established by the City Code of Ordinances and applicable Iowa law. The temporary moratorium does not conflict with existing zoning standards but temporarily modifies the administration of development approvals related to Data Centers while the City completes additional review. The proposed ordinance establishes clear applicability, definitions, exceptions, enforcement provisions, duration, and procedures consistent with the City's legislative authority to adopt zoning regulations and temporary development controls.

2. **The request substantially conforms to any associated prior approval for the development, including, but not limited to, a special use permit, Master Development Plan, or Site Plan.**

   The proposed amendment is a citywide text amendment establishing temporary regulations. It is not associated with a specific development application, special use permit, Master Development Plan, or Site Plan.

3. **The administrative body has considered the recommendation of staff.**

   Staff has reviewed the proposed temporary moratorium ordinance and recommends approval. Staff recommends the Planning and Zoning Commission forward a recommendation of approval to the City Council.

4. **The request is consistent with applicable policies of the Comprehensive Plan and applicable utility plans and capital improvements plans; or, if it addresses a topic that is not contained or not fully developed in the Comprehensive Plan, the request does not impair the implementation of the Comprehensive Plan.**

   The City's Comprehensive Plan does not contain detailed policies specifically addressing Data Centers. The proposed temporary moratorium does not establish permanent restrictions or development standards but provides time for the City to evaluate whether additional policies or regulations are necessary.

   The temporary nature of the ordinance allows the City to study potential impacts associated with Data Centers, including utility capacity, infrastructure demands, and compatibility with surrounding land uses, without impairing implementation of the Comprehensive Plan.

5. **The request promotes the purposes of this Zoning Ordinance as established in § 156.A.002, Purposes, and in other applicable purpose statements in this chapter.**

   Section 156.A.002 establishes the purposes of the Zoning Ordinance, including promoting orderly growth and development, protecting public health, safety, and welfare, encouraging compatible land uses, protecting property values, and ensuring efficient use of public infrastructure.

   The temporary moratorium supports these purposes by allowing the City to evaluate appropriate regulations before processing future Data

Center development applications. The ordinance promotes orderly development by ensuring future regulations are informed by adequate planning analysis.

6. **The proposed amendment helps to mitigate adverse impacts of the use and development of land on the natural or built environments, including, but not limited to, mobility, air quality, water quality, noise levels, stormwater management, wildlife protection, and vegetation; or will be neutral with respect to these issues;**

   The proposed temporary moratorium is intended to provide sufficient time to evaluate potential impacts associated with Data Centers before additional development approvals are processed.

   Areas identified for further review include:

   - Electrical infrastructure demand;
   - Water consumption;
   - Wastewater impacts;
   - Stormwater management;
   - Noise generation;
   - Emergency response needs;
   - Transportation impacts;
   - Environmental impacts; and
   - Compatibility with surrounding land uses.

   The temporary ordinance itself is not expected to create adverse environmental impacts and provides an opportunity to evaluate mitigation measures that may be incorporated into future regulations.

7. **The proposed amendment is necessary to address a changing condition that was not anticipated in the Comprehensive Plan or this Zoning Ordinance; and**

   Data Center development represents an emerging land use with unique infrastructure, environmental, and community considerations. The City's current Zoning Ordinance does not contain specific regulations or development standards addressing Data Centers.

   The temporary moratorium provides time for the City to evaluate this emerging use and determine whether amendments to the Comprehensive Plan, Zoning Ordinance, or other development regulations are appropriate.

8. **The proposed amendment advances the strategic objectives of the City Council, such as fiscal responsibility, efficient use of infrastructure and public services, and other articulated city objectives.**

   The proposed temporary moratorium advances City Council's direction to evaluate Data Center development and establish appropriate regulations.

   The review period will allow the City to evaluate potential impacts on infrastructure, utilities, public services, and community resources while balancing economic development opportunities with responsible planning practices.

Attachments: Draft Ordinance

ORDINANCE 15123

AN ORDINANCE OF THE CITY OF MARSHALLTOWN, IOWA, ESTABLISHING A TEMPORARY MORATORIUM ON THE ACCEPTANCE, PROCESSING, AND ISSUANCE OF PERMITS, APPLICATIONS, AND DEVELOPMENT APPROVALS RELATING TO DATA CENTERS

WHEREAS, the City Council of the City of Marshalltown, Iowa, has identified a need to review the City's Comprehensive Plan, Zoning Ordinance, subdivision regulations, building and development standards, and other applicable City policies and ordinances as they relate to the development and operation of Data Centers;

WHEREAS, Data Centers may involve land-use and infrastructure considerations that warrant further study, including, but not limited to, electrical demand, water consumption, wastewater discharge, stormwater management, noise, emergency response, public safety, transportation impacts, utility capacity, environmental impacts, building design, site security, compatibility with surrounding land uses, and the potential need for additional development standards or conditions;

WHEREAS, the City Council desires to provide City staff, the Planning and Zoning Commission, and other appropriate boards, commissions, consultants, agencies, and utility providers sufficient time to study these matters and recommend any amendments to the City's Comprehensive Plan, Zoning Ordinance, development regulations, or other policies deemed appropriate;

WHEREAS, the City Council further desires to preserve the status quo while this review is underway and prevent the approval of applications that could be inconsistent with amendments or regulations developed as a result of that review;

WHEREAS, the City Council finds the following temporary regulations to be in the best interests of the City of Marshalltown and the general public.

NOW THEREFORE BE IT ORDAINED BY THE COUNCIL OF THE CITY OF MARSHALLTOWN, IOWA:

That the City of Marshalltown, Iowa, hereby adopts the following temporary ordinance designated the "TEMPORARY DATA CENTER MORATORIUM ORDINANCE," which because of its temporary nature shall not be printed as part of the permanent Code of Ordinances. The Temporary Data Center Moratorium Ordinance is as follows:

Section 1. Definitions.

For purposes of this Ordinance, the following definitions shall apply:

"Data Center" shall mean a building, structure, facility, or portion thereof primarily used to house computer servers, data-processing equipment, digital storage equipment, networking equipment, telecommunications equipment, or related systems for the remote storage, management, processing, transmission, or distribution of digital data, including associated mechanical, electrical, cooling, security, power-generation, and utility infrastructure.

The term "Data Center" shall not include incidental server or computer equipment customarily maintained as an accessory use within an office, commercial, governmental, educational, medical, industrial, or other principal use when such equipment is used primarily to support activities conducted on the same premises.

"Development Approval" shall include, but not be limited to, any permit, application, zoning amendment, rezoning request, site plan, development plan, subdivision approval, special-use permit, variance, building permit, grading permit, utility-extension approval, certificate of occupancy, or other City authorization related to the establishment, construction, expansion, or operation of a Data Center.

"Substantial Expansion" shall mean any expansion, modification, redevelopment, or conversion that increases the operational capacity, physical footprint, utility demand, or primary function of a facility for Data Center purposes.

Section 2. Temporary Moratorium Established.

A temporary moratorium is hereby established on the acceptance, processing, and issuance of permits, applications, and other Development Approvals for the establishment, construction, Substantial Expansion, or operation of Data Centers within the corporate limits of the City of Marshalltown.

During the effective period of this Ordinance, no City department shall accept, process, approve, or issue any Development Approval for a Data Center unless specifically authorized pursuant to this Ordinance.

Section 3. Applicability.

This Ordinance shall apply to:

1. New Data Centers proposed within the City of Marshalltown; and
2. The Substantial Expansion, redevelopment, or conversion of existing structures or uses for Data Center purposes.

This Ordinance shall not prevent ordinary repair, maintenance, or replacement work necessary to protect public health, safety, or welfare at a lawfully existing facility, provided such work does not constitute the establishment, expansion, or conversion of a facility into a Data Center.

Section 4. Pending Applications.

This moratorium shall apply to applications that have not received final City approval as of the effective date of this Ordinance, except to the extent the City Attorney determines that an applicant has acquired vested rights under applicable law.

Section 5. Exceptions.

The City Council may authorize an exception to this Ordinance by resolution when it determines that such exception is necessary to:

A. Protect public health, safety, or welfare;
B. Comply with applicable state or federal law;
C. Avoid substantial injustice; or
D. Address a unique circumstance where strict application of this Ordinance would result in an unreasonable hardship or otherwise be contrary to the public interest.

Section 6. Effective Date and Duration.

The provisions of this Ordinance shall take effect upon passage and publication as provided by law and shall remain in full force and effect until November 23, 2026, unless earlier repealed, extended, or otherwise modified by action of the City Council.

The City Council may extend the effective period of this Ordinance one time, by resolution, for a period not to exceed sixty (60) additional days if the City Council determines that additional time is necessary to complete review of potential amendments to the City's Comprehensive Plan, Zoning Ordinance, development regulations, or other applicable policies relating to Data Centers.

Any extension authorized pursuant to this Section shall identify the extended expiration date of the moratorium and shall not be construed as authorizing an indefinite continuation of the moratorium without further legislative action by the City Council.

Section 7. Violations.

Any person, corporation, partnership, or other entity that submits, attempts to submit, constructs, establishes, expands, or operates a Data Center in violation of this Ordinance shall be deemed in violation of this Ordinance.

Each day that a violation continues shall constitute a separate offense.

Violations of this Ordinance shall constitute a municipal infraction and shall be subject to the penalties established by the City of Marshalltown Code of Ordinances and applicable Iowa law.

Section 8. Repealer.

All previous ordinances or parts of such ordinances in conflict with the provisions of this Ordinance are hereby repealed for the effective period of this Ordinance.

Section 9. Severability Clause.

If any section, provision, or part of this Ordinance shall be adjudged to be invalid or unconstitutional, such adjudication shall not affect the validity of the Ordinance as a whole, or any section thereof, not adjudged invalid or unconstitutional.

Passed this _____ day of ____________, 2026, and signed this _____ day of ____________, 2026.

CITY OF MARSHALLTOWN, IOWA

____
Mike Ladehoff, Mayor

ATTEST:

____
Alicia Hunter, CMC, City Clerk

I, Alicia Hunter, City Clerk of the City of Marshalltown, Iowa, do hereby certify that the foregoing ORDINANCE was passed and approved by the City Council of the City of Marshalltown, Iowa, on the _____ day of ____________, 2026, and was published in the Marshalltown Times-Republican, a newspaper of general circulation in the City of Marshalltown, Iowa, on the _____ day of ____________, 2026.

________________________________ Alicia Hunter, City Clerk

Data Center Research References and Source Attribution

Working Draft: This reference list is a working document and will be updated throughout the data center ordinance review process as additional publications and materials are received by City staff. The list is intended to document resources considered during the legislative review and is not yet comprehensive.

Disclaimer: Inclusion of a publication does not imply endorsement of its contents or conclusions by the City.

Source Attribution: The "Provided By" column identifies the individual or organization that submitted or provided each reference material. Affiliations are included for informational purposes to provide context regarding the source of the material. Inclusion of a reference does not indicate endorsement by the City of Marshalltown or reflect a determination regarding the accuracy or applicability of the material.

Source Record No.	Citation	Provided By
1	Knittel, Christopher R., Juan Ramon L. Senga, and Shen Wang. *Flexible Data Centers and the Grid: Lower Costs, Higher Emissions?* NBER Working Paper No. 34065. National Bureau of Economic Research, July 2025.	David E. Frisvold, Professor, Department of Economics, University of Iowa
2	Muller, Nicholas Z. *Measuring the Impact of Data Centers in the United States Economy: Monetary Damage from Air Pollution and Greenhouse Gas Emissions.* NBER Working Paper No. 35100. National Bureau of Economic Research, April 2026.	David E. Frisvold, Professor, Department of Economics, University of Iowa
3	Alvarez, Fernando E., David Argente, Joyce Chow, and Diana Van Patten. *Data Centers and Local Economies in the Age of AI: A Shift-Share Approach.* NBER Working Paper No. 35194. National Bureau of Economic Research, May 2026.	David E. Frisvold, Professor, Department of Economics, University of Iowa

4	Nichols, Charlie, AICP. *The Physical Footprint of Artificial Intelligence. Zoning Practice*, Vol. 42, No. 10. American Planning Association, October 2025.	Travis Kraus, Professor, School of Planning & Public Affairs; Director, Iowa Initiative for Sustainable Communities, University of Iowa
5	Higdon, Jake, and Joseph Parilla. *How Advanced Industrial Zones Can Help States Thrive in the Electro-Industrial Economy.* Brookings Institution, July 9, 2026.	Heidi Pierson Dalal, Executive Director Marth-Ellen Tye Foundation & Local Resident

MARSHALLTOWN
IOWA
HOUSING & COMMUNITY DEVELOPMENT

Mike Ladehoff, Mayor
Carol Webb, City Administrator
Deb Millizer, Director
Clayton Ender, Assistant Director
24 North Center Street
Marshalltown, IA 50158-4911
Tel - (641) 754-5756
Fax - (641) 754-5717

TO: Planning and Zoning Commission
FROM: Clayton Ender, Assistant Housing & Community Development Director
DATE: August 6th, 2026
RE: Data Center Zoning Study – Work Session No. 1

Background

Data centers have emerged as a significant land use issue for communities throughout Iowa and the United States due to increasing demand for cloud computing, artificial intelligence (AI), and digital infrastructure. While these facilities can provide economic development opportunities, they also present unique planning considerations related to electrical infrastructure, water consumption, emergency services, noise, land use compatibility, and long-term community impacts.

The City of Marshalltown's Zoning Ordinance does not currently define or regulate data centers as a distinct land use. As interest in data center development has increased nationally and within Iowa, the City Council directed staff and the Planning and Zoning Commission to evaluate whether zoning regulations should be developed to address this emerging land use.

This work session marks the beginning of that regulatory review process. Rather than focusing immediately on specific zoning regulations, staff intends to follow a structured planning process that begins with identifying community goals, gathering and evaluating research, analyzing how that information applies to Marshalltown, developing potential policy options, and ultimately preparing draft regulations for Commission and City Council consideration.

Planning Process

Staff recommends that the city use the following framework throughout the ordinance review process:

- Goals – Identify the community objectives and desired outcomes that should guide future data center regulations.
- Research – Gather information from academic literature, industry publications, utility providers, peer communities, and other stakeholders.
- Analysis – Evaluate how the research applies to Marshalltown's existing infrastructure, land use patterns, economic development objectives, and Comprehensive Plan.

- Policy – Develop and evaluate potential zoning regulations and performance standards for Commission consideration.
- Implementation – Prepare draft ordinance language, conduct public hearings, and consider adoption by the City Council.

The review process is intended to ensure that any proposed regulations are based on identified community objectives and informed by available research rather than predetermined outcomes.

Work Session Objectives

This introductory work session is intended to begin the first two phases of the planning process by:

- Discussing the opportunities and challenges associated with data center development.
- Identifying community goals and planning issues that should guide future discussions.
- Providing an overview of the research staff has begun collecting.
- Receiving initial feedback from the Planning and Zoning Commission regarding additional topics that should be explored.

No formal action is requested at this meeting.

Topics for Discussion

As the review process progresses, staff anticipates discussing topics including, but not limited to:

- Overview of modern data center facilities.
- Types and scales of data centers.
- Economic development opportunities.
- Electrical infrastructure requirements.
- Water supply and cooling technologies.
- Noise generated by mechanical equipment and backup generators.
- Emergency power systems and fuel storage.
- Building size, height, and site design considerations.
- Construction activity and long-term traffic impacts.
- Compatibility with surrounding land uses.
- Environmental considerations.
- Examples of zoning regulations adopted by peer communities.

Additional topics may be identified by the Planning and Zoning Commission or through ongoing research.

Public Engagement and Research

Staff has begun gathering information from a variety of sources and will continue to conduct research throughout the ordinance review process. Publications and reference materials identified through staff research or provided by outside organizations and individuals will be compiled in a working reference list. The list will be updated throughout the review process and is intended to document the resources considered during ordinance development.

In addition to reviewing published research, staff has contacted 19 faculty members at Iowa's three public universities (the University of Iowa, Iowa State University, and the University of Northern Iowa) to request research and technical resources related to data center planning and regulation. Staff has received preliminary information from several faculty members, and those materials are included as attachments to this agenda item. Additional materials received during the review process will be incorporated into the working reference list and shared with the Planning and Zoning Commission.

Next Steps

Following this work session staff will continue gathering research, engaging stakeholders, and presenting information to the Planning and Zoning Commission over the coming months. As community goals become better defined and research is completed, staff will begin evaluating potential regulatory approaches and preparing draft zoning regulations for Commission review.

NBER WORKING PAPER SERIES

FLEXIBLE DATA CENTERS AND THE GRID: LOWER COSTS, HIGHER EMISSIONS?

Christopher R. Knittel
Juan Ramon L. Senga
Shen Wang

Working Paper 34065
http://www.nber.org/papers/w34065

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2025

This research was supported by funding from the Future Energy Systems Center of the MIT Energy Initiative (MITEI). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Flexible Data Centers and the Grid: Lower Costs, Higher Emissions?
Christopher R. Knittel, Juan Ramon L. Senga, and Shen Wang
NBER Working Paper No. 34065
July 2025
JEL No. D61, L94, Q41, Q48

ABSTRACT

Data centers are among the fastest-growing electricity consumers, raising concerns about their impact on grid operations and decarbonization goals. Their temporal flexibility—the ability to shift workloads over time—offers a source of demand-side flexibility. We model power systems in three U.S. regions: Mid-Atlantic, Texas, and WECC, under varying flexibility levels. We evaluate flexibility's effects on grid operations, investment, system costs, and emissions. Across all scenarios, flexible data centers reduce costs by shifting load from peak to off-peak hours, flattening net demand, and supporting renewable and baseload resources. This load shifting facilitates renewable integration while improving the utilization of existing baseload capacity. As a result, the emissions impact depends on which effect dominates. Higher renewable penetration increases the emissions-reduction potential of data center flexibility, while lower shares favor baseload generation and may raise emissions. Our findings highlight the importance of aligning data center flexibility with renewable deployment and regional conditions.

Christopher R. Knittel
Massachusetts Institute of Technology
MIT Sloan School of Management
Department of Economics
and NBER
knittel@mit.edu

Juan Ramon L. Senga
Massachusetts Institute of Technology
Center for Energy and Environmental
Policy Research,
and Climate Policy Center
jsenga@mit.edu

Shen Wang
Massachusetts Institute of Technology
swang16@mit.edu

1 Introduction

Data centers are among the fastest-growing electricity consumers, with their energy demand projected to increase over the coming years [1, 2]. In the U.S., that projection is an increase of 7-12% by 2030 [3]. This surge is driven by advances in artificial intelligence and the prevalence of cloud computing, which poses challenges for grid reliability [4] and decarbonization efforts [5]. The additional load could put stress on the grid and increase the usage of existing thermal power plants, which may increase carbon emissions. For example, in PJM, the forecasted increase of 32 GW (20% increase) in summer peak load mostly comes from data centers and is equivalent to adding another mid-sized state's demand to the system [6]. However, opportunities exist to operate data centers more flexibly as demand response resources, potentially mitigating large load impacts. One of these strategies takes advantage of a latent demand response resource we call *data center temporal flexibility*—the ability of data centers to change its load profile by shifting workload across time [7]. Data centers do not operate at full capacity all the time and typically maintain utilization rates of around 80% [8]. This is especially true for AI training, which has a relatively flat workload pattern. This gives a headroom of 20% of data center capacity to accommodate additional shifted data center load. Operationally, there could be benefits to shift tasks to hours when renewable availability is high or prices are low [9, 10]. This may not only save operating costs for data centers, but also provide flexibility and increase reliability for the power system while also meeting climate goals [11, 12].

Prior work hints at these benefits—curtailment relief, renewable firming, and even 24/7 carbon-free alignment when spatial shifting across data center networks is layered on [13, 14, 15, 16, 5, 17, 18]. Google's carbon-aware scheduler offers a high-profile proof-of-concept [10], and market-based signals appear decisive in whether load-shifting cuts or *raises* emissions [19]. Strikingly, even modest flexibility could offset most of the projected U.S. data center growth without a single new power plant [20].

But, several critical research questions remain unanswered: First, it is unclear how data center flexibility affects power system planning and operations. The ability to shift demand could significantly impact investment decisions, plant retirements, and operational strategies. This may alter the trajectory of capacity expansion and reliability planning for regional operators. Second, the potential grid benefits that flexible data centers bring are not yet understood for different levels of flexibility. While some portion of the data center load is flexible, the degree to which it can be shifted is constrained over time. Tasks cannot be postponed indefinitely, and certain tasks may not be shifted at all. Thus, understanding the combinations of flexibility levels (in terms of duration and shifting potential) that can lower cost and emissions is critical. Furthermore, the impact of data center flexibility on different regions may vary depending on the characteristics of the regional grid.

To address these questions, we use the GenX capacity expansion model (CEM). GenX is a least cost model that co-optimizes generation investment, retirement, and operational decisions in the power system for a representative year of operation [21]. It has been used extensively to assess policy and technology impacts on the grid (see [22, 23, 24, 25] for examples). We modified the GenX code-base to accommodate different data center temporal flexibility scenarios (see Methods in Section 8). We model combinations of scenarios that vary the *shifting horizon*—

the time window in which loads can be shifted—from 1 to 24 hours, and the *share of flexible workload*—the fraction of total shiftable demand (20% of total gross demand)—from 1% to 100%. The final GenX model then includes decisions on the hourly shifting of flexible data center load, limited by the level of flexibility in each scenario. We also include a baseline case without flexibility for comparison and assume that without flexibility, data centers have constant load throughout the year.

Our testbeds—Texas, the Mid-Atlantic, and the Western Interconnect (WECC)—collectively host 82% of the nation's projected 2030 data center demand [5, 17]. Figure 1 visualizes zonal loads and transmission corridors.

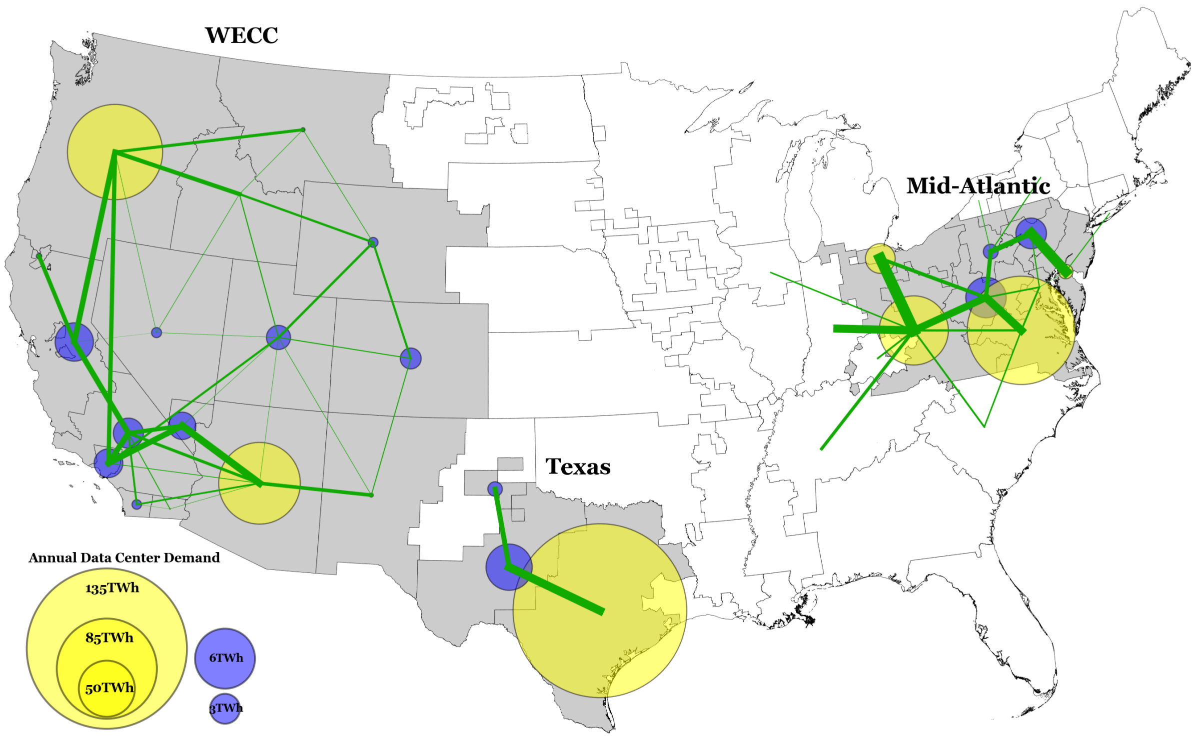

Figure 1: Model Zones with Transmission Lines and Data Center Load. Yellow and Blue circles represent annual data center demand per zone. Green line segments indicate transmission links between zones in a region where the width is scaled to the capacity of the line. Detailed information on load assumptions can be found in Supplementary Note 9.1.

Our findings show that data center temporal flexibility can significantly change a power system's operations and generation mix. Higher flexibility levels enable net load shifting from peak to off-peak hours, flattening the net load profile[1]. This reduces reliance on peaker or ramping plants and promotes more stable operation of base load generators. When renewables are sufficiently cost-competitive—as in Texas, where wind and solar are projected to supply 54% of generation—high levels of data center flexibility results in up to 40% lower CO_2 emissions and

[1]Net Load is defined as demand net of VRE (Variable Renewable Energy) generation and curtailment

accelerate retirements of coal and nuclear plants. This reverses in the Mid-Atlantic and WECC: renewable penetration is lower, coal units that survive retirements can run more uniformly, and system-wide emissions rise by as much as 3%, even though costs still fall.

We confirm this cost sensitivity in a counterfactual experiment that raises renewable investment and fixed O&M costs in Texas to 1.3 times baseline values. Renewable share collapses to 21%, coal plants remain on the system, and the emissions advantage of flexibility disappears—demonstrating that data center load shifting substitutes for baseload when clean energy is economical.

Across all regions and price scenarios, however, temporal flexibility *always* lowers total system costs—by up to 5% in Texas—while steering new investment toward renewables (wind in Texas, solar in WECC and the Mid-Atlantic) and crowding out battery storage. Flexible data center operations thus emerge as a robust, low-cost reliability resource whose climate value hinges on the underlying economics of clean power.

To our knowledge, this is the first study to trace *end-to-end* consequences—from grid build-out to hourly dispatch—of data center flexibility. The results show that it can either accelerate decarbonization or it can entrench fossil fuels that it seeks to displace.

2 The Impact of Data Center Flexibility on Data Center and System Operations

2.1 Data Center and Grid Operations

We first look at how data centers shift their load given different combinations of temporal flexibility. Fig. 2 and 3 show the entire year's data center load shifting operations for the Mid-Atlantic and Texas, respectively while fig. S5 show WECC's. We also show in Fig. 4 and 5 the interaction of data center load shifting and power system dispatch across the three regions for average summer and winter conditions, respectively.

We generally observe that data center load is temporally shifted out of daily peak load hours so that it can flatten net load. For example, in all regions, we find consistent patterns of shifting from early morning hours and early night hours to midday during the winter (Fig. 2, 3, fig. S5). This aligns with reducing the two peaks during these periods and shifting demand to midday hours with high solar availability (Fig. 4). Notably, the presence of storage and hydro (purple and dark blue, respectively, in Fig. 4I) in WECC further complements flexible load by reducing net load volatility. In all cases, flexible data center load reduces system ramping requirements. We see this from the flattening of the net load curve when comparing the dashed and solid red lines in Fig. 4.

Comparing the "No Flexibility" subplots in Fig. 5 (Fig. 5A, 5D, 5G) with those incorporating 1-hour and 24-hour shifting horizons (middle Fig. 5B, 5E, 5H, and right Fig. 5C, 5F, 5I columns), we observe a clear flattening of the net load curve (red lines) as flexibility increases. This operational shift leads to notable changes in resource utilization: peaking gas units (NGCT) are dispatched less frequently, while baseload and mid-merit units—such as NGCC, coal, and nuclear—operate more uniformly. The alignment of data center load with solar generation also increases the use of solar. Battery dispatch also becomes less prominent as

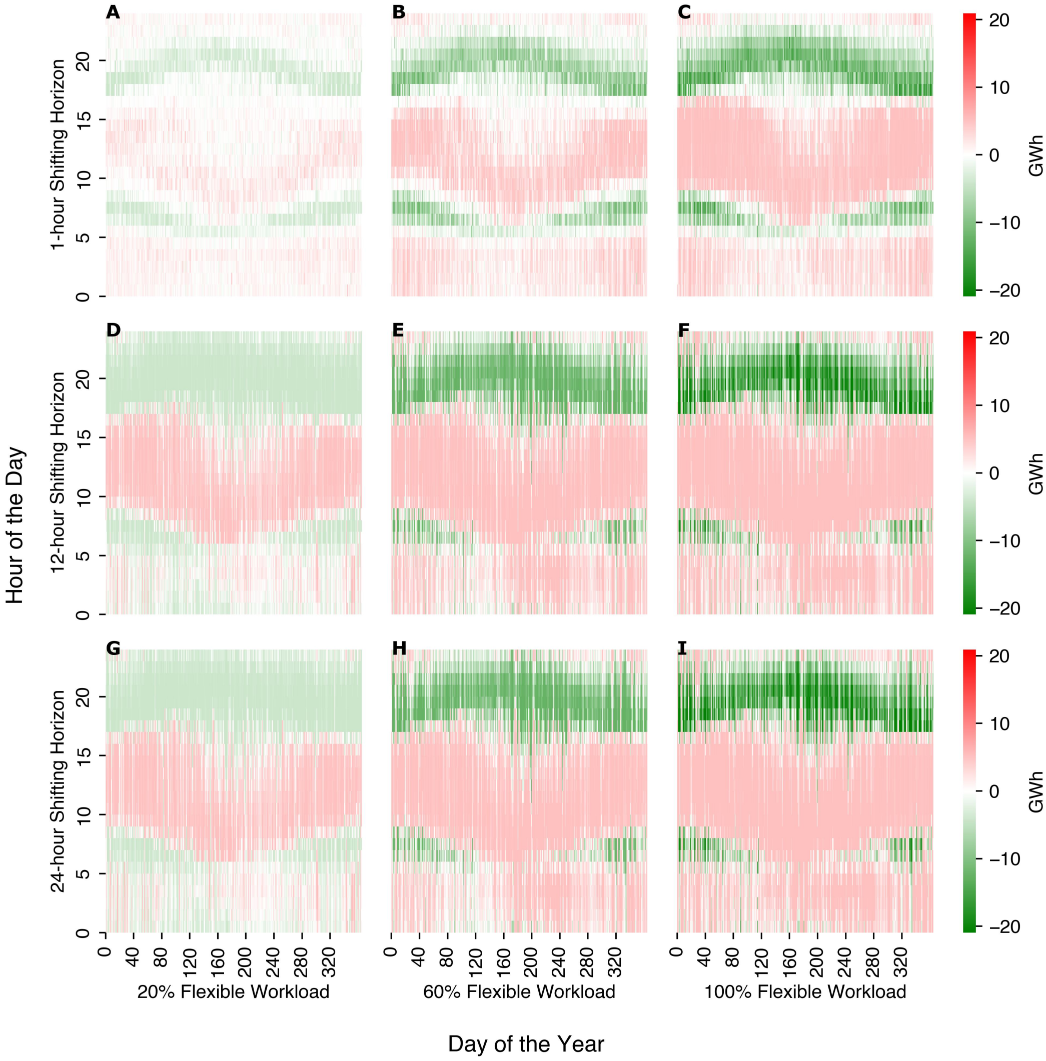

Figure 2: Data Center Load Shifting Operations for the Mid-Atlantic. Each panel displays the net hourly workload shiftted (in GWh) across an entire year, with the x-axis representing days (1–365) and the y-axis representing hours of the day (0–23). Positive values (red) indicate workload shifted into a given hour; negative values (green) represent workload shifted out. Columns show increasing flexible workload shares—4%, 12%, and 20%—based on 20%, 60%, and 100% of a shiftable portion capped at 20% of total capacity. Rows indicate shifting horizons of 1 hour, 12 hours, and 24 hours, reflecting the maximum time a task can be advanced or delayed. This illustrates how varying flexibility levels and temporal windows influence both intra-day load scheduling and broader seasonal shifting patterns.

flexible load partially substitutes its role in balancing variability. Across regions, Texas notably differs from the Mid-Atlantic and WECC. While workloads are shifted from nighttime to mid-day during the summer for the Mid-Atlantic and WECC (Fig. 2, 5C, 5I, fig. S5), data center load is frequently shifted away from the midday in Texas (Fig. 3, 5E). This difference is driven by Texas' midday net load peaks (dashed red line, Fig. 5D), which is unique in the three regions as Texas has high cooling demand during these hours and limited baseload generation (Fig. 5E, 5F). In contrast, more baseload nuclear and coal capacity in the Mid-Atlantic and WECC leads to net load peaks that are later in the day. Flexible data center loads then tend to be shifted into midday hours, leveraging lower marginal costs from solar generation and avoiding evening ramp pressures (Fig. 5C, 5I). We also observe a reduction in coal and nuclear generation in

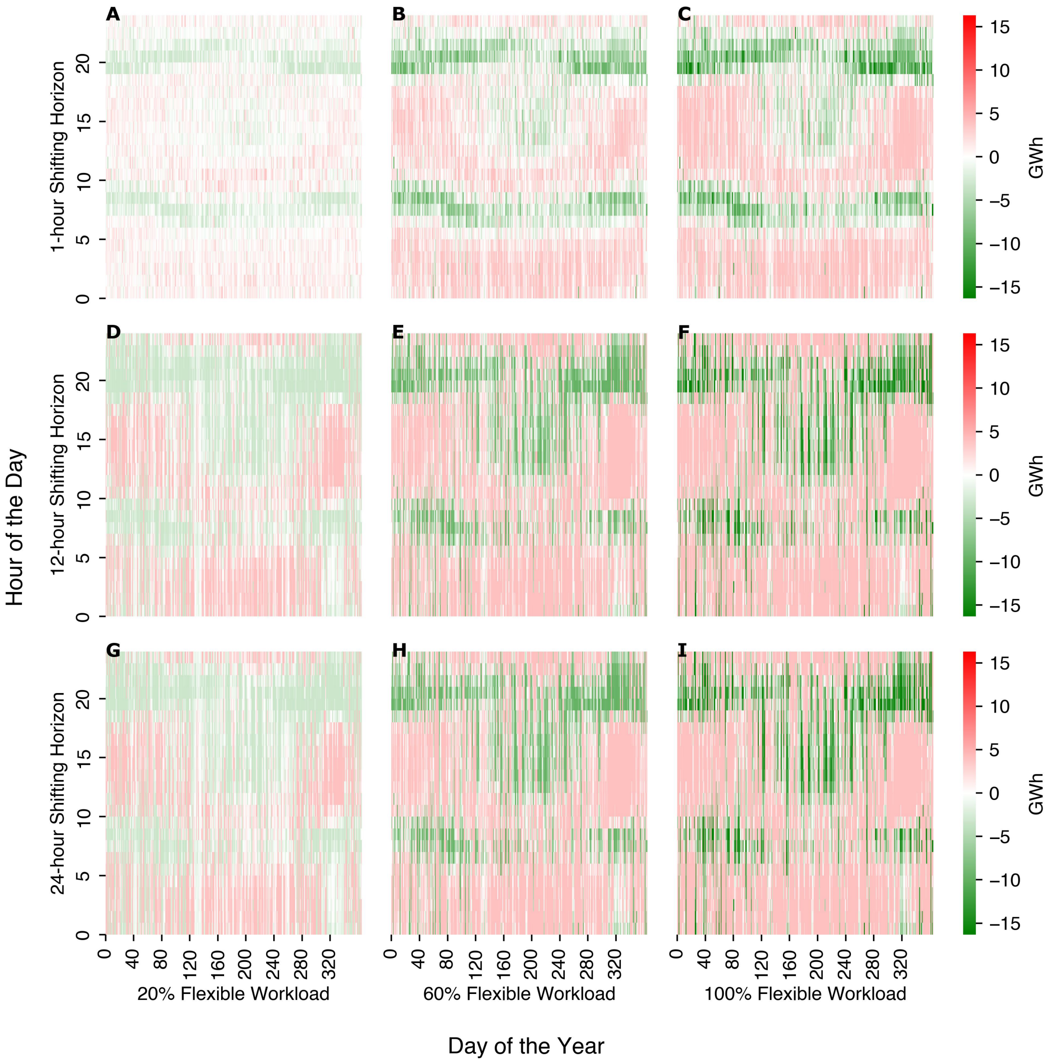

Figure 3: Data Center Load Shifting Operations for Texas. Each panel displays the net hourly workload shiftted (in GWh) across an entire year, with the x-axis representing days (1–365) and the y-axis representing hours of the day (0–23). Positive values (red) indicate workload shifted into a given hour; negative values (green) represent workload shifted out. Columns show increasing flexible workload shares—4%, 12%, and 20%—based on 20%, 60%, and 100% of a shiftable portion capped at 20% of total capacity. Rows indicate shifting horizons of 1 hour, 12 hours, and 24 hours, reflecting the maximum time a task can be advanced or delayed. This illustrates how varying flexibility levels and temporal windows influence both intra-day load scheduling and broader seasonal shifting patterns.

Texas with more flexibility, leading up to almost no baseload generation at a 24-hour shifting horizon (Fig. 5D, 5E, 5F). Regional system characteristics, such as resource mix and load shape, thus influence optimal data center shifting operations.

Finally, we note that the load shifting is more localized in the 1-hour shifting horizon case as load is constrained to a smaller time window, but becomes more pronounced with a 12- or 24-hour shifting horizon (for example, Fig. 2A, 2D, 2G for the Mid-Atlantic). This also implies that the 24-hour and 100% share of flexible workload scenario (Fig. 2I) shows the most extensive redistribution of data center load. However, even modest flexibility (12-hour shifting horizon and 60% share of flexible workload; Fig. 2E) in data center operations can already lead to significant grid re-balancing.

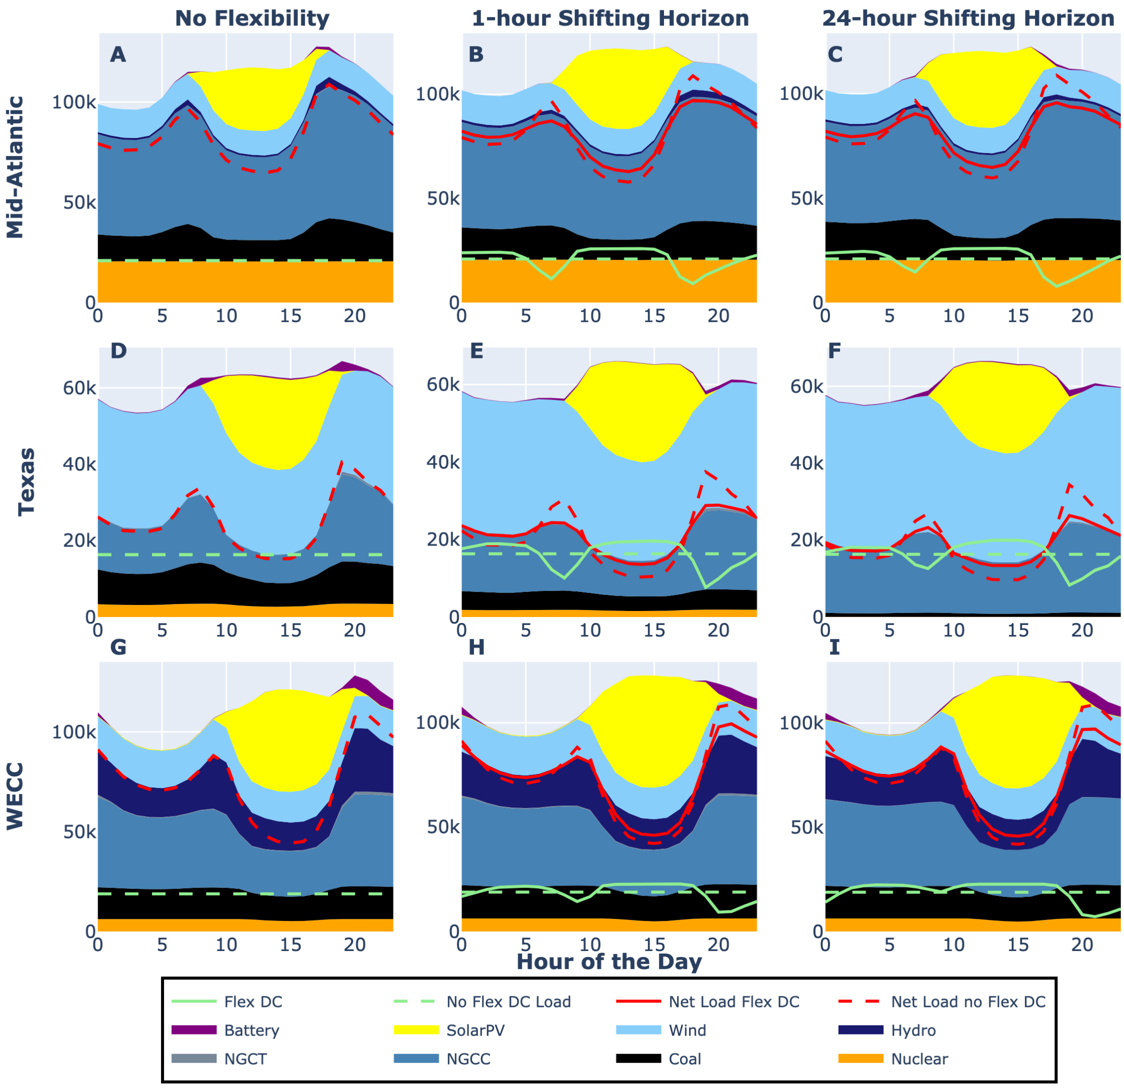

Figure 4: Average Winter Operations per Region. Each subplot shows the average generation per technology, net load, and data center load shifts in MWh per hour of the day during the Winter season. Panels correspond to three levels of operational flexibility: No Flexibility **(A, D, G)** 1-hour **(B, E, H)** and 24-hour **(C, F, I)** Shifting Horizon with a 100% share of flexible workload.

2.2 Capacity and Generation Mix

Data center flexibility and the redistribution of net load naturally affect the capacity and generation mix of a power system. Fig. 6 shows the capacity and generation per technology at varying levels of data center flexibility for each region.

Overall, we observe two main effects. First, flexibility supports renewable investments. By shifting demand into hours with high renewable availability, data center flexibility increases the economic value of wind and solar generation. This leads to larger solar investment in the Mid-Atlantic (and to a lesser extent in WECC), and larger wind investment in Texas. The preference for either is driven by which resources are more abundant in the region. Texas has strong wind potential, while the Mid-Atlantic and WECC are better suited to solar. Second, data center flexibility supports baseload operations. By flattening net load profiles, data center

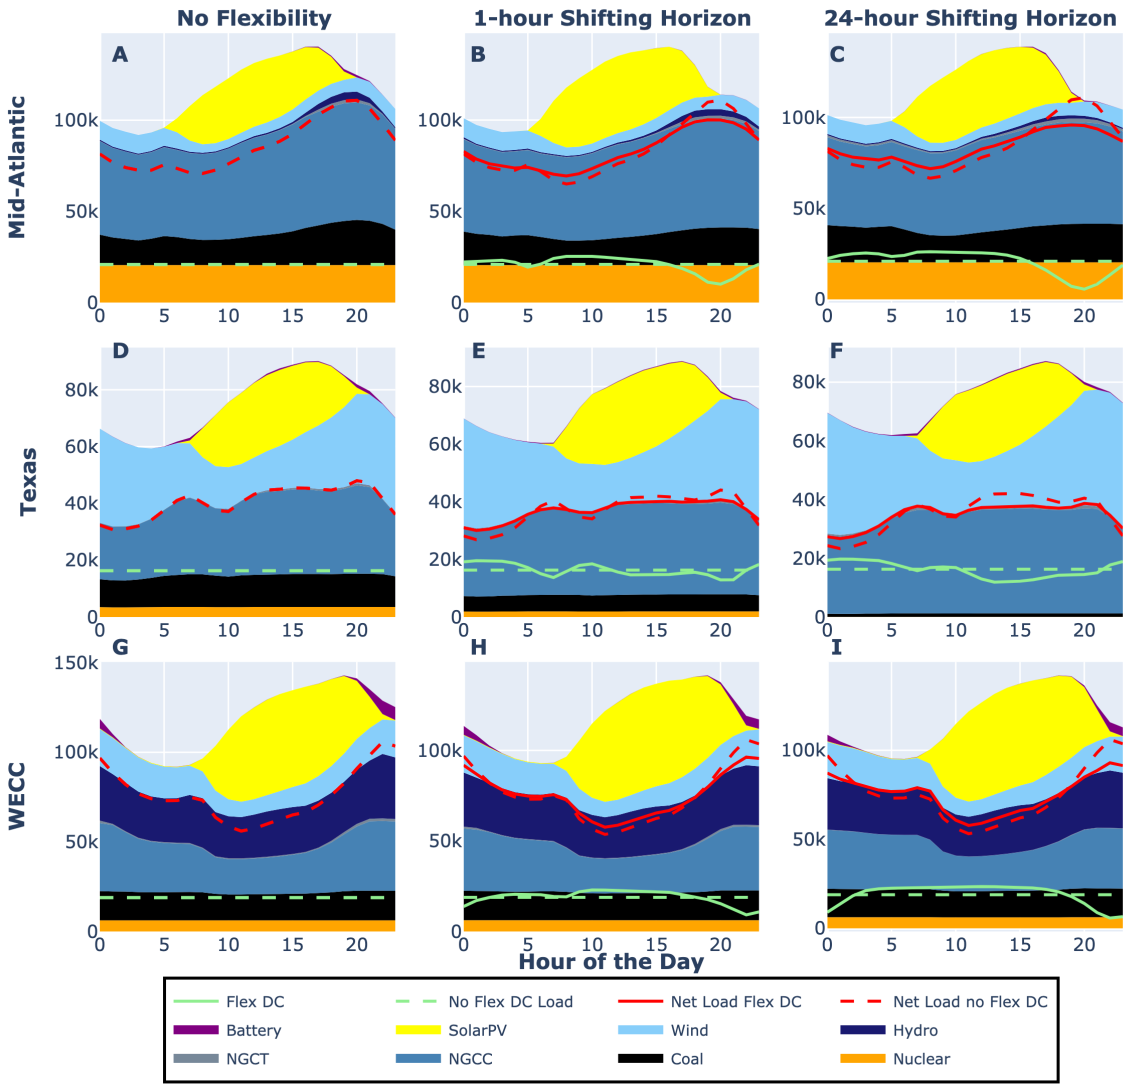

Figure 5: Average Summer Operations per Region. Each subplot shows the average generation per technology, net load, and data center load shifts in MWh per hour of the day during the Summer season. Panels correspond to three levels of operational flexibility: No Flexibility **(A, D, G)** 1-hour **(B, E, H)** and 24-hour **(C, F, I)** Shifting Horizon with a 100% share of flexible workload. The graphs don't include the regions' net electricity imports. This is relevant for the Mid-Atlantic where we assume deterministic hourly net imports. Details on net import data can be found in Supplementary Note 9.4.

flexibility makes it more cost-effective to run inflexible baseload plants like coal with fewer ramping requirements.

Whether natural gas capacity and generation increase or decrease depends on which of these two effects dominates. In the Mid-Atlantic and WECC, the support for baseload is stronger. This reduces the need for flexible natural gas capacity as coal generation becomes more economically viable. In contrast, in Texas, the support for renewables dominates due to the high share of renewable generation of around 54% (39% wind, 15% solar) of total mix, compared to 22% (10% wind and 12% solar) in the Mid-Atlantic and 33% (14% wind and 19% solar) in WECC. This drives an increase in natural gas generation that serves as a flexible, fast-ramping complement to wind. As a result, Texas sees less reliance on baseload plants like

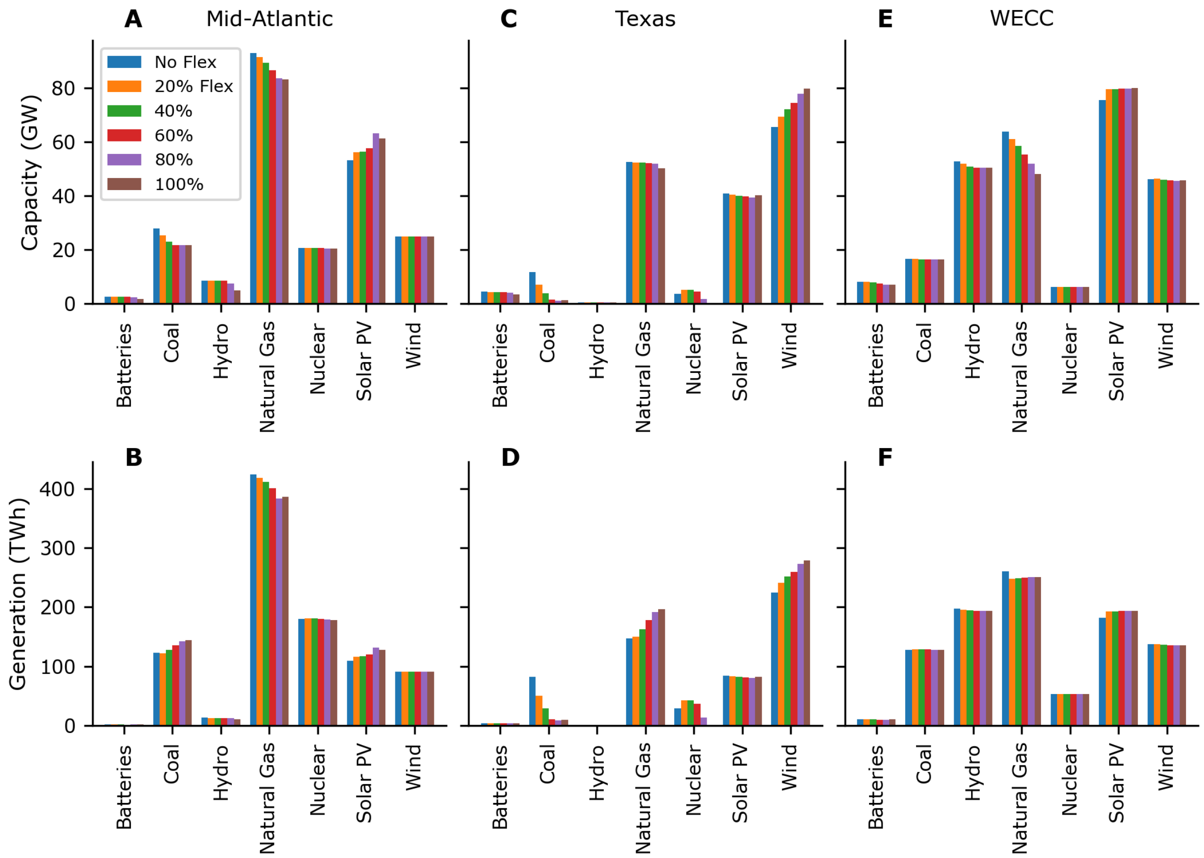

Figure 6: Capacity and Generation per Region with 24-hour shifting horizon. Top row **(A, C, E)** shows total installed capacity by technology, accounting for new investments and retirements. Bottom row **(B, D, F)** shows corresponding total generation by technology. Results are shown across increasing flexible workload shares (20% to 100% in 20% increments), assuming a 24-hour shifting horizon, alongside a baseline scenario without flexibility. Coal, nuclear, and hydro are excluded from new capacity additions but remain eligible for retirement. Data on capacity retirements and investments at all flexibility combinations can be found in Supplementary Notes 10 and 11.

coal and nuclear. This is reflected in coal and nuclear retirements and reduced generation. Interestingly, in the Mid-Atlantic, coal retirements also increase, but the baseload support of data center flexibility leads to higher coal generation.

3 The Impact of Data Center Flexibility on Carbon Emissions

The results of our modeling show that projected 2030 data center load growth relative to a system with no data center growth increases annual CO_2 emissions by 47 Mmt (20%), 55 Mmt (58%), and 46 Mmt (24%) for the Mid-Atlantic, Texas, and WECC regions, respectively. The increase in projected emissions emphasizes the urgency of identifying strategies to reduce data centers' environmental impact, particularly in evaluating if data center flexibility can lead to emissions mitigation. However, the environmental consequences of the shifts in generation and capacity induced by data center flexibility are not straightforward. As section 2.2 showed, flexibility can simultaneously promote both renewable deployment and greater utilization of inflexible baseload generators. This dual effect raises a natural question: does data center flexibility reliably reduce emissions, or can it, under certain conditions, lead to the opposite?

While the prevailing view is that any form of demand flexibility complements the use of clean energy, our results suggest this intuition may not always hold. In this section, we examine how emissions outcomes depend on the interplay between flexibility and the underlying generation mix. We find that temporal flexibility in data centers does not always reduce total annual CO_2 emissions relative to a system without flexibility.

Fig. 7D, 7E, and 7F show the percentage reduction in emissions of systems with flexible data centers compared to systems without that flexibility. In Texas, emissions fall significantly by up to 40%. But in the Mid-Atlantic, we observe a counterintuitive result: greater data center flexibility leads to higher CO_2 emissions. In systems with high renewable penetration and limited remaining coal capacity, flexibility mostly enables greater renewable utilization and emissions reductions (i.e., Texas). Data center flexibility can have a significant impact on emissions when coupled with high renewable penetration, such that even with projected growth in data center load, a flexible system can achieve lower emissions than a system without either data center growth or flexibility (Fig. 7H). In contrast, in systems with a large share of existing coal and relatively limited VRE availability, flexibility tends to shift load toward cheap, carbon-intensive baseload generation, which raises emissions even as costs fall (Fig. 7D).

This trade-off is evident in the Mid-Atlantic. Initially, with flexibility, load shifts to hours with high VRE generation, which reduces thermal dispatch and variable O&M costs—the first effect. However, once VRE potential is exhausted, the second effect emerges with additional load shifted to hours when cheap thermal generation is available, particularly baseload coal (Fig. 5A and 5C). With full flexibility, average hourly coal utilization in the Mid-Atlantic rises from 50% to 59%. Fig. 8A and 8B show heatmaps of hourly coal utilization in the Mid-Atlantic. Without flexibility, coal ramps up in the evening and ramps down during the day when solar generation is high (Fig. 8A). The summer evenings see the largest utilization of coal as this time period coincides to the highest net loads. With flexibility, coal's output gets distributed more evenly from the evening to the early morning (Fig. 8B). The high summer evening utilization of coal is lowered, and utilization throughout the rest of the day is increased as evening data center load gets shifted to the morning (Fig. 2). A duration curve reveals that coal operates between 68 to 78% utilization for 4,526 hours (52% of the year) with flexibility, compared to just 898 hours (10%) without (Fig. 8C). This shift to steadier coal operations contributes to higher total emissions .

To further illustrate that the emissions impact is driven by the availability of renewables, we simulate an alternative set of Texas scenarios where the investment and fixed O&M costs of renewables are increased. In these systems with less economically viable renewables, we expect and find that the share of renewables decreases. More importantly, we see that the emissions-reducing effect of flexibility is reversed. Data center flexibility increases emissions relative to an inflexible system by up to 5% at renewables that cost 1.3 times the base prices (Fig. 8I). In this system, as in the Mid-Atlantic, insufficient renewable capacity to absorb flexible demand results in coal not being retired and instead becoming more heavily utilized with less ramping (Fig. 8D, 8E).

These findings show that the generation mix determines whether data center flexibility reduces or increases emissions. This also parallels results on battery storage where it can

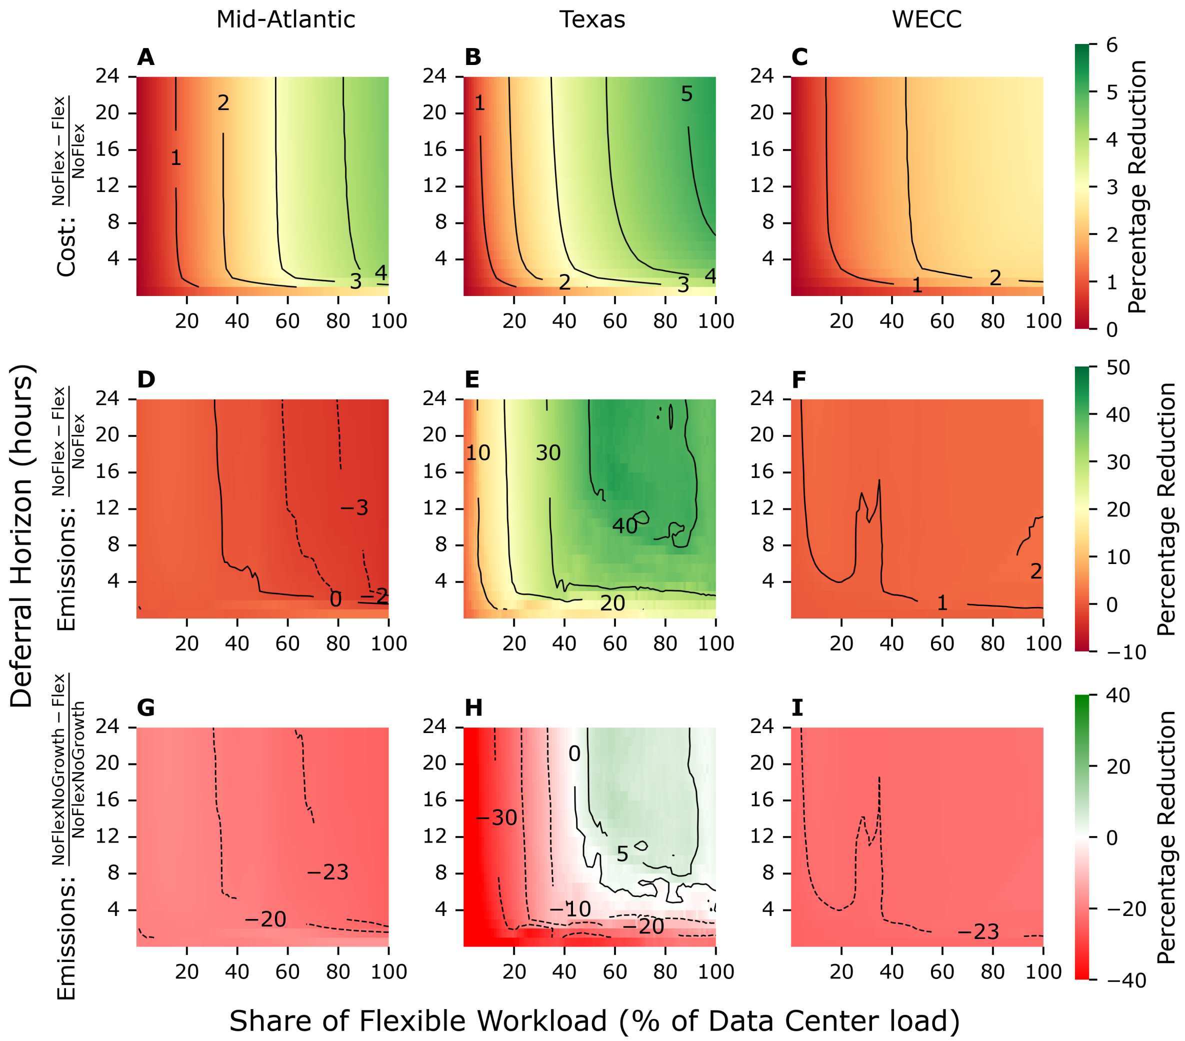

Figure 7: Cost and Emissions Reduction. Top row **(A, B, C)**: Heatmaps show the percentage reduction in total system cost from introducing flexible data centers, relative to a system without flexibility, across combinations of shifting horizon and share of flexible workload. Middle row **(D, E, F)**: Heatmaps of percentage reduction in system CO_2 emissions under the same flexibility configurations, relative to the no flexibility baseline. Bottom row **(G, H, I)**: Heatmaps of percentage reduction of system CO_2 emissions relative to a reference system with no data center flexibility and no data center load growth. Green (red) color indicates a decrease (increase) in CO_2 emissions relative to the no growth scenario. The "no growth" baseline assumes data center load in 2030 maintains the same share of total system load as in 2022. Details on load assumptions are found in Supplementary Note 9.1. Note: Color scales vary across rows.

produce counterintuitive emissions outcomes under certain conditions [26]. In systems with high renewable penetration and potential like Texas, flexibility typically aligns with wind and solar, displacing thermal generation and lowering emissions. As a consequence, the lower emissions outcome is not guaranteed and is contingent on the availability and build-out of renewable resources. The emissions impact of data center flexibility is thus not inherent to flexibility itself, but rather depends on the surrounding resource mix and investment environment. Flexibility consistently reduces system costs, but without adequate clean generation, it can inadvertently increase system emissions by reinforcing baseload coal operations.

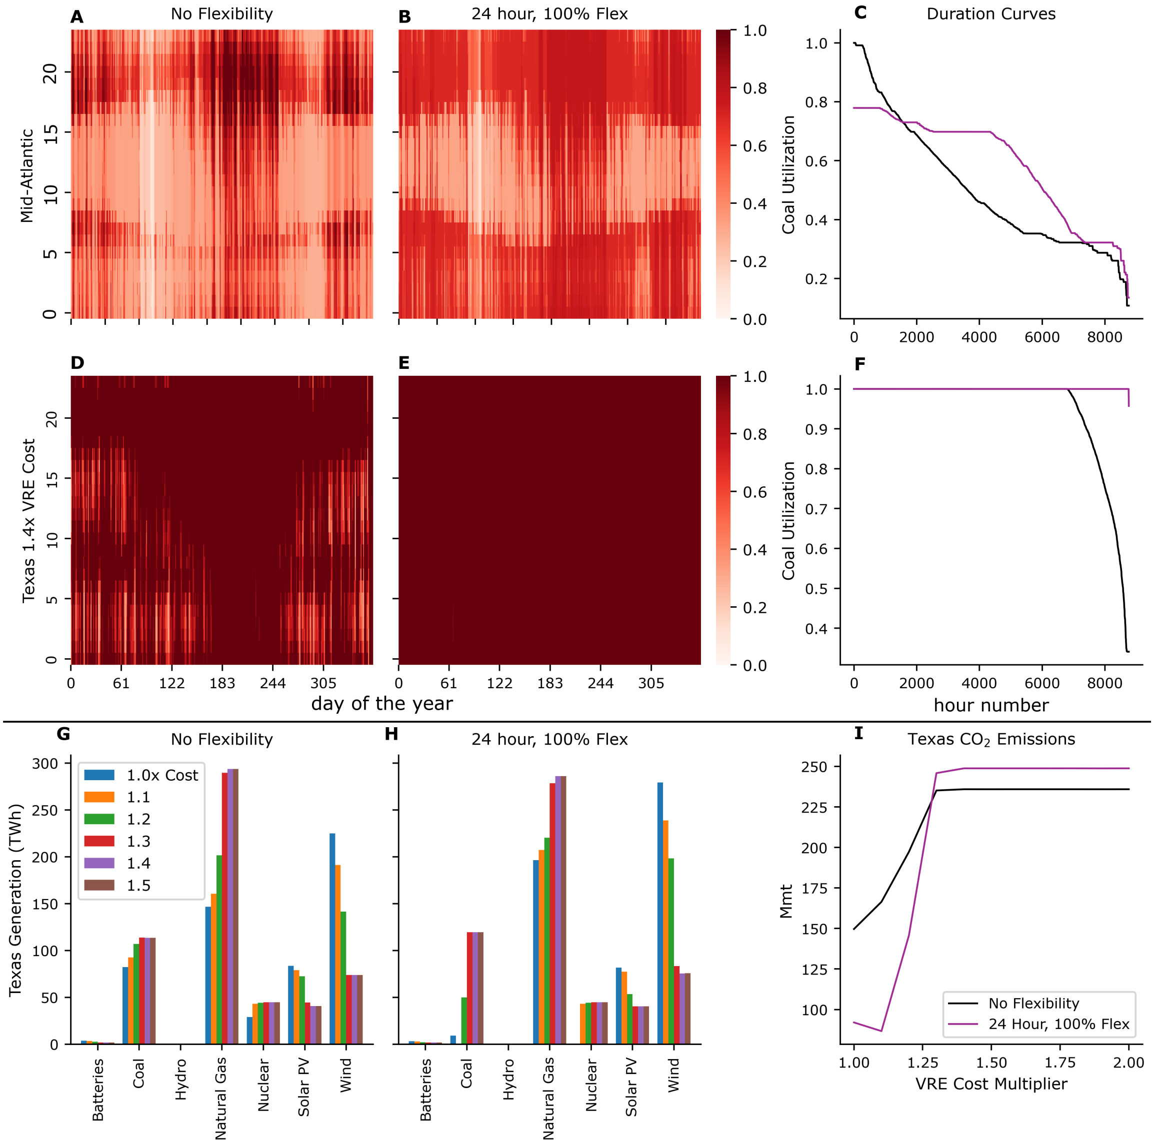

Figure 8: Coal utilization and emissions under varying VRE costs and data center flexibility. Panels **(A-B)** show heatmaps of coal utilization in the Mid-Atlantic under no flexibility and with a 24-hour shifting horizon with 100% flexible workload. Panels **(D-E)** show analogous results for Texas under a 40% increase in VRE costs. Coal utilization is normalized using a common capacity baseline—the higher of the two scenarios' remaining coal capacity—to ensure consistent comparison. Panels **(C)** and **(F)** display duration curves ranking hourly coal output in descending order across all 8,760 hours. Panels **(G-H)** show average generation by technology in Texas under varying VRE cost multipliers, which affect investment and fixed O&M costs for new wind and solar. Panel **(I)** reports total system CO_2 emissions across the same cost multipliers.The base model is represented by a cost multiplier of 1.0.

4 The Impact of Data Center Flexibility on Cost

A key output of the GenX model is the optimal total annual system cost. This consists of investment, fixed and variable operating and maintenance (O&M), fuel, startup costs, and applicable tax credits. Fig. 7A, 7B, and 7C are heatmaps of the percentage reduction in total system cost of a system with Flexible Data Centers relative to the No Flexibility scenario for the Mid-Atlantic, Texas, and WECC regions, respectively.

Data center temporal flexibility leads to lower total system costs. Increasing the shifting horizon and the share of flexible workload also increases the cost savings, with reductions of up to 4% in the Mid-Atlantic, 5% in Texas, and 2% in WECC compared to a no flexibility scenario. The cost reduction is primarily constrained by the share of flexible workload rather than the shifting horizon. The contour lines indicate that achieving specific cost reduction levels is only possible over certain ranges of the share of flexible workload. For example, in Texas, a 2% reduction in costs cannot be achieved if only 10% of data center load can be shifted to other hours, no matter how long the shifting horizon is. To achieve that 2% reduction, the share of flexible workload has to be increased to a value between 21 to 49%. The relationship between flexibility and cost reduction is also nonlinear, with an additional share of flexible workload yielding diminishing returns.

While data center flexibility reduces cost at all levels of flexibility, the source of these savings varies per region. fig. S7, S8, and S9 show the cost difference per cost component relative to a No Flexibility scenario for the Mid-Atlantic, Texas, and WECC, respectively. In the Mid-Atlantic, the savings come from a reduction of investments in new natural gas (Fig. 6A, fig. S2D) and the retirement of a portion of the existing coal plants (Fig. 6A, fig. S1B). With the lower investment in new natural gas capacity, the system avoids a generation mix that would need to spend on fuel costs that it would otherwise incur in an inflexible system. These lower investments and the additional retirements of coal generation also avoid fixed O&M costs. Note, however, that the increase in operation of the remaining non-retired coal plants with additional flexibility (fig. 6B) increases the variable O&M costs.

In Texas, the cost benefits of temporal flexibility rely on aligning data center load with cheap VRE resources. When there is a larger share of load that can be shifted to hours with high VRE availability factors, it results in lower spending on fuel and variable O&M for thermal generators. However, this increased share of renewables is only possible by building new wind capacity, increasing the investment cost (Fig. 6C, fig. S3B). At relatively small shares of flexible workload (i.e., $\leq 50\%$), Texas systems with flexibility increase fixed O&M cost, because of the non-retirement of nuclear plants (Fig. 6C, fig. S1D). Similar insights to those in Texas can be found in WECC. The only difference is that the increase in investments primarily comes from solar rather than wind, which slightly goes down as flexibility increases. This reduction in wind investments (Fig. 6E, fig. S4B) and an increase in natural gas retirements (Fig. 6E, fig. S1I) with more flexibility is what drives the lower fixed O&M costs.

5 Policy Implications

This study examines the role of data center temporal flexibility in shaping power system outcomes amid rapidly growing electricity demand. Policymakers should consider mechanisms to incentivize or require flexible data center operations, such as dynamic pricing, demand response programs, or performance-based incentives tied to load-shifting capabilities. Flexibility enables cost-effective grid management by reducing peak demand, facilitating renewable integration, and lowering system costs. However, its emissions impact is highly context-dependent: in grids with abundant renewables, flexibility supports decarbonization, while in fossil-heavy systems, it may increase emissions by extending the operational life of baseload coal. Therefore, unlocking the full benefits of data center flexibility requires coupling it with strong clean energy policies such as carbon pricing, investment incentives, or renewable portfolio standards, to ensure emissions reductions accompany cost savings.

6 Conclusion

Data centers are projected to account for a substantial share of U.S. electricity demand in the coming years. This study examines how the temporal flexibility of data center load can alter power systems in response to this rapid demand growth. Specifically, we analyze how flexibility affects grid operations, investment and retirement decisions, system costs, and emissions outcomes. Our results show that temporal flexibility can substantially reshape the power system and mitigate the challenges associated with rising data center demand.

Flexible data centers are able to adjust their load profiles in response to system conditions, enabling more efficient power system operations. We find that the cost-optimal load-shifting strategy tends to flatten the net load curve by reducing demand during peak hours and shifting it toward periods of low net load. This operational adjustment reduces reliance on peaker plants and enables better utilization of low-cost renewable energy, resulting in lower system costs compared to inflexible demand. The magnitude of these savings depends on both the share of flexible workload and the shifting horizon, with more flexibility leading to lower cost. To capture these benefits, policy-makers should consider the necessary legislative support or market-based incentives that promote temporal flexibility. Our analysis provides a framework for identifying combinations of flexibility parameters that yield equivalent cost savings, which could inform effective policy design and implementation.

Data center flexibility also has important implications for long-term capacity investment and retirement decisions. These impacts depend on the existing generation mix, renewable resource availability, and the evolving costs of clean energy technologies. In general, temporal flexibility encourages investment in wind and solar by shifting demand to hours with low marginal costs, which often align with high renewable production. In high renewable systems like Texas, where renewables comprise 50% of generation, flexibility leads to increased investment in VRE capacity, accelerated retirement of baseload capacity, and increased reliance on flexible thermal generation to manage intermittency. In contrast, in regions with lower renewable shares, such as the Mid-Atlantic, flexibility can increase the utilization of baseload plants like coal even if some retirements still occur.

As a result, the emissions impact of data center flexibility is highly context-dependent. In systems with abundant and cost-competitive renewables, flexibility supports decarbonization by aligning demand with clean energy availability. However, in fossil-heavy grids, flexibility can increase emissions by increasing coal generation. Complementary policies that accelerate renewable deployment or lower clean energy costs are needed alongside flexibility incentives to ensure that the emissions-reduction potential of data center temporal flexibility is fully realized.

Overall, we find that the rapid growth in data center load has significant implications for power system planning and operations. Its inherent capability to temporally shift load can be used to mitigate the costs and infrastructure needs associated with this increased demand. To ensure that this flexibility supports decarbonization, it must be deployed alongside strong clean energy policies that accelerate renewable deployment. When aligned with such policies, data center flexibility can act as a valuable grid asset, lowering peak capacity, reducing cost, and facilitating the integration of variable renewable resources.

7 Limitations

The model optimizes a representative year of operation where the investment, dispatch, and data shifting decisions are made by a centralized entity. Decentralized decision making of carbon-aware loads is not accounted for but is studied by [19]. We assume that data center load is constant and that it can freely shift its load within the hourly horizon without the need for ramping and without disruption to operations. The costs associated with data center load shifting are not modeled, as we focus on workload advancement and delay, where direct costs are typically small or difficult to quantify due to their dependence on user preferences and application-specific requirements. The stochastic variability of VRE and load is not considered by assuming exogenous, deterministic time series inputs. We also assume that there are no restrictions on the amount of new capacity that can be built and that this capacity can be built by the model year. The model can then be thought of as a stylized and idealized U.S. power system with data centers.

8 Methods

8.1 GenX

The analysis uses the Capacity Expansion model, GenX. Details and documentation on the GenX model can be found in `https://genxproject.github.io/GenX.jl/dev/`. The specific implementation used for this work can be found in the Supplementary Code. GenX is a least-cost mixed integer linear programming (MILP) model that co-optimizes generation and transmission investments and dispatch decisions among pre-defined zones within the power system. The optimization accounts for capital, operational, and fuel costs, generator technical operating characteristics, capacity factors for renewables, and demand information. The objective is to minimize annual system cost, which is the sum of investment in generation and storage, fixed and variable operating and maintenance costs, new transmission investment costs, fuel and startup costs, accounting for tax credits and other incentives, if any. GenX assumes a representative year of operation and perfect foresight of hourly demand and capacity factor data for renewables supply in its dispatch decisions. We source input data from PowerGenome [27] which is a data processing software that aggregates data from publicly available sources such as NREL's ATB for cost data [28], NREL's EFS for demand data [29, 30] and EIA's Form-860 [31] for existing generator data.

8.2 Modeling Data Center Temporal Flexibility

We modified the GenX code-base to include temporal flexibility of data centers. GenX already has the capability to represent flexible demand resources built into the base model, but we included constraints to constrain the amount that can be shifted to and from an hour, given the capacity of a data center. For our model, data centers that can shift load temporally are modeled similarly to storage assets that can "store" demand from each hour and must be deployed elsewhere within the shifting horizon.

Consider a representative data center with capacity C and data center load L_t at hour t. Temporal flexibility is defined by two parameters: First is the shifting horizon h, which is the number of hours before or after the originally scheduled hour in which data center load can be deferred and advanced. Second, the share of flexible workload s, $(0 \leq s \leq 1)$, which is the fraction of total data center demand that can be shifted to other hours.

Let $Y_t \in \mathbb{R}$ be the amount of data center demand that is yet to be satisfied at hour t. A positive value means that there exists demand that is delayed, and a negative value means that demand has been advanced. Let $S_t \geq 0$ be the amount of shifted data center load that is satisfied at hour t. Finally, let $D_t \geq 0$ be the amount of data center load that was originally allocated to hour t but is deferred and will be satisfied in a future hour t', where $t < t' \leq t + h$. We include the following constraint to model data center flexibility:

Data Center Load Balancing Constraint: the amount of data center demand that is yet to be satisfied is equal to the amount from the previous hour, less what is satisfied in the current hour, plus any deferrals.

$$Y_t = Y_{t-1} - S_t + D_t \tag{1}$$

Maximum Time to Delay Demand Constraint: The amount of data center load that is satisfied in the next h hours from time t must be greater than or equal to the amount that is yet to be satisfied by time t.

$$\sum_{i=t+1}^{t+h} S_i \geq Y_t, \forall t \tag{2}$$

Maximum Time to Advance Demand Constraint: The amount of data center load that is deferred in the next h hours from time t must be greater than the advanced demand (negative of Y_t) in hour t.

$$\sum_{i=t+1}^{t+h} D_i \geq -Y_t, \forall t \tag{3}$$

Maximum amount of demand Deferred Constraint: The amount of demand that can be deferred in each hour must be less than the share of flexible workload.

$$D_t \leq sL_t, \forall t \tag{4}$$

Maximum Data Center Load Satisfied : Shifted data center load that is satisfied during an hour must be less than or equal to the capacity of the data center net of the deferred demand.

$$S_t \leq C - L_t + D_t \tag{5}$$

References

[1] E. Masanet, A. Shehabi, N. Lei, S. Smith, J. Koomey, Recalibrating global data center energy-use estimates. *Science* **367** (6481), 984–986 (2020), doi:10.1126/science.aba3758, https://www.science.org/doi/10.1126/science.aba3758.

[2] IEA, *Energy and AI*, Tech. rep., International Energy Agency (IEA) (2025), https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai.

[3] A. Shehabi, *et al.*, *United States Data Center Energy Usage Report*, Tech. Rep. LBNL-2001637 (2024), doi:10.2172/1372902, https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf.

[4] NERC, *Incident Review Considering Simultaneous Voltage-Sensitive Load Reductions*, Tech. rep., North American Electric Reliability Corporation (NERC) (2025), https://www.nerc.com/pa/rrm/ea/Documents/Incident_Review_Large_Load_Loss.pdf.

[5] EPRI, *Powering Data Centers: U.S. Energy System and Emissions Impacts of Growing Loads*, Tech. Rep. 3002031198 (2024), https://www.epri.com/research/products/000000003002031198.

[6] PJM, *PJM Long Term Load Forecast Report*, Tech. rep., Pennsylvania-New Jersey-Maryland (PJM) Interconnection (2025), https://www.pjm.com/-/media/DotCom/library/reports-notices/load-forecast/2025-load-report.pdf.

[7] A. Wierman, Z. Liu, I. Liu, H. Mohsenian-Rad, Opportunities and challenges for data center demand response, in *International Green Computing Conference* (IEEE, DALLAS, TX, USA) (2014), pp. 1–10, doi:10.1109/IGCC.2014.7039172, http://ieeexplore.ieee.org/document/7039172/.

[8] Y. Zhou, A. Paredes, C. Essayeh, T. Morstyn, AI-focused HPC Data Centers Can Provide More Power Grid Flexibility and at Lower Cost (2024), doi:10.48550/arXiv.2410.17435, http://arxiv.org/abs/2410.17435, arXiv:2410.17435 [eess].

[9] A. Radovanovic, Inside Google: Our data centers now work harder when the sun shines and wind blows (2020), https://blog.google/inside-google/infrastructure/data-centers-work-harder-sun-shines-wind-blows/.

[10] A. Radovanović, *et al.*, Carbon-Aware Computing for Datacenters. *IEEE Transactions on Power Systems* **38** (2), 1270–1280 (2023), doi:10.1109/TPWRS.2022.3173250, https://ieeexplore.ieee.org/document/9770383/.

[11] T. Han, *et al.*, Designing and regulating clean energy data centres. *Nature Reviews Clean Technology* (2025), doi:10.1038/s44359-025-00062-0, https://www.nature.com/articles/s44359-025-00062-0.

[12] W. Zhang, V. M. Zavala, Remunerating space–time, load-shifting flexibility from data centers in electricity markets. *Applied Energy* **326**, 119930 (2022), doi:10.1016/j.apenergy.2022.119930, https://linkinghub.elsevier.com/retrieve/pii/S0306261922011874.

[13] L. Rao, X. Liu, M. D. Ilic, J. Liu, Distributed Coordination of Internet Data Centers Under Multiregional Electricity Markets. *Proceedings of the IEEE* **100** (1), 269–282 (2012), doi:10.1109/JPROC.2011.2161236, http://ieeexplore.ieee.org/document/5989839/.

[14] K. Kim, F. Yang, V. M. Zavala, A. A. Chien, Data Centers as Dispatchable Loads to Harness Stranded Power. *IEEE Transactions on Sustainable Energy* **8** (1), 208–218 (2017), doi:10.1109/TSTE.2016.2593607, http://ieeexplore.ieee.org/document/7517380/.

[15] J. Zheng, A. A. Chien, S. Suh, Mitigating Curtailment and Carbon Emissions through Load Migration between Data Centers. *Joule* **4** (10), 2208–2222 (2020), doi:10.1016/j.joule.2020.08.001, https://linkinghub.elsevier.com/retrieve/pii/S2542435120303470.

[16] L. Liu, X. Shen, Z. Chen, Q. Sun, R. Wennersten, Optimal Energy Management of Data Center Micro-Grid Considering Computing Workloads Shift. *IEEE Access* **12**, 102061–102075 (2024), doi:10.1109/ACCESS.2024.3432120, https://ieeexplore.ieee.org/document/10606267/.

[17] EPRI, *Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption*, Tech. Rep. 3002028905 (2024), https://www.epri.com/research/products/3002028905.

[18] I. Riepin, T. Brown, V. M. Zavala, Spatio-temporal load shifting for truly clean computing. *Advances in Applied Energy* **17**, 100202 (2025), doi:10.1016/j.adapen.2024.100202, https://linkinghub.elsevier.com/retrieve/pii/S2666792424000404.

[19] W. Jiang, O. Huber, M. C. Ferris, L. Roald, Can Carbon-Aware Electric Load Shifting Reduce Emissions? An Equilibrium-Based Analysis (2025), doi:10.48550/arXiv.2504.07248, http://arxiv.org/abs/2504.07248, arXiv:2504.07248 [eess].

[20] T. H. Norris, T. Profeta, D. Patino-Echeverri, A. Cowie-Haskell, *Rethinking Load Growth Assessing the Potential for Integration of Large Flexible Loads in US Power Systems*, Tech. rep., Durham, NC: Nicholas Institute for Energy, Environment & Sustainability, Duke University. (2025), https://nicholasinstitute.duke.edu/publications/rethinking-load-growth.

[21] J. D. Jenkins, N. A. Sepulveda, Enhanced decision support for a changing electricity landscape: The GenX configurable electricity resource capacity expansion model (2017), https://energy.mit.edu/wp-content/uploads/2017/10/Enhanced-Decision-Support-for-a-Changing-Electricity-Landscape.pdf.

[22] A. Manocha, G. Mantegna, N. Patankar, J. D. Jenkins, Reducing transmission expansion by co-optimizing sizing of wind, solar, storage and grid connection capacity. *Environmental Research: Energy* **2** (1), 015011 (2025), doi:10.1088/2753-3751/adafab, https://iopscience.iop.org/article/10.1088/2753-3751/adafab.

[23] A. Botterud, C. R. Knittel, J. E. Parsons, J. R. L. Senga, S. D. Story, Policy-Driven Transmission Expansion in the United States: Implications for Costs, Emissions, and Reliability (2025), doi:10.21203/rs.3.rs-6148704/v1, https://www.researchsquare.com/article/rs-6148704/v1.

[24] N. A. Sepulveda, J. D. Jenkins, A. Edington, D. S. Mallapragada, R. K. Lester, The design space for long-duration energy storage in decarbonized power systems. *Nature Energy* **6** (5), 506–516 (2021), doi:10.1038/s41560-021-00796-8, https://www.nature.com/articles/s41560-021-00796-8.

[25] N. A. Sepulveda, J. D. Jenkins, F. J. De Sisternes, R. K. Lester, The Role of Firm Low-Carbon Electricity Resources in Deep Decarbonization of Power Generation. *Joule* **2** (11), 2403–2420 (2018), doi:10.1016/j.joule.2018.08.006, https://linkinghub.elsevier.com/retrieve/pii/S2542435118303866.

[26] J. E. Bistline, D. T. Young, Emissions impacts of future battery storage deployment on regional power systems. *Applied Energy* **264**, 114678 (2020), doi:10.1016/j.apenergy.2020.114678, https://linkinghub.elsevier.com/retrieve/pii/S0306261920301902.

[27] G. Schivley, Power Genome, https://github.com/PowerGenome/ (2023).

[28] NREL, NREL Annual Technology Basline 2022, https://atb.nrel.gov/electricity/2022/data (2022).

[29] T. T. Mai, *et al.*, *Electrification Futures Study: Scenarios of Electric Technology Adoption and Power Consumption for the United States*, Tech. Rep. NREL/TP–6A20-71500, 1459351 (2018), doi:10.2172/1459351, http://www.osti.gov/servlets/purl/1459351/.

[30] T. Mai, *et al.*, Electrification Futures Study Load Profiles. *National Renewable Energy Laboratory* (2020), doi:10.7799/1593122.

[31] EIA, Form EIA-860 detailed data with previous form data, https://www.eia.gov/electricity/data/eia860/ (2022).

[32] EPA, *Documentation for EPA's Power Sector Modeling Platform v6 - Summer 2021 Reference Case*, Tech. rep., United States Environmental Protection Agency (2021), https://www.epa.gov/power-sector-modeling/documentation-epas-power-sector-modeling-platform-v6-summer-2021-reference.

[33] C. T. M. Clack, A. Choukulkar, B. Coté, S. A. McKee, *Dataset Overviews: Renewable Generation, Electric Demand, Transmission Line Ratings & Losses, and Climate Change*, Tech. rep., Vibrant Clean Energy, LLC (2020).

Acknowledgements

The authors would like to acknowledge J.E. Parsons for his insights and discussion on models and assumptions. All views expressed in this paper are those of the authors and do not necessarily reflect the views of acknowledged individuals or affiliated institutions.

Funding

This research was supported by funding from the Future Energy Systems Center of the MIT Energy Initiative (MITEI).

Author Contributions

Conceptualization: All authors contributed to the conceptualization of the study. Data Curation: J.R.L.S. and S.W. curated and gathered data. Software: J.R.L.S. and S.W. created the GenX model version used in performing the analysis. Methodology: J.R.L.S. and S.W. designed the methodology. Investigation: All authors contributed to the analysis of the data and results. Visualization: J.R.L.S. and S.W. created the visualization and plots. Writing – original draft/review and editing: All authors contributed to the writing of the manuscript.

Competing Interests

The authors declare no competing interests.

Additional Information

Supplementary information and Data Availability

We provide a Supplementary Document that discusses additional methods and assumptions. Data necessary to replicate the results in the paper are provided in: https://github.com/JRLSenga/DataCenter_TemporalFlex

Code Availability

Replication software is provided in: https://github.com/JRLSenga/DataCenter_TemporalFlex

Correspondence and Request for Materials

Requests and questions should be directed to Juan Ramon L. Senga

Supplementary Materials for
Flexible Data Centers and the Grid: Lower Costs, Higher Emissions?

Christopher R. Knittel, Juan Ramon L. Senga, Shen Wang

9 Input Assumptions

9.1 Regions and Load

Table S1 shows the number of zones per region along with the total and hourly data center load, and the total non-data center load per zone. Zones are IPM regions from the EPA. Data center load is sourced from [5] while non-data center load is sourced from [27] who source the base hourly demand from NREL's EFS [30]. We use the "High Growth" scenario as our base case.

Table S1: Base Data Center and Non-Data Center Load per Model Zone (in MWh)

Region	Zone	Zone Number	Data Center Load (No Growth)	MWh per Hour (Base)	Data Center Load (Base)	Non-Data Center Load	Total Load (Base)	Share of Data Center Load
Texas	ERC_REST	1	19,196,070	15,415	135,033,849	399,018,968	534,052,817	25.3%
Texas	ERC_WEST	2	852,970	685	6,000,174	17,730,244	23,730,418	25.3%
Texas	ERC_PHDL	3	268,607	216	1,889,506	5,583,406	7,472,912	25.3%
Mid-Atlantic	PJM_AP	1	1,838,296	600	5,253,312	56,335,623	61,588,935	8.5%
Mid-Atlantic	PJM_ATSI	2	1,415,551	2,634	23,076,374	88,176,254	111,252,627	20.7%
Mid-Atlantic	PJM_Dom	3	30,409,937	9,605	84,137,386	88,425,299	172,562,685	48.8%
Mid-Atlantic	PJM_EMAC	4	5,470,400	1,280	11,212,385	170,081,177	181,293,562	6.2%
Mid-Atlantic	PJM_PENE	5	678,669	221	1,939,439	20,798,210	22,737,648	8.5%
Mid-Atlantic	PJM_SMAC	6	107,757	10	90,179	67,240,255	67,330,434	0.1%
Mid-Atlantic	PJM_WMAC	7	1,394,599	455	3,985,355	42,738,269	46,723,625	8.5%
Mid-Atlantic	PJM_West	8	15,408,103	6,131	53,703,310	170,446,981	224,150,291	24.0%
WECC	WEC_CALN	1	3,618,438	568	4,977,247	94,177,173	99,154,420	5.0%
WECC	WEC_LADW	2	2,710,231	426	3,727,987	70,539,254	74,267,241	5.0%
WECC	WEC_SDGE	3	899,730	141	1,237,600	23,417,298	24,654,898	5.0%
WECC	WECC_SCE	4	2,838,787	446	3,904,820	73,885,197	77,790,016	5.0%
WECC	WECC_MT	5	641,553	59	518,599	16,650,987	17,169,585	3.0%
WECC	WEC_BANC	6	510,699	80	702,479	13,291,990	13,994,469	5.0%
WECC	WECC_ID	7	141,588	13	118,052	24,698,492	24,816,543	0.5%
WECC	WECC_NNV	8	1,038,416	149	1,303,616	10,911,131	12,214,747	10.7%
WECC	WECC_SNV	9	2,827,400	405	3,549,487	29,708,849	33,258,336	10.7%
WECC	WECC_UT	10	3,030,388	358	3,133,353	36,427,790	39,561,143	7.9%
WECC	WECC_PNW	11	11,987,277	8,449	74,011,855	165,076,398	239,088,252	31.0%
WECC	WECC_CO	12	1,886,168	310	2,715,211	69,022,399	71,737,610	3.8%
WECC	WECC_WY	13	1,626,323	146	1,278,063	23,664,710	24,942,773	5.1%
WECC	WECC_AZ	14	6,871,657	7,242	63,441,221	85,613,629	149,054,850	42.6%
WECC	WECC_NM	15	612,040	58	506,741	40,742,002	41,248,743	1.2%
WECC	WECC_IID	16	53,403	8	73,457	1,389,919	1,463,376	5.0%

The EPRI report [17] indicates what % of each state's 2023 electricity demand was for Data Centers. We assume that without data center growth, load will have the same % share of data center load in 2030. We then calculate the additional data center load on top of the % share in the base case. We provide an illustrative example for Texas Zone 1:

1. 4.59% of Texas' load in 2023 is for Data Centers.
2. 2030 NREL EFS Demand for Texas Zone 1 is 418.2 TWh. 418 TWh $\times$ 4.59% = 19.2 TWh of Base Data Center Load in 2030
3. 418.2 TWh $-$ 19.2 TWh = 399 TWh of non-Data Center Load in 2030

4. There is a high growth forecast of 25.28% of Texas 2030 load will come from Data Centers
5. 399 TWh / (100% − 25.28%) = 534 TWh total Texas Zone 1 Load.
6. 534 TWh × 25.28% = 135 TWh of Data Center load in 2030.
7. Since we assume that data center load is constant per hour, we divide the 135 TWh by 8760 hours.

9.2 Generators

Our model includes existing capacity generators as well as a set of new technologies that can be deployed. Existing generation capacity is sourced from EIA Form-860 and aggregated through PowerGenome [27]. Details can be found in Table S2. Investment, operating, and maintenance costs for new generators can be found in Table S3. Fixed O&M costs, CAPEX, and WACC for new capacity are taken as average values from NREL ATB 2022 from the years 2023 to 2030 [28]. The investment costs vary based on regional multipliers. Meanwhile, cost assumptions for existing plants use the basis year 2020, with variation assumptions from PowerGenome depending on the start year of operation. Production and tax credits associated with the Inflation Reduction Act are also implemented in the model.

Table S2: Capacity of Existing Generators per Technology in each Region (in GW)

Technology	Mid-Atlantic	Texas	WECC
Batteries	0.25	4.40	13.69
Conventional Hydroelectric	3.35	0.54	50.25
Conventional Steam Coal	39.24	13.63	22.08
Hydroelectric Pumped Storage	5.21	0.00	5.05
Natural Gas Fired Combined Cycle	55.90	41.75	52.75
Natural Gas Fired Combustion Turbine	22.26	11.16	23.57
Nuclear	22.80	5.12	7.42
Onshore Wind Turbine	5.97	34.05	32.22
Solar Photovoltaic	10.79	20.83	38.27

Table S3: New Technology Investment and Operation Cost Assumptions in 2030

	Capex ($/MW)	Capital Recovery Period (years)	WACC	Investment Cost ($/MW-yr)	Fixed O&M ($/MW-yr)	Variable O&M ($/MWh)
Natural Gas Combined Cycle	932,813	15	3.56%	81,708	28,000	2
Solar Photovoltaic	913,819	20	2.50%	58,794	22,623	-
Onshore Wind Turbine	1,131,578	20	3.06%	76,816	40,367	-
Battery	250,489	20	2.50%	16,116	6,262	-

9.3 Transmission

We source current transfer capabilities per line between each IPM zone from the EPA's Power Sector Modeling Platform v6—2021 Summer Reference Case [32]. We assume a pipeline flow model such that the amount of transmission that can flow between two zones is only restricted by the capacity of the line.

9.4 Net Imports

Texas and WECC are fairly isolated as model regions within the continental U.S.. The impacts of electricity exchange with neighboring regions on these two regions via transmission lines is therefore minimal. However, the Mid-Atlantic is extensively connected to other neighboring regions such as the Midwest, Southeast, and New York. To account for this in the model, we sourced hourly net import data for the Mid-Atlantic from EIA's Grid Monitor Dashboard for 2022 (https://www.eia.gov/electricity/gridmonitor/dashboard/electric_overview/US48/US48). Within the dataset, Mid-Atlantic's (MIDA) net imports are aggregated to hourly exchange with CAR (Carolinas), MIDW (Midwest), NY (New York), and TEN (Tennessee). To allocate the net import to model zones, we first determine whether the model zone has existing transmission capacity with the EIA regions. If there is, we calculate the percentage allocation as the total load of the model zone divided by the total load of all model zones connected to the region (see Table S4). Each model zone's net import is thus the hourly net import from the EIA data set multiplied by this allocation percentage.

Table S4: Net Import Allocation Percentage

	Zone Number	CAR	MIDW	NY	TEN
PJM_AP	1	0%	0%	0%	0%
PJM_ATSI	2	0%	0%	0%	0%
PJM Dom	3	43%	0%	0%	0%
PJM_EMAC	4	0%	0%	89%	0%
PJM_PENE	5	0%	0%	11%	0%
PJM_SMAC	6	0%	0%	0%	0%
PJM_WMAC	7	0%	0%	0%	0%
PJM_West	8	57%	100%	0%	100%

9.5 CO_2 Emissions Factors

Emission factors are available for Natural Gas and Coal. CO_2 is generated per MMBtu of fuel consumed. We assume 0.09552 mtCO_2/MMBtu and 0.05306 mtCO_2/MMBtu for coal and natural gas, respectively. Table S5 shows the average heat rates for existing generators.

Table S5: Average Heat Rates of Existing Generators (in MMbtu/MWh)

	Mid-Atlantic	Texas	WECC
Conventional Steam Coal	12.27	11.28	10.97
Natural Gas Fired Combined Cycle	8.19	8.69	8.07
Natural Gas Fired Combustion Turbine	13.12	12.10	12.08
Nuclear	10.45	10.45	10.45

9.6 Supply Curves

Supply curves for renewables are sourced from PowerGenome [27], who source the data from Vibrant Clean Energy's data sets [33].

9.7 Fuel Costs

Fuel costs are sourced from EIA's Annual Energy Outlook (AEO) 2022 for the year 2030. The individual zones are matched to the AEO regions through PowerGenome. Fuel cost information can be found below in Table S6.

Table S6: Fuel Cost (in \$/MMbtu)

Model Region	Fuel	AEO Region	Fuel Name	Model Zone	$/MMBtu
Mid-Atlantic	Coal	South Atlantic	south_atlantic_coal	PJM_AP, PJM_Dom	2.40
		East North Central	east_north_central_coal	PJM_ATSI, PJM _West	1.84
		Middle Atlantic	middle_atlantic_coal	PJM_EMAC, PJM_PENE, PJM_SMAC, PJM_WMAC	2.25
	Natural Gas	South Atlantic	south_atlantic_naturalgas	PJM_AP, PJM_Dom	4.10
		East North Central	east_north_central_naturalgas	PJM_ATSI, PJM _West	3.41
		Middle Atlantic	middle_atlantic_naturalgas	PJM_EMAC, PJM_PENE, PJM_SMAC, PJM_WMAC	3.19
	Uranium	South Atlantic	south_atlantic_uranium	PJM_Dom	0.71
		East North Central	east_north_central_uranium	PJM_ATSI, PJM _West	0.71
		Middle Atlantic	middle_atlantic_uranium	PJM_EMAC, PJM_SMAC, PJM_WMAC	0.71
Texas	Coal	West South Central	west_south_central_coal	ERC_REST	1.73
	Natural Gas	West South Central	west_south_central_naturalgas	ERC_REST	3.49
	Uranium	West South Central	west_south_central_uranium	ERC_REST, ERC_PHDL, ERC_WEST	0.71
WECC	Coal	Mountain	mountain_coal	WECC_AZ, WECC_CO, WECC_MT, WECC_NM, WECC_NNV, WECC_UT, WECC_WY	1.55
		Pacific	pacific_coal	WECC_PNW, WECC_SCE	2.02
	Natural Gas	Mountain	mountain_naturalgas	WECC_AZ, WECC_CO, WECC_ID, WECC_MT, WECC_NM, WECC_NNV, WECC_SNV, WECC_UT, WECC_WY	4.00
		Pacific	pacific_naturalgas	WECC_IID, WECC_PNW, WECC_SCE, WEC_BANC, WEC_CALN, WEC_LADW, WEC_SDGE	3.88
	Uranium	Mountain	mountain_uranium	WECC_AZ	0.71
		Pacific	pacific_uranium	WECC_PNW, WEC_CALN	0.71

10 Capacity Retirements

Fig. S1 shows the impact of flexible data centers on retirement decisions for nuclear, coal, and natural gas generators across the three regions. In the Mid-Atlantic and WECC, retirement decisions appear largely insensitive to data center flexibility across all three fuel types, although coal and natural gas show an increase in retirements for the Mid-Atlantic (Fig. S1B) and WECC (Fig. S1I, respectively. Nuclear retirements for both regions (Fig. S1A, Fig. S1G) and natural gas retirements for the Mid-Atlantic (Fig. S1C) remain similar regardless of flexibility levels. This suggests that the generation mix and system constraints in these regions limit the ability of flexible demand to displace firm capacity.

In contrast, Texas sees different retirement patterns. As both the share of flexible workload and the shifting horizon increase, significant generator retirements are observed, particularly for nuclear (Fig. S1D) and coal (Fig. S1E) resources. At high flexibility levels and long shifting horizons (e.g., $\geq 80\%$ flexible workload and 24-hour shifting horizon), nuclear and coal retirements approach or exceed 80% and 90%, respectively. This indicates that flexible data center

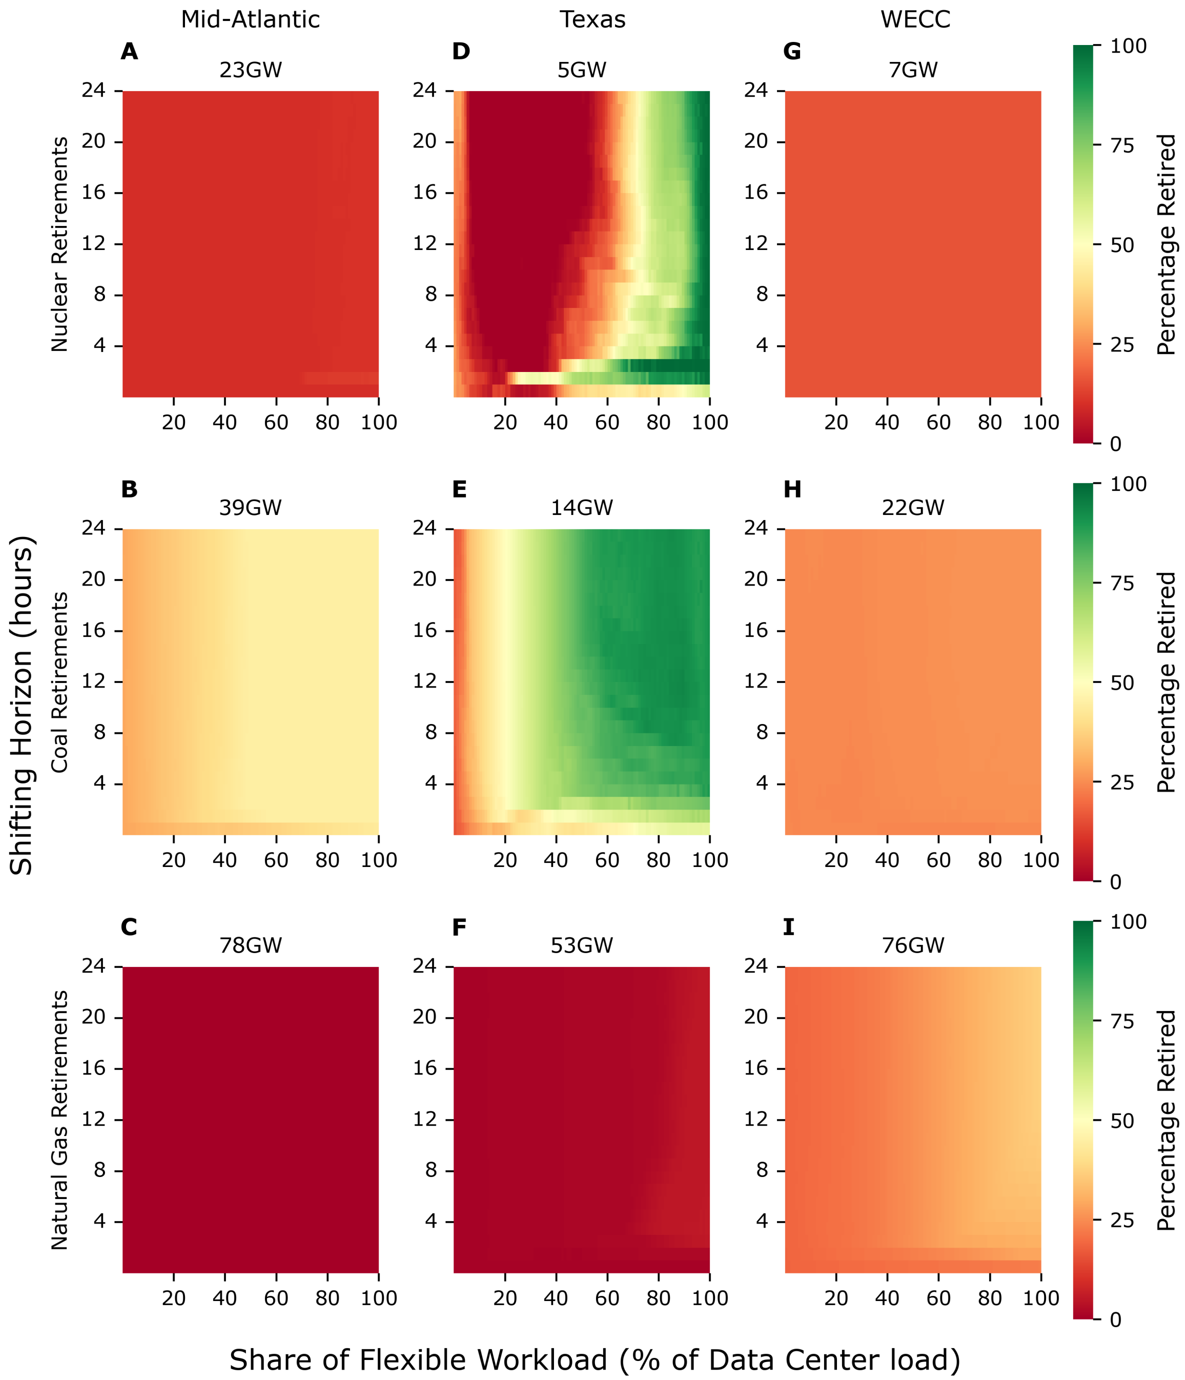

Figure S1: Percentage Retirement per Technology for each Region. Panels show the percentage of existing capacity retired for nuclear **(A, D, G)**, coal **(B, E, H)**, and natural gas **(C, F, I)** across combinations of data center shifting horizon (1 to 24 hours) and flexible workload share (1% to 100%). Values displayed at the top of each heatmap indicate the initial installed capacity for the corresponding technology in each region.

demand in Texas has a capacity substitution effect, particularly for baseload resources. This is due to the region's high penetration of high-quality renewables. Natural gas retirements in Texas (Fig. S1F) remain low overall, with only marginal increases under the highest flexibility levels.

11 Capacity Investments

Fig. S2, S3, and S4 show the effect of flexible data center operations on new capacity investments in the Mid-Atlantic, Texas, and WECC, respectively.

In the Mid-Atlantic, solar (Fig. S2A) investments increase significantly with higher levels of data center flexibility, particularly when both the share of flexible workload exceeds 60% and the shifting horizon extends beyond 2 hours. Under these conditions, solar capacity reaches over 52 GW. This reflects the ability of the system to align flexible data center demand with solar output. This increases the value of solar in balancing load within a day, which leads to more investments. Investments in wind capacity (Fig. S2B), in contrast, remain unchanged across the different combinations of flexibility. This suggests that the temporal characteristics of flexible demand do not substantially affect wind investment decisions in the Mid-Atlantic. Battery investments (Fig. S2C) remain relatively limited, with total new capacity not exceeding 2 GW. Notably, battery deployment decreases once the flexible workload share exceeds 60%. This decline can be attributed to functional competition between batteries and flexible data center loads, as both serve similar roles in providing temporal flexibility to the power system. As flexible data center operations become more prominent, they can displace the need for additional storage by shifting load in response to system conditions. Investments in new natural gas capacity (Fig. S2D) decrease as data center flexibility increases. With high levels of both the share of flexible workload and long shifting horizons, natural gas investment drops from over 14 GW to below 6 GW, indicating that flexible demand can substitute for peaking gas capacity by reducing peak load and system ramping needs.

In Texas, a higher level of data center flexibility leads to an increase in wind investments from approximately 46 GW to over 58 GW (Fig. S3B). Solar capacity (Fig. S3A) shows a more modest and stable pattern with only a slight increase from 19 GW to 22.5 GW. The flatter gradient suggests that while solar remains valuable, its incremental benefit diminishes in the presence of high data center flexibility. This is due to the temporal mismatch between peak solar output and peak system stress in Texas. Battery and natural gas investments (Fig. S3C, S3D) remain negligible across the entire flexibility space, with capacities barely exceeding 0.05 GW. Thus, in Texas, data center temporal flexibility can strongly incentivize wind deployment, supporting a more renewable-heavy system configuration. Similar to PJM, this indicates that flexible data center operations are effectively substituting for both short-duration storage and fast-ramping thermal resources.

In WECC, both solar and wind capacity exhibit noticeable increases as data center flexibility increases. Solar investments (Fig. S4A) increase steadily from 38 GW to over 43 GW, particularly when the share of flexible workload exceeds 40% and the shifting horizon is greater than 8 hours. Similarly, wind capacity (Fig. S4B) shows an upward trend, growing from 13 GW to 14 GW under higher flexibility. Just like in the Mid-Atlantic and Texas, these patterns also suggest that flexible data center demand in WECC increases the economic viability of variable renewables. In contrast, there are no new battery investments (Fig. S4C) across all flexibility scenarios. Natural gas (Fig. S4D) investments are small even without flexibility at around 1.75 GW, and decline to almost no investments as flexibility increases. This indicates that the load flexibility is sufficient to meet system balancing needs, diminishing the marginal value of new natural gas plants.

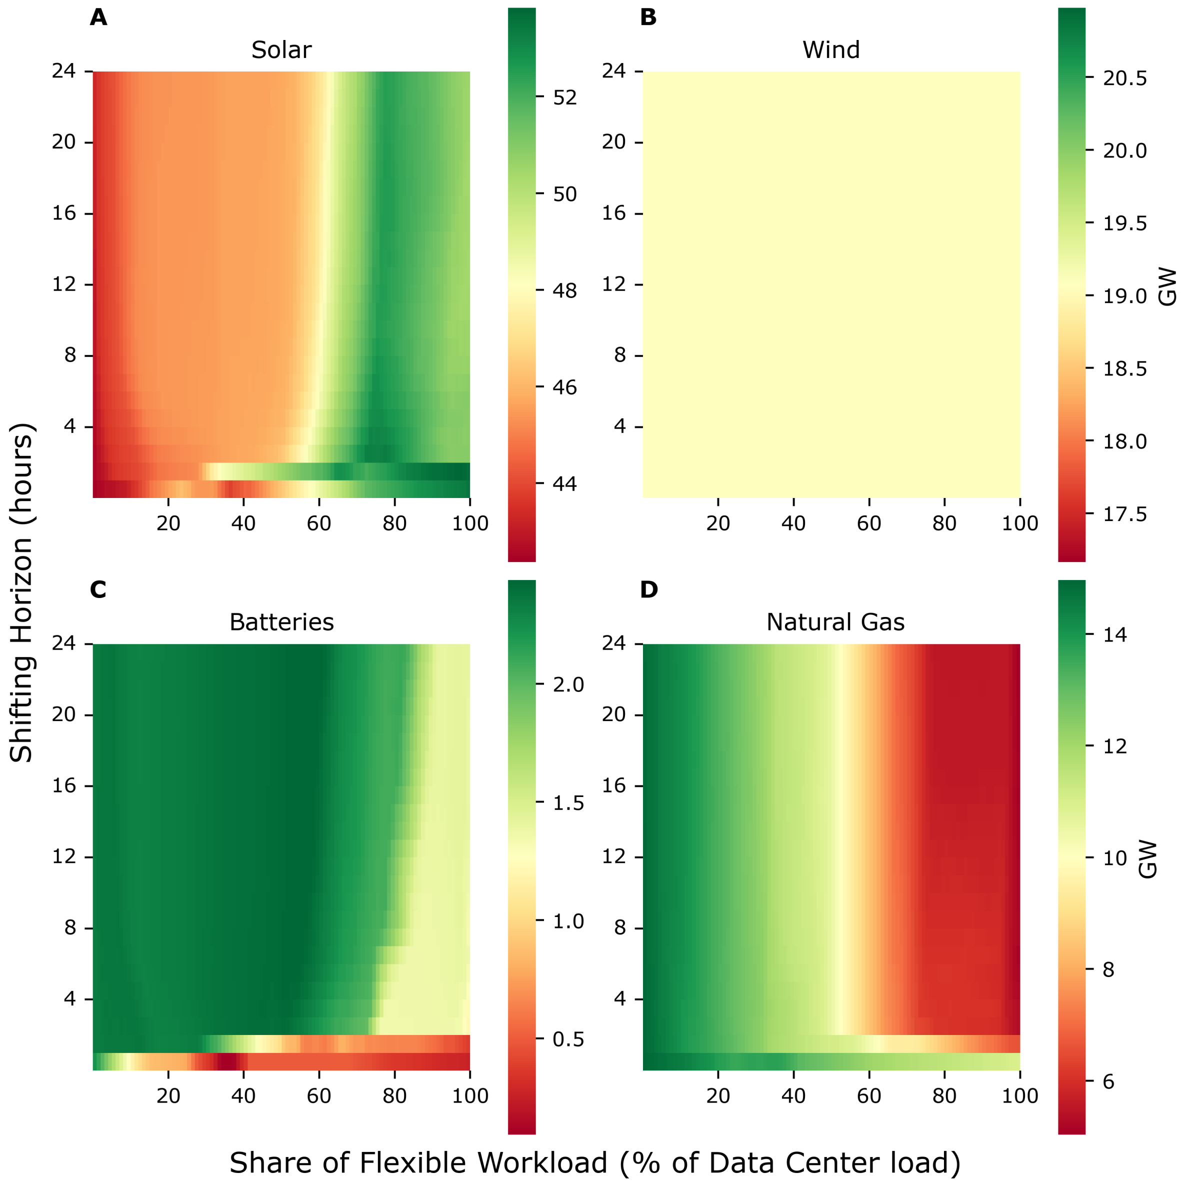

Figure S2: Capacity Investments per Technology in the Mid-Atlantic. Panels show new capacity additions for solar **(A)**, wind **(B)**, batteries **(C)**), and natural gas **(D)** across combinations of data center shifting horizon (1 to 24 hours) and flexible workload share (1% to 100%). Note that color scales vary across subplots.

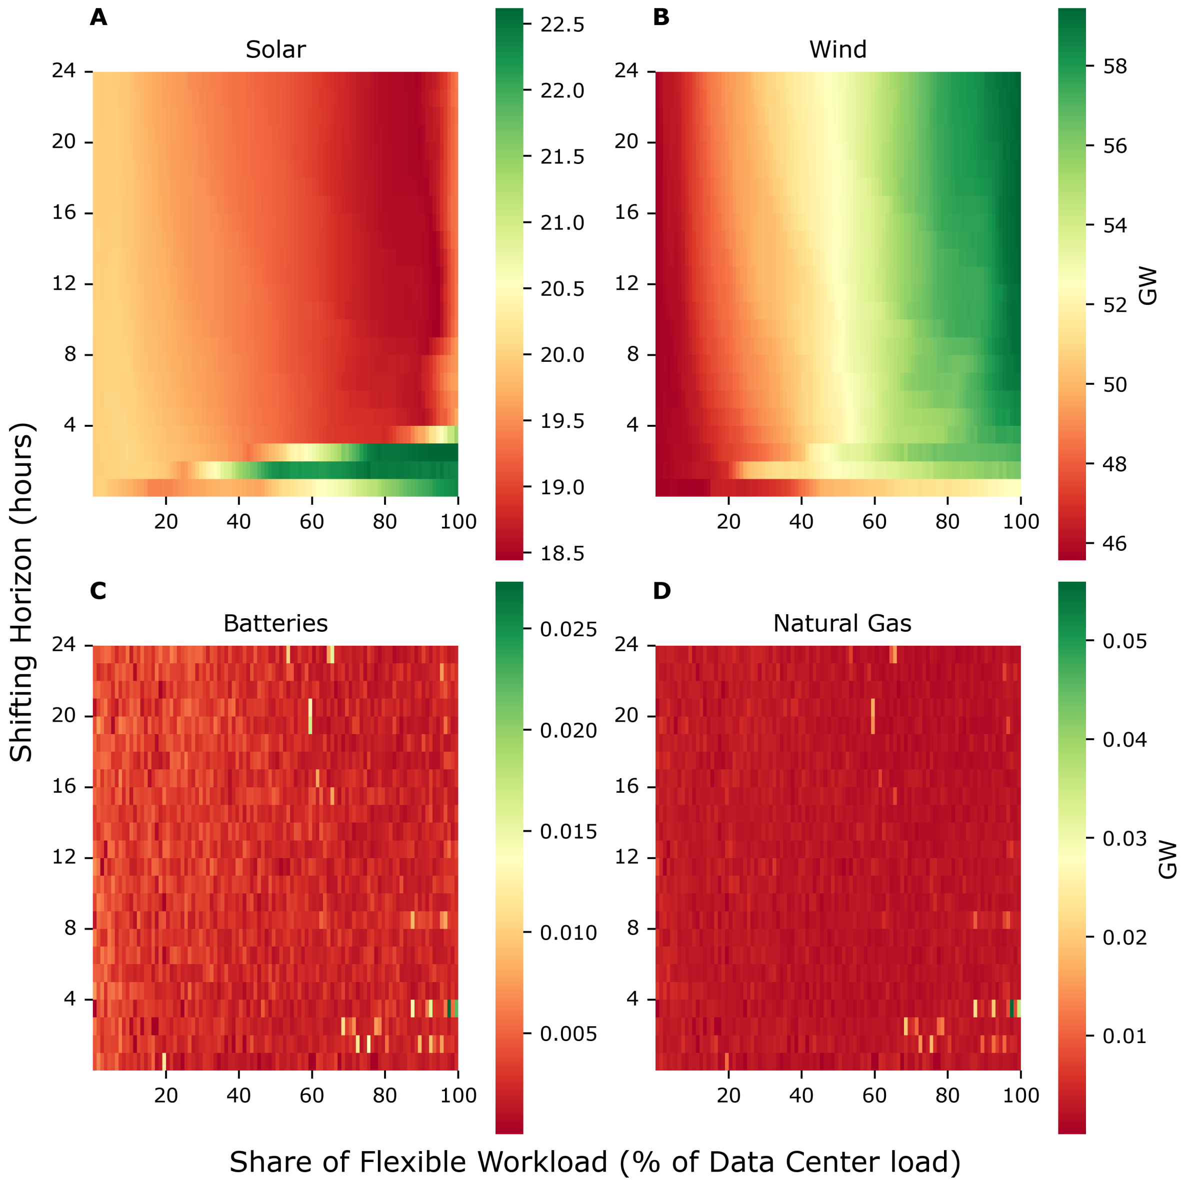

Figure S3: Capacity Investments per Technology in Texas. Panels show new capacity additions for solar **(A)**, wind **(B)**, batteries **(C)**), and natural gas **(D)** across combinations of data center shifting horizon (1 to 24 hours) and flexible workload share (1% to 100%). Note that color scales vary across subplots.

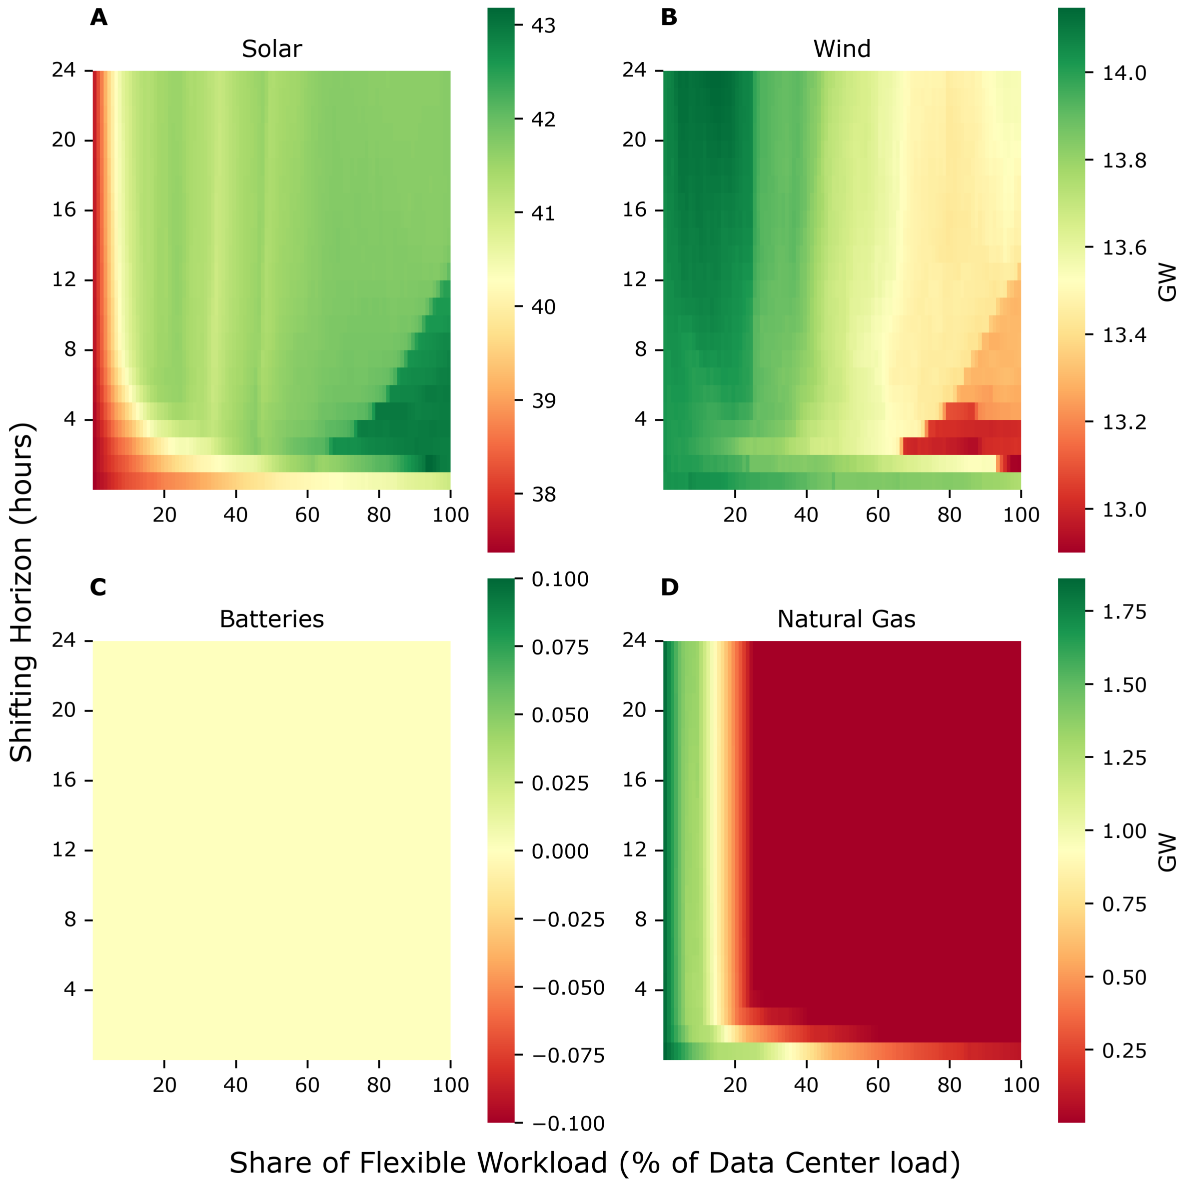

Figure S4: Capacity Investments per Technology in WECC. Panels show new capacity additions for solar **(A)**, wind **(B)**, batteries **(C)**), and natural gas **(D)** across combinations of data center shifting horizon (1 to 24 hours) and flexible workload share (1% to 100%). Note that color scales vary across subplots.

12 WECC Data Center Shifting Operations

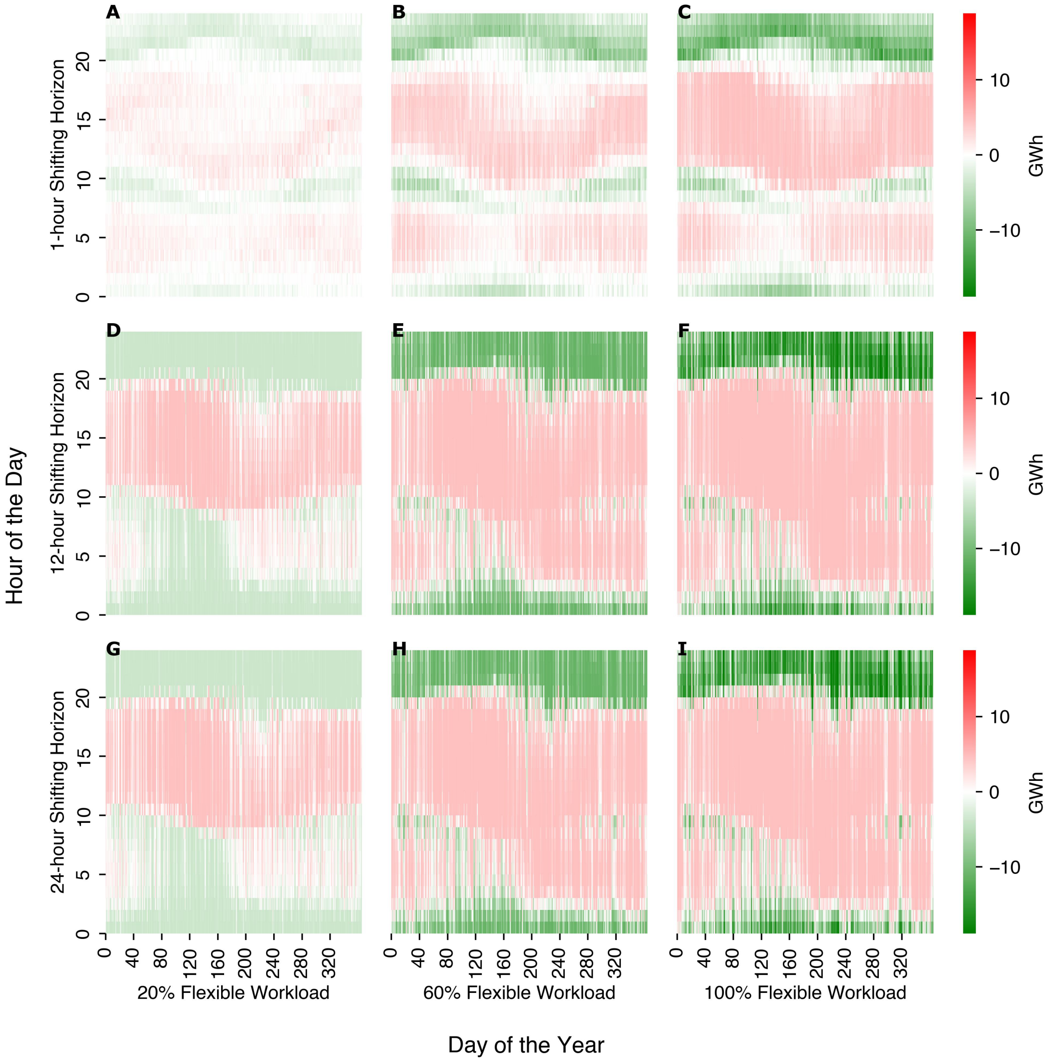

Figure S5: Data Center Load Shifting for WECC. Each panel displays the net hourly workload shifted (in GWh) across an entire year, with the x-axis representing days (1–365) and the y-axis representing hours of the day (0–23). Positive values (red) indicate workload shifted into a given hour; negative values (green) represent workload shifted out. Columns show increasing flexible workload shares—4%, 12%, and 20%—based on 20%, 60%, and 100% of a shiftable portion capped at 20% of total capacity. Rows indicate shifting horizons of 1 hour, 12 hours, and 24 hours, reflecting the maximum time a task can be advanced or delayed. This illustrates how varying flexibility levels and temporal windows influence both intra-day load scheduling and broader seasonal shifting patterns.

13 Additional Capacity Information

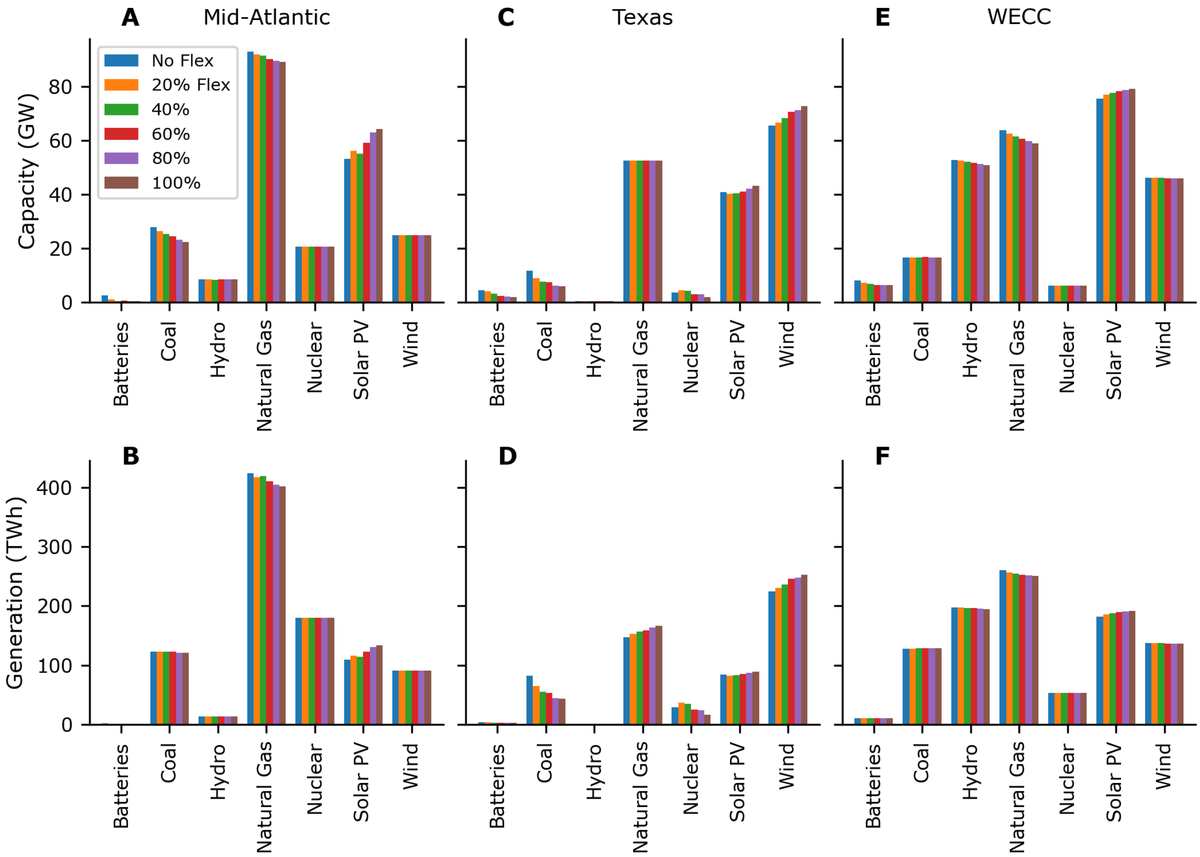

Figure S6: Capacity and Generation per Region with a 1-hour shifting horizon. Top row **(A, C, E)** shows total installed capacity by technology, net of new investments and retirements, for each region. Bottom row **(B, D, F)** presents corresponding total generation by technology. Results are shown for flexible workload shares ranging from 20% to 100% in 20% increments, alongside a baseline scenario without flexibility. All scenarios assume a 1-hour shifting horizon. No new capacity investments can be made in Coal, Nuclear, and Hydro. All technology types can be retired.

14 Cost Differences per Component

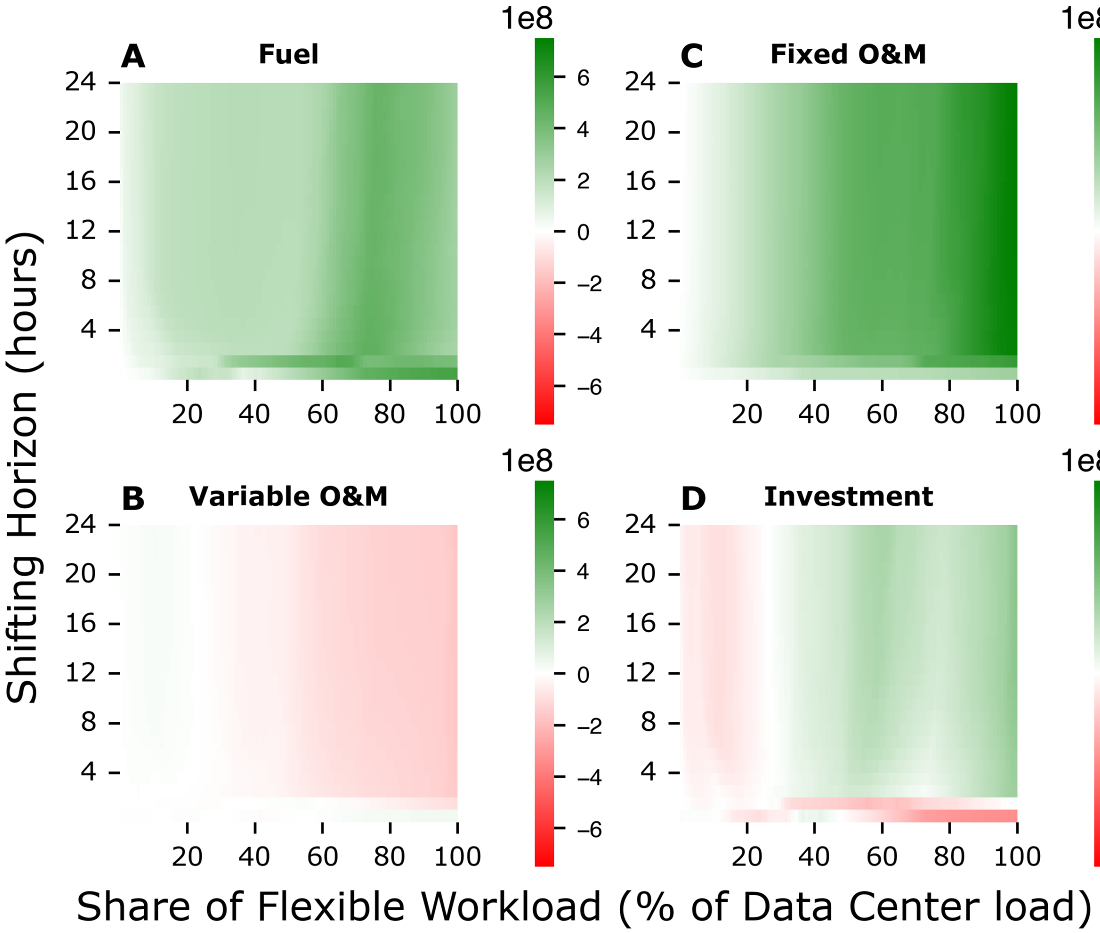

Figure S7: Cost Difference per Component for the Mid-Atlantic. Panels show the change in system costs between scenarios with and without data center flexibility for fuel **(A)**, fixed O&M **(B)**, variable O&M **(C)**, and generation investment **(D)**, across combinations of shifting horizon and flexible workload share. Green indicates a cost reduction with flexibility; red indicates an increase.

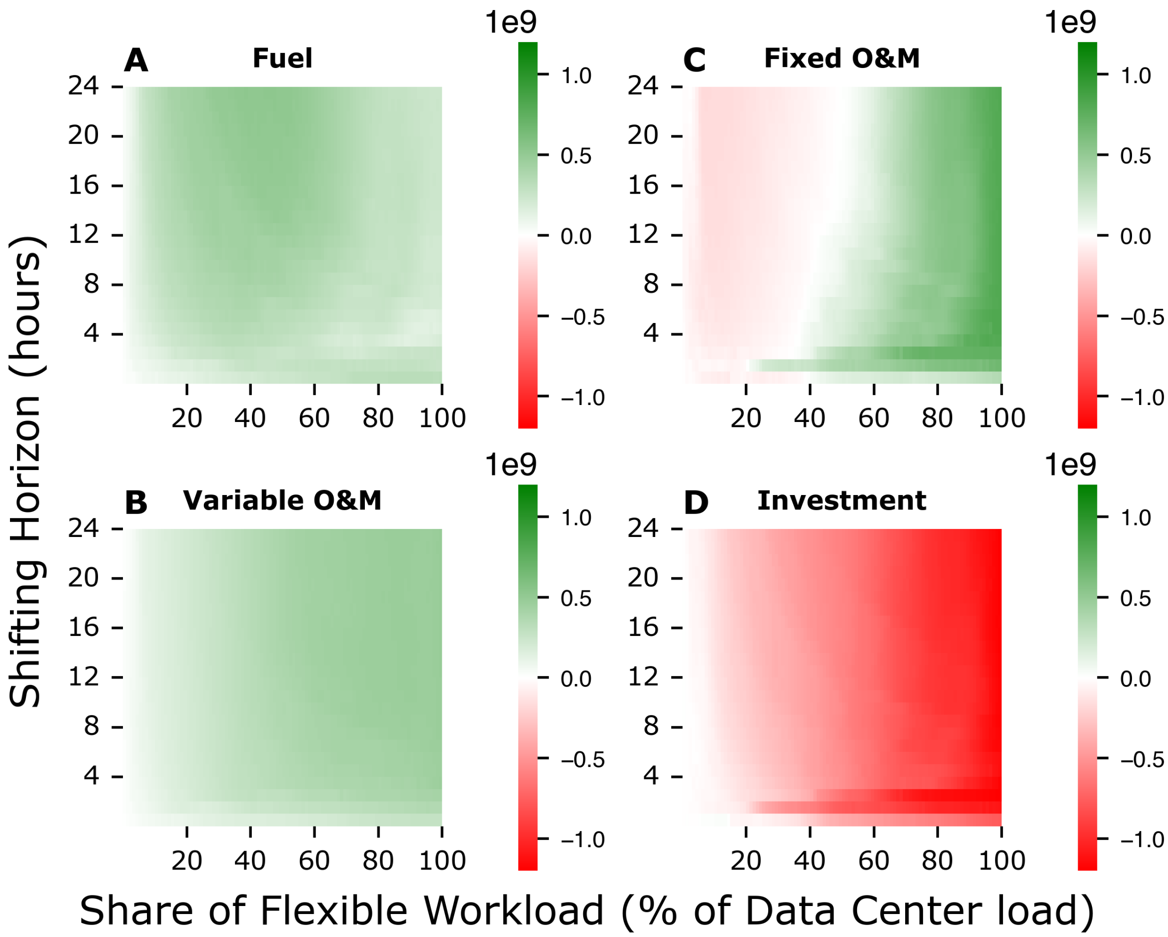

Figure S8: Cost Difference per Component for Texas. Panels show the change in system costs between scenarios with and without data center flexibility for fuel **(A)**, fixed O&M **(B)**, variable O&M **(C)**, and generation investment **(D)**, across combinations of shifting horizon and flexible workload share. Green indicates a cost reduction with flexibility; red indicates an increase.

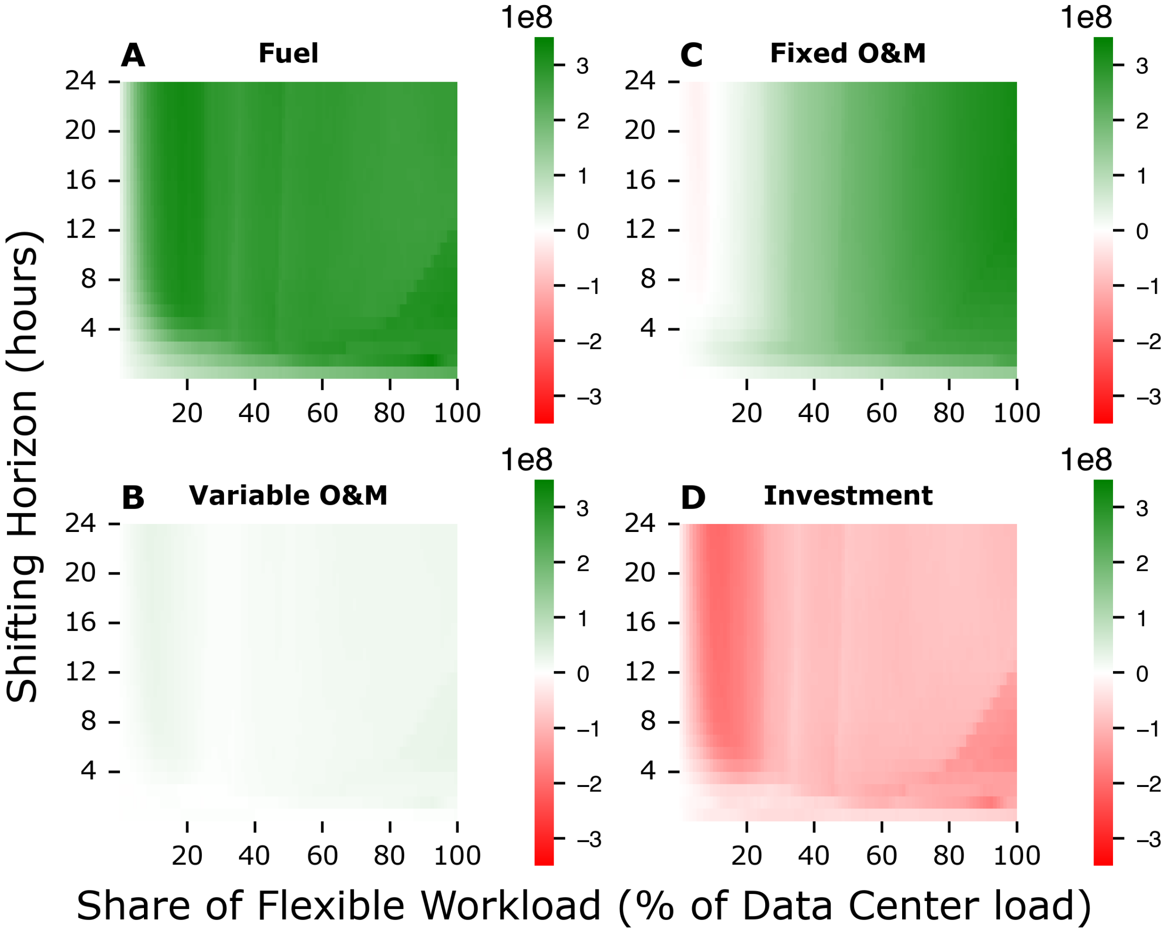

Figure S9: Cost Difference per Component for WECC. Panels show the change in system costs between scenarios with and without data center flexibility for fuel **(A)**, fixed O&M **(B)**, variable O&M **(C)**, and generation investment **(D)**, across combinations of shifting horizon and flexible workload share. Green indicates a cost reduction with flexibility; red indicates an increase.

NBER WORKING PAPER SERIES

MEASURING THE IMPACT OF DATA CENTERS IN THE UNITED STATES ECONOMY: MONETARY DAMAGE FROM AIR POLLUTION AND GREENHOUSE GAS EMISSIONS.

Nicholas Z. Muller

Working Paper 35100
http://www.nber.org/papers/w35100

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2026

The author recognizes generous financial support for this research provided by the Tepper School of Business at Carnegie Mellon University. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Measuring the Impact of Data Centers in the United States Economy: Monetary Damage from Air Pollution and Greenhouse Gas Emissions.
Nicholas Z. Muller
NBER Working Paper No. 35100
April 2026
JEL No. Q41, Q51, Q53, Q54

ABSTRACT

This paper quantifies the environmental externalities associated with electricity consumption by data centers in the United States, focusing on damages from local air pollution and greenhouse gas (GHG) emissions. Using facility-level data for approximately 2,800 operational data centers in 2025, combined with electricity grid characteristics and emissions data, the analysis estimates pollution impacts through the AP4 integrated assessment model and applies the social cost of carbon for GHG valuation. Results indicate that data centers consume roughly 250 TWh of electricity—about 5–6% of U.S. generation—and generate approximately $25 billion in gross external damages (GED), with a range of $10–$33 billion. These damages are highly geographically concentrated, with Texas and Virginia accounting for 30%of the national total. While GED comprises about 5% of industry GDP, this ratio varies widely across states, exceeding GDP in some regions. Planned data center expansion could increase electricity demand and associated damages by up to 85% in the near term. Despite these environmental costs, preliminary comparisons suggest that the damages attributable to AI-related energy use are small relative to potential productivity gains.

Nicholas Z. Muller
Carnegie Mellon University
Tepper School of Business
Department of Engineering and Public Policy
and NBER
nicholas.muller74@gmail.com

I. Introduction

The rapid growth of data centers in the United States (U.S.) is raising concerns because of the potentially wide-ranging impacts of this growth on the energy system and on environmental conditions in local communities (Blunt and Hiller, 2026; Hudson, 2026; Rust and Parker, 2026). At a systems level, data centers are electricity-intensive, and the surge in demand for power is straining regional electricity grids, complicating long-term planning and potentially increasing electricity prices. Quick expansion of a sector with large loads means that in some regions, utilities must accelerate investments in generation, transmission, and distribution infrastructure (Brown, 2026). In other cases, older generation assets that had been retired are being brought back online (Plumer, 2024). The disruptive change associated with data centers may have significant implications for ratepayers and it may create regulatory challenges (Wall Street Journal, 2026). At the local level, development of large data centers generates tensions over land use, water consumption for cooling, noise and air quality (Hudson, 2026). This research provides an up-to-date assessment of electricity use by data centers and it quantifies the social costs from pollution emitted during electricity production to meet this load. Thus, the paper contributes one critical dimension to the analysis of the net impact of data centers in the U.S.

This research focuses on measuring the social costs from criteria air pollutants and greenhouse gases (GHGs) caused by electricity production used to power data center operations. Together, these pollutants cause damages that amount to an appreciable share of economic output (Mohan et al., 2023). These calculations rely on several data sources. The location and power usage of existing and planned data centers as of 2025 is provided by S&P Global Insights (S&P Global, 2025). Critical to determining the pollution impacts of data centers and electricity used by data centers, is the nature of power plants that produce the electricity that supports operations. For example, social costs from emissions differ markedly depending on whether power plants burn coal, natural gas or rely on renewable or nuclear technology. The locations, technologies, and fuel used to generate electricity that powers these facilities is provided in publicly available data issued by the United States

Department of Energy (USDOE) and the United States Environmental Protection Agency (USEPA), (USDOE, 2026; CAMPD, 2026; eGRID, 2026).

With data characterizing the electricity grids and associated emissions, the next step in the empirical estimation of social cost is to connect emissions to impacts. For local air pollution this study relies on the AP4 model (Tschofen et a., 2019; Dennin et al., 2025). This integrated assessment model connects emissions of several economically significant local air pollutants to public health consequences, and social costs. Emissions of GHGs are valued using the social cost of carbon (SCC). Monetary damages from air pollution and GHGs are combined to calculate Gross External Damages (GED), a monetary measure of impacts from environmental pollutants used in the prior literature (Muller, Mendelsohn, and Nordhaus, 2011). This analysis then compares the GED to measures of the economic benefit that data centers provide. These include Gross Domestic Product (GDP) and personal income from the data processing, hosting, and related services industry.

There has been considerable research on energy consumption by data centers[1]. Most related to this analysis is an unpublished working paper by Han et al., (2024). This research connects data center energy use to emissions and health consequences in a similar fashion to the present paper. There are several key differences. First, Han et al., (2024) use energy consumption data from 2023 which, though only two years old, is outdated given rapid industry growth. In contrast, this analysis uses energy consumption information from data centers in 2025. Second, Han et al., (2024) use state-level energy consumption data whereas the present paper uses Quarter 3 (Q3) 2025 data for the fleet of roughly 2,800 operational data centers at the facility level. Third, Han et al., (2024) track electricity load from operational data centers in 2023. The present analysis also encompasses expected emissions and social costs from planned facilities. Fourth, Han et al., (2024) emphasize their 2030 forecast impacts. The focus of this paper is to quantify impacts in 2025 and to compare the GED to GDP to gauge the magnitude of impact.

[1] EPRI, 2024; Koomey, 2008; Mytton and Ashtine, (2022)

The central results of this analysis are summarized as follows. First, damages from electricity consumption at data centers in 2025 amount to $25 billion ($10 to $33 billion)[2]. The GED vary considerably across space. In one sense this should be obvious, a priori, as data center development exhibits clustering which reflects regional demand, electricity prices, and regulatory contexts. However, the extent of this variation and the underlying factors causing it are noteworthy. The GED from data centers in Virginia and Texas amount to 30% of the national total. Commercial electricity prices in these states were well below the national average in 2025. Further, data centers in just 10 states produce 70% of the total GED. Second, to make sense of the magnitude of these cost estimates, the GED are compared to GDP and personal income from the information technology data processing, hosting, and related services industry. For all data centers, the GED comprise just 5% of GDP. In North Dakota and Wyoming, the GED from electricity generation to power data centers are more than 5 times and 2 times larger than GDP from the data processing, hosting, and related services industry, respectively. Third, the expected damages associated with planned data centers suggest rapid near-term increases in power consumption and social costs. Estimated GED from planned facilities totals $21 billion which is roughly equivalent to GED from operational facilities in Q3 2025. Particularly large increases in load and social costs are projected to occur in Virginia, Ohio, and Texas. Finally, in a series of provisional calculations, the magnitude of GED from AI-related activities at data centers appears to be small relative to estimates of the associated productivity benefits. For example, if one assumes that AI exhibits a 0.5% boost to GDP, the associated GED comprise just 2.4% of the productivity gains.

The remainder of this report proceeds as follows. Section II covers data and methods. Section III reports results, and Section IV concludes.

[2] All monetary values are reported in $2025.

II. Methods and Data

This section begins with the methods and data used to calculate electricity consumption by data centers before covering the model used to estimate the impact of local air pollution. The section concludes with the approach to estimating damages from GHGs.

S&P Global Insights provides both the location and detailed physical attributes of every data center in the U.S. (S&P Global, 2025). Among these, there are approximately 2,800 operational facilities. For each data center, net utilized power is reported in kilowatts (kW). We convert this to megawatt hours (MWh) of electricity consumed as shown in equation 1. This calculation multiplies net utilized power, power usage efficiency (*PUE*), the number of hours in a year, and a load factor, which governs the fraction of total facility capacity in use that draws power from the grid. *PUE* is the standard metric for measuring data center energy efficiency. This metric is unique to each facility, and it is reported by S&P (S&P Global, 2025) It is calculated as the ratio of total facility power to IT equipment power. Leveraging guidance from S&P Global, the load factor ranges from 50 to 60% and is set to 55% in the default scenario. The load factor is varied in a sensitivity analysis.

$$Mwh_{f,t} = \left(\frac{kW_{f,t}}{1{,}000}\right) \times PUE_{f,t} \times 8{,}760 \times LF_{f,t} \tag{1}$$

The next step in the empirical analysis connects facility-level electricity load and emissions. The requires information on locating facility in an electricity grid region. The paper adopts eGRID subregions which were developed by USEPA to estimate the mix of power plants supplying electricity to each region in a given year (eGRID, 2026). The EIA form 923 (USDOE, 2026) reports generation data and the USEPA Clean Air Markets Program Data (CAMPD, 2026) reports emissions. Both data sets are resolved at the power plant level for the year 2025. CAMPD includes plant level information on emissions of GHGs (carbon dioxide - CO_2) and local air pollution including sulfur dioxide (SO_2) and nitrogen oxides (NO_x). Emissions of primary fine particulate matter ($PM_{2.5}$), volatile organic compounds, and ammonia (NH_3) are provided by eGRID for 2021. Newer data for these species are not available.

The analysis calculates uniform emission rates within each sub-region. Data center load is multiplied times these subregion-specific emission rates to estimate emission levels for each data center (f) as shown in (2). The denominator in (2) includes all generation in a given region, inclusive of emitting and non-emitting generators.

$$E_{f,s,t} = MWh_{f,t} \times \left[\left(\frac{\sum_{p=1}^{P} E_{p,s,t,g}}{\sum_{p=1}^{P} MWh_{p,s,t,g}}\right)\right] \tag{2}$$

where: $E_{f,s,t}$ = emissions of pollution species (s), from data center (f), during year (t).

This approach implicitly assumes that data centers comprise average load which is met predominantly by baseload generation assets.

This study uses the AP4 integrated assessment model to connect emissions of local and regional air pollution from electric power generation to their consequences on society (Dennin et al., 2025). The AP4 model[3] facilitates detailed facility level estimates of pollution cost, through air quality modeling, epidemiology, and economic valuation modules. AP4 is the latest version of the APEEP family of models (Muller and Mendelsohn, 2007; 2009; Muller, Mendelsohn, and Nordhaus, 2011). Earlier versions of the AP4 model have been subject to rigorous comparison with other similar models[4]. AP4's predictions have been analyzed relative to the network of monitors used by the USEPA to enforce air quality standards (Dennin et al., 2025).

AP4 begins with the National Emissions Inventory (NEI), which is an inventory of local air pollution emissions for all reported emissions in the contiguous United States (USEPA, 2026). Using this inventory, AP4 estimates baseline concentrations of $PM_{2.5}$ accounting for all emissions in the US economy. Importantly, AP4 models the connection between emissions of precursor species and ambient $PM_{2.5}$. These precursor species include SO_2, NO_x, NH_3, volatile organic compounds, and primary $PM_{2.5}$. AP4 models these relationships through a reduced-complexity chemistry module that accounts for interactions among

[3] Publicly available here: https://nickmuller.tepper.cmu.edu/APModel.aspx
[4] Gilmore et al., (2018).

inorganic species. Each of these precursor species contributes to the concentrations of total $PM_{2.5}$ which AP4 predicts at the county levels.

Using these estimated concentrations, the model then uses detailed population estimates to calculate exposures for every county in the contiguous US. Populations are disaggregated by age groups because prior epidemiological research has shown that susceptibility to pollution exposure is a function of baseline health status (Krewski et al., 2009). The model contains spatially detailed information on baseline health status including mortality risks by county and age group. These data are provided by the Centers for Disease Control and Prevention (CDC, 2025).

Then, AP4 uses peer-reviewed concentration response functions to link exposure to $PM_{2.5}$ to elevated premature mortality risk (Krewski et al., 2009; Lepeule et al., 2012). AP4 focuses on premature mortalities for two reasons. First, prior research has shown that premature mortality risk comprises the vast bulk of damages from local air pollution (Muller, Mendelsohn, Nordhaus, 2011; USEPA, 2011). Second, if the model included mortality and morbidity health states in the damage estimates, one might be concerned about double counting if, say, short run health effects like emergency room visits from asthma are followed eventually by a premature death.

The model calculates the deaths attributable to $PM_{2.5}$ exposure ($M_{a,j,t}$). This calculation is shown in equation (3).

$$M_{a,j,t} = Pop_{a,j,t}\gamma_{a,j,t}\left(1 - \frac{1}{exp(\beta \Delta PM_{j,t})}\right) \tag{3}$$

where: $Pop_{a,j,t}$ = population of age-group (a), in county (j), year (t).
$\gamma_{a,j,t}$ = baseline mortality risk of age-group (a), in county (j), year (t).
β = statistically estimated parameter from Krewski et al., (2009).
$\Delta PM_{j,t}$ = change in ambient $PM_{2.5}$ in county (j), year (t).

Once the mortality risk from air pollution is estimated using equation (3), AP4 attributes a monetary value to these risks. In both academic literature and in policy analyses, the standard way to monetize mortality risk uses the value of a statistical life (VSL). The VSL is the marginal rate of substitution between income and mortality risk. It is a measure of the willingness to pay to reduce mortality risk. There are two ways in which the VSL is estimated empirically. In revealed preference studies, the VSL is estimated from the relationship between wages and on the job mortality risk. In stated preference studies, the VSL is estimated based on individual responses to highly structured surveys. This analysis uses what had been the USEPA's preferred VSL of $10 million (USEPA, 2025), which has been adjusted to $2025. In this study, as is standard practice, the VSL is applied uniformly to persons of all ages and income levels. USEPA's preferred VSL is a combination of both revealed and stated preference studies.

AP4 is applied in two ways in this paper. First, marginal damage, or $/ton damages, for each of the local air pollutants are calculated. These facility level damages are averaged by state to match the eGRID subregions. The $/ton damages are multiplied times the estimated tonnage of emissions produced from electricity generation to power data centers. Second, AP4 is run with and without data center load, and the associated emissions. This strategy is used to produce figures 3, 4, A2, and A3.

In order to model the social cost from greenhouse gas emissions, this study uses the SCC. The SCC represents the present discounted value of damages from carbon dioxide equivalents on a per ton basis. While there is deep structural uncertainty associated with the SCC, numerous peer reviewed estimates have been produced by environmental economists and policy analysts. Two different SCC values are used in this study. The default value is approximately \$190/ton CO_2, and it was reported by the Biden Administration's Interagency Working Group on the Social Cost of Carbon (USEPA, 2023). The second value is an earlier estimate also from the Biden Administration resulting from a meta-analysis of multiple models. The second value is approximately \$50/ton CO_2 (White House, 2021). These SCC estimates are multiplied times the CO_2 emitted by the power plants producing electricity to power data centers.

The data provided by S&P Global (2025) include information on both operational and planned data centers. The central results presented here focus on operational data centers in the quarter 3 of 2025. A second set of calculations focus on the social costs from planned facilities. This is intended to provide a sense of growth in both power consumption and social costs from power consumption in the near term as these new facilities come online.

Throughout this report we refer to the damages, or social costs, associated with local air pollution and greenhouse gases as the gross external damages, or GED. The GED was introduced by Muller, Mendelsohn, and Nordhaus (2011). This term reflects the fact that damages from air pollution and greenhouse gases comprise an externality. In the context of datacenter power consumption, the external costs from power generation are borne by consumers exposed to $PM_{2.5}$. Damages from greenhouse gases manifest many years following emission, and hence, reflect an externality borne by future generations. Further, the social costs from pollution emissions are *gross* (not net) since they do not account for payments made to reduce pollution removal or abatement, nor do the costs account for payments made to purchase tradable allowances in cap-and-trade programs.

In addition to the data discussed above, the analysis makes use of other publicly available datasets. These specifically include data from the state and regional national income and product accounts (NIPAs) reported by the US Bureau of Economic Analysis (USBEA, 2026a) In particular, the data used are gross domestic product, or GDP, by industry, with a particular focus on information technology: data processing, hosting, and related services. These data are matched by state to the location of data centers. The GED are compared to the GDP by state, aggregating over facilities within each state. An analogous comparison is made to personal income for information technology: data processing hosting and related services (USBEA, 2026b). It is important to note the limitations associated with making these comparisons. First, GDP is an imperfect measure of the value of services provided by these facilities. By definition, GDP does not capture spillover effects or positive externalities associated with data hosting facilities. In addition, personal income statistics do not measure multiplier effects associated with households and consumers in each state having incrementally more income, and the benefits associated with additional expenditures on state and regional economies. Nonetheless these comparisons gauge the magnitude of the damages relative to conventionally used measures of the benefits of these facilities for state and regional economies.

III. Results

The bottom row in Table 1 reports the total power consumption and damages across all datacenters in the contiguous U.S. in 2025 (S&P Global, 2025). The datacenters consumed approximately 250 terawatt hours (TWh) of electricity. This was between 5 and 6% of total electricity produced in the US[5]. The associated GED is estimated to be $24.6 billion. A sensitivity analysis reported in table A2 suggests a range of the GED from $10 billion to $33 billion. Specifically, when using a higher run rate at the data centers (75% rather than the default of 55%) the GED increase to $33.6 billion. Employing the lower SCC of $50/ton CO_2, the GED decrease to $9.9 billion. And employing the higher PM-mortality dose-response function increases the GED to $30.7 billion.

Four findings emerge from table 1. First, data center activity is highly concentrated. Facilities in just 10 states contribute 70% of total GED. Within these 10 states, facilities in Texas and Virginia cause 30% of the GED. Second, the GED from GHGs comprise 80% of the total, though this share is sensitive to modeling assumptions. Third, the GED/MWh are large relative to commercial electricity prices. In 6 of the 10 states in table 1, the GED/MWh *exceeds* the state average commercial electricity price. Fourth, electricity prices in all 10 states in table 1 are less than the national average commercial price in 2025 which was $0.13 per kWh[6]. In Texas and Virginia, prices are 30% lower than the national average.

Air Quality Impacts. Figure 2 displays the total power consumption by state. Texas and Virginia stand out with consumption levels over 25 TWh. Several other states in the Midwest and on the West coast exhibit large electricity loads from data centers. Figure 3 displays the change in ambient $PM_{2.5}$ from data center power consumption. Exposures to these changes comprise the GED from air pollution. Figure 3 indicates areas of elevated $PM_{2.5}$ occur in Texas, Virginia and the Carolinas, Pennsylvania and Ohio, Wyoming, Nebraska, Iowa, and

[5] This estimate is roughly in line with forecasts made by EPRI (2024).
[6] https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=table_5_03

North Dakota. Intuitively, the regions that incur elevated $PM_{2.5}$ have (or are downwind from areas with) large data center load as reported in table 1. The regionally widespread changes in $PM_{2.5}$ manifest because power plants that meet data center load may not be proximal to data centers, and second, because air pollutants disperse often long distances once they are emitted. Further, the electricity grid is more pollution-intensive east of the Rocky Mountain states.

Figure 4 shows where the GED from air pollution exposures occur. Recall, as shown in equation (3), the air pollution GED are calculated in each county using the estimated incremental concentrations of $PM_{2.5}$ (shown in figure 3) together with county level population and mortality rate data. Figure 4 makes clear the importance of population. That is, the counties exhibiting the highest damage are in large cities. For example, Houston, Dallas, and San Antonio in Texas, each incur an estimated GED from local air pollution above $10 million. Large cities in the Midwest also exhibit elevated $PM_{2.5}$ damages, including Chicago, Pittsburgh, Detroit, Columbus, and Cleveland, Ohio. In addition, even very small changes in ambient $PM_{2.5}$ that manifest in populous metropolitan areas result in significant damages. This is evident in Los Angeles, Phoenix, Las Vegas, and Boston.

Total GED. Figure 5 displays the total GED by state, inclusive of GHGs and local air pollution emissions. In this figure, the GED are attributed to the state in which data centers operate, not where the GED occur. Table 1 shows the top 10 states ranked by the GED caused by power consumption from data centers in each state. Figure 5 and table 1 clearly indicate that data center activity in Texas and Virginia cause the largest GED. Even among the 10 states shown in table 1, Virginia and Texas are outliers. For example, the GED in Texas and Virginia are about 2 times larger than the GED from data centers in Ohio, the state causing the third highest GED. Commercial electricity prices in Virginia and Texas are under $0.10/KWh, which is well below the national average of $0.13/KWh. Because of the importance of electricity costs to the overall cost structure faced by data centers, these low prices are likely an important determinant of facility location.

GED, GDP, and Personal Income. Table 2 relates the GED to state-level GDP contributed by information technology (IT) facilities that provide data processing, hosting, and related services. The states in this table are ranked in descending fashion according to the share the GED comprises of GDP from this industry by state. Importantly, the most recent GDP data available is for 2024. Nationally, GED from all data centers amounts to $25 billion, while GDP for this industry was $544 billion in 2024. This relatively small share masks significant heterogeneity.

Data centers in North Dakota generate $665 million in GED and just $121 million in GDP from IT data hosting services. GED was over 5 times larger than the value of economic production from IT data hosting services in this state. North Dakota is clearly an outlier. The next largest share of GED relative to GDP is in Wyoming. In this state GED from electricity load for data centers amounts to $244 million which is about 2 times larger than GDP contributed by these facilities. Data centers in Iowa and Oklahoma cause GED that is between 50% and 80% of state GDP. The remaining 6 states in table 2 contain data centers that produce GED under 40% of GDP. Virginia, as noted above, contains data centers that consume the most electricity load and that contribute the largest share of GED. The GED from data centers in Virginia comprises nearly 18% of total GED. However, the state GDP from data centers in Virginia amounts to just under 3% of national GDP from this industry. In contrast, data centers in California contribute 37% of national GDP and just 2% of the GED.

It is important to consider how to Interpret these comparisons of GED to GDP. Specifically, the GDP tracked here is that reported for data processing, hosting, and related services in 2024. By definition, GDP attributed to this industry does not measure positive externalities associated with data hosting services such as productivity effects that spill over to other areas in the economy. Nor does GDP attribute local stimulus (job creation or expenditures by employees) to GDP for this industry. Thus, the comparison between GED and GDP is intended as an illustrative scaling exercise. Further, as mentioned above, GDP reported here is for 2024. Considerable growth in this industry between 2024 and 2025 suggested GDP in 2025 may have been considerably higher than in 2024.

It is also worth considering the geographic scope of the GED and GDP measures. The GED is associated with electricity consumption by data centers located in each state. For air pollution, we know that the consequences of exposures tend to extend well beyond state boundaries. As such, the damages encompass regional or national exposures, and these are attributed back to the location of electricity generation. For greenhouse gases, of course, the GED is global. Thus, the fraction of GED from greenhouse gases accruing in the state where the data centers are located is likely quite small. In contrast, GDP is reported for each state inclusive of economic flows that manifest within each state.

In pursuit of another way to characterize the magnitude of the GED produced from electricity generation to power data centers, Table 3 reports the GED as a fraction of personal income from data processing, hosting, and related services. Table 3 ranks the states in descending order by the fraction that GED comprises of personal income. Like Table 2, the states at the top of Table 3 include North Dakota, Wyoming, and Iowa. The use of personal income to gauge the magnitude of GED is motivated by the fact that personal income from data processing, hosting, and related services is a measure of local economic stimulus from these economic activities. Much like Table 2, Table 3 shows that North Dakota is an outlier in terms of the magnitude of GED relative to personal income. Specifically, GED is 15 times larger than personal income from data hosting facilities. Using this measure, 5 of the remaining 9 states in table 3 exhibit GED in excess of personal income. Together, Tables 2 and 3 show, albeit imperfectly, that the GED associated with power generation to meet data center load amounts to an appreciable share of the measurable economic benefit of these facilities.

Planned Facilities. The next set of results examines damages from planned facilities (as of Q3 2025). Table 4 shows the 10 states with the largest GED from planned data centers. One should view these estimates with caution since power consumption is, by definition, estimated and forward-looking. Total planned capacity is estimated to consume 210 TWh with an associated GED of $21 billion. These estimates, coupled with operational facilities, suggest increases in electricity load and the GED of up to 85% in the near term. Virginia exhibits the largest expected GED from planned data centers of just under $7 billion. This is

60% more than the GED from operational data centers in Virginia and it is more than one-quarter of total damages from all operational data centers in 2025. In Ohio, the GED from planned data centers is estimated to be just over $2 billion. This reflects roughly a 40% increase over the GED from existing data centers operating in 2025. Table 4 suggests that GED from data center activity in Illinois, Pennsylvania, and Arizona may double. Substantial increases are also estimated to occur in Texas, Missouri, Georgia, Iowa, and North Dakota. These comparisons, while provisional, are indicative of the rate of growth in data center capacity, electricity use, and in the social costs from this high growth industry. Figures A2 and A3 in the appendix demonstrate the expected change in ambient $PM_{2.5}$ and the air pollution GED from planned facilities. The areas exhibiting the highest ambient $PM_{2.5}$ occur along the East coast from the Carolinas up through Pennsylvania. Much like the results for operational data centers, the air pollution GED from planned facilities manifest in large cities, with large, exposed populations.

Data Center Types. Table A1 in the appendix decomposes the electricity use and GED by the type of data centers in the S&P Global Insights data (S&P Global, 2025). There is a total of 2,816 operational facilities in the data. Of these, 1,509 are retail and 525 are wholesale facilities. Together, these data centers contribute about $7 billion (29%) of the GED and 34% of electricity consumption. There are 432 hyperscale facilities. These (large) facilities contribute 41% of the GED and consume 40% of the electricity load. Facilities devoted to cryptocurrency mining are about as large as hyperscale facilities, but the power consumption per square foot is 80% larger. As a result, even though there are only 121 crypto mining centers, they produce one quarter of the GED while consuming 20% of the total power for all facilities.

AI: The core value of AI lies in up leveling productivity, through model building, task deployment, and data acquisition. At this early stage of global AI deployment, precisely estimating the impact of AI on economic performance is challenging and largely beyond the scope of the present research. Further, a precise understanding of the share of GED reported in this paper that is attributable to AI training and inference activities remains out of reach. However, in order to obtain a preliminary estimate of the GED from AI activities and to gauge the magnitude of this estimate, the following provisional calculations are executed. First, IEA (2025) reports that roughly 15% of all energy use at data centers is due to AI activities. Using the results in Table 1, this suggests that 38 TWh and $3.7 billion in GED are due to AI activities. Recent estimates of the impact of AI on GDP are generally under 10% by 2030 (Goldman Sachs, 2023). Figure 6 combines these data points to demonstrate the comparison between the GED and GDP effects from AI. In sum, the GED effects appear very small relative to potential increased output. For example, if AI boosts GDP by 1%, the GED amount to 1% of the value of enhanced output. If AI boosts GDP by 0.5%, the GED constitute just 2.4% of increased GDP. Finally, if AI adds 0.1% to GDP, a very small increase in output, the GED would in this case comprise about 12% of the additional output from AI.

IV. Conclusions

This paper estimates the damages from local air pollution and GHG emissions from electricity powering data centers in the U.S. in 2025. The GED, which total $24 billion, comprise just 5% of the GDP contributed by data storage and data hosting facilities as measured using the national income and product accounts. However, the GED are highly concentrated; 30% of the GED stem from facilities in Virginia and Texas. Lower than average commercial electricity prices in these states may play a central role in facility location decisions. Planned facilities may increase electricity load and associated GED by up to 85% in the near term. This will likely have substantial impacts on regional electricity grids and public health. Finally, unless the effect of AI activities on GDP are exceedingly small, the GED from AI-based energy use within data centers appears to be very small relative to the possible benefits of this technology.

Tables and Figures

Table 1: Power Consumption and Damages from Operational Data Centers in 2025.

Data Center State	Elec. Load[A] (Gwh)	Total GED[B]	GED GHG	GED LAP	GED ($/Mwh)	Electricity Price[C] ($/Mwh)
VA	46,400	4,290	3,420	869	92	95
TX	36,200	3,230	2,760	468	89	86
OH	13,500	1,740	1,330	419	129	116
IA	9,459	1,440	1,050	391	152	110
GA	13,700	1,420	1,240	182	104	115
IL	10,700	1,410	1,050	363	132	131
OR	18,200	1,210	1,060	150	66	106
AZ	13,600	1,090	987	101	80	124
TN	4,855	717	546	172	148	129
NE	5,098	712	568	145	140	88
All States	251,000	24,600	20,000	4,590	101	134

A = electricity load from S&P Global.

B = GED in ($millions), author's calculations.

C = state average commercial electricity price.

Table 2: Power Consumption and Damages from Operational Data Centers and State GDP.

State	Elec. Load[A] (Gwh)	GED ($ million)	GDP ($ million)	GED/ GDP
ND	6,511	665	121	5.51
WY	2,167	244	114	2.15
IA	9,459	1,440	1,904	0.76
OK	3,199	456	805	0.57
NE	5,098	712	1,876	0.38
AL	3,297	493	1,329	0.37
KY	3,045	466	1,277	0.37
NM	2,684	208	641	0.32
OH	13,500	1,740	5,619	0.31
VA	46,400	4,290	13,952	0.31
All States	251,000	24,600	544,148	0.05

A = electricity load from S&P Global.
GED = authors calculations.
GDP = GDP for Data processing, hosting, and related services. These are industries in the Data Processing, Hosting, and Related Services subsector group establishments that provide the infrastructure for hosting and/or data processing services. (USBEA, 2026)

Table 3: Power Consumption and Damages from Operational Data Centers and Personal Income.

State	Elec. Load[A] (Gwh)	GED ($ million)	Personal Income ($ million)	GED/ Personal Income
ND	6,511	665	44	15.18
WY	2,167	244	45	5.45
IA	9,459	1,440	581	2.48
NE	5,098	712	350	2.03
OK	3,199	456	230	1.99
NM	2,684	208	179	1.16
AL	3,297	493	524	0.94
OH	13,500	1,740	1,854	0.94
KY	3,045	466	519	0.90
VA	46,400	4,290	4,841	0.89
All States	251,000	24,600	116,390	0.21

A = electricity load from S&P Global.
GED = authors calculations.
Personal income = Private nonfarm earnings, for data processing, hosting, and related services. These are industries in the Data Processing, Hosting, and Related Services subsector group establishments that provide the infrastructure for hosting and/or data processing services. (USBEA, 2026)

Table 4: Power Consumption and Damages from Planned Data Centers.

State	Elec. Load Planned (Gwh)	GED Planned ($ million)	GED Planned/ GED 2025
VA	75,000	6,940	1.62
OH	19,200	2,490	1.43
TX	21,700	1,920	0.59
MO	8,817	1,390	8.97
IL	9,637	1,270	0.90
GA	10,600	1,110	0.78
AZ	12,900	1,030	0.95
PA	6,420	614	0.95
IA	3,939	601	0.42
ND	5,244	533	0.80
All States	210,000	20,800	0.85

A = electricity load from S&P Global for data centers marked as planned (not operational as of Q3 2025).
GED = authors calculations.

Figures

Figure 1: Location of Data Centers in the Contiguous United States (https://www.datacentermap.com/usa/).

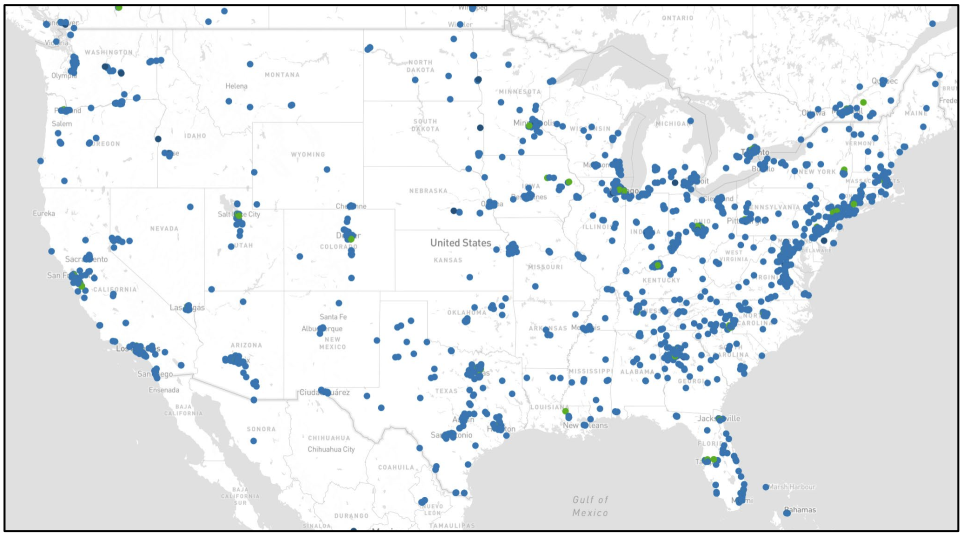

Figure 2: Electricity Consumption from Data Centers by State.

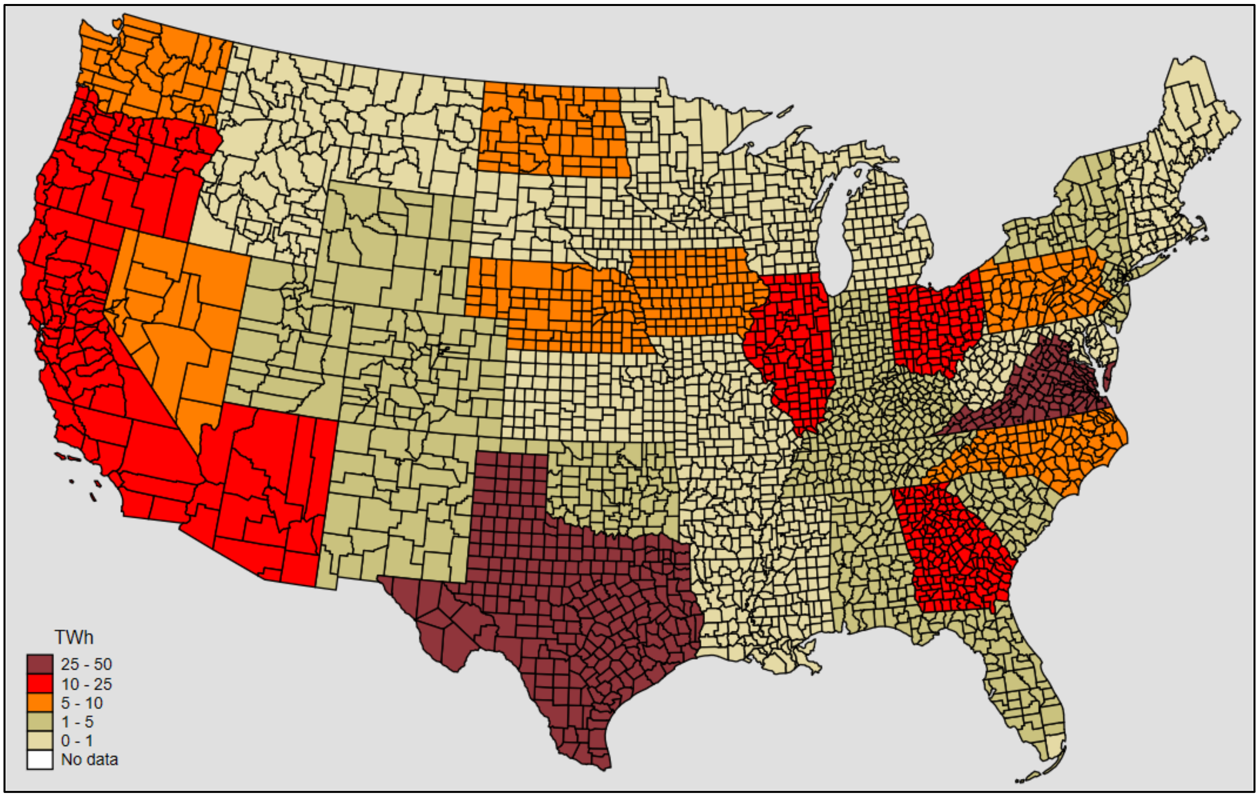

Source: S&P Global (2025) and author's calculations.

Figure 3: Change in $PM_{2.5}$ Concentrations from Data Center Power Consumption.

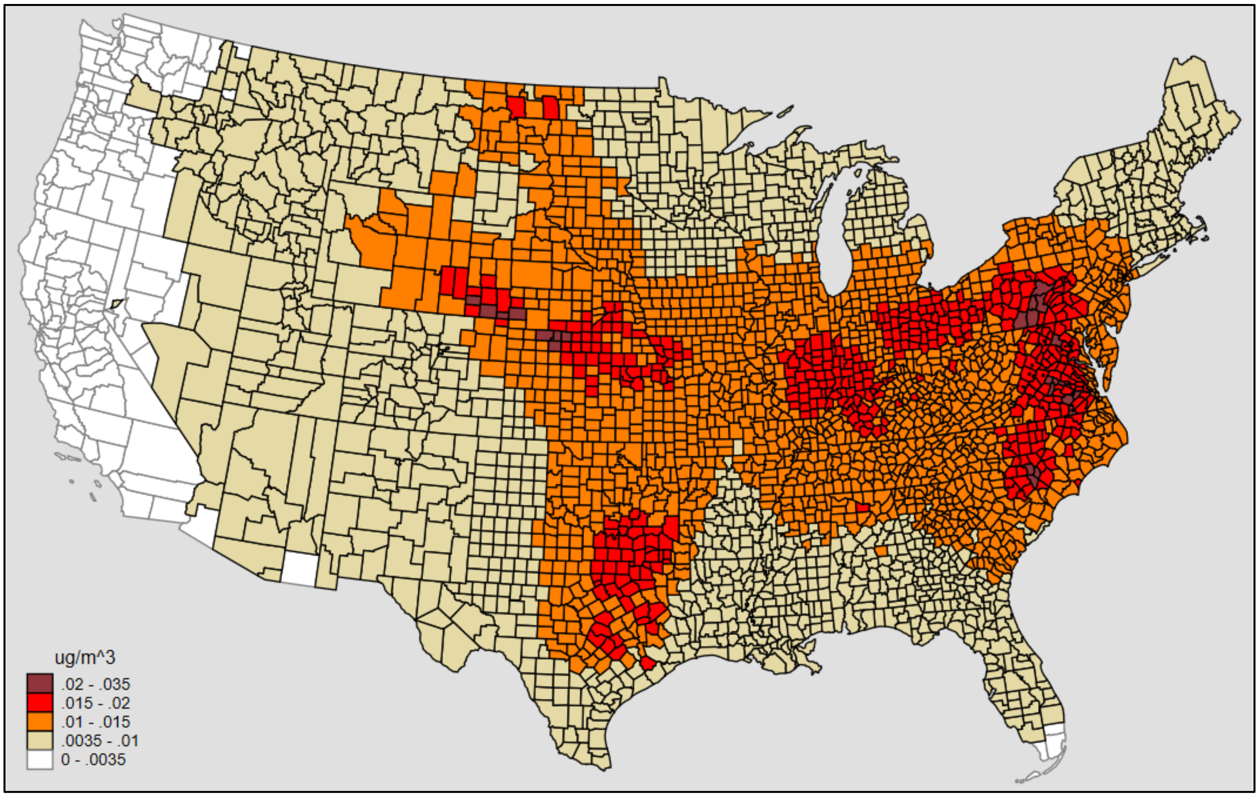

Source: Author's calculations.

Figure 4: Air Pollution GED by County of Incidence from Data Center Power Consumption.

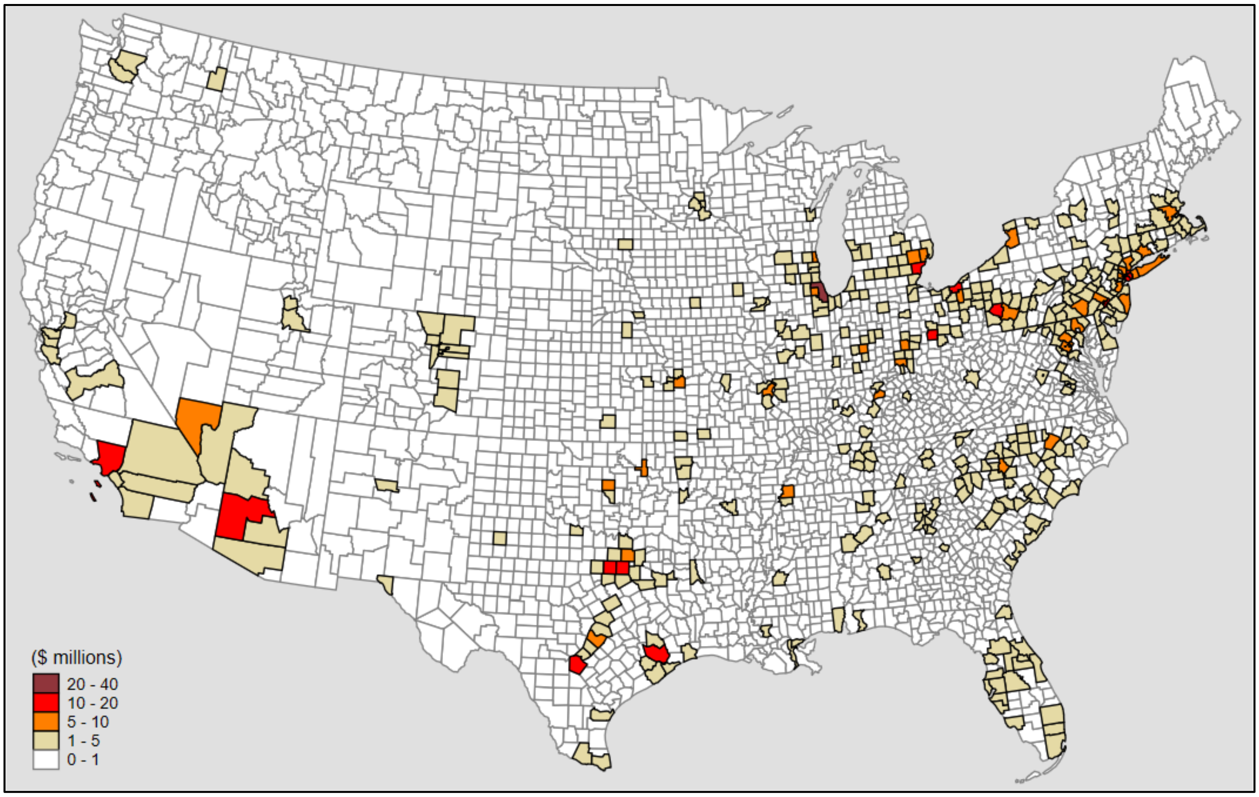

Source: Author's calculations.

Figure 5: GED from Data Centers by State Location of Data Centers.

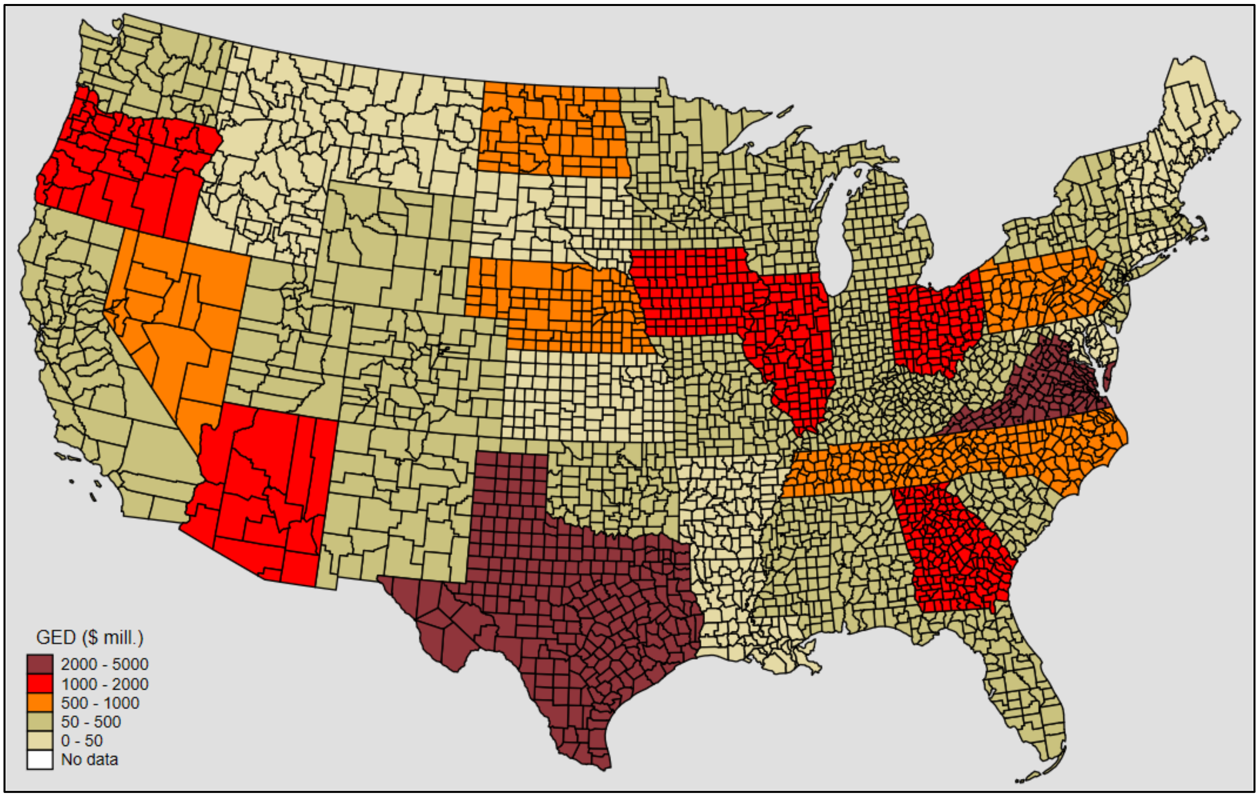

Source: Author's calculations. The GED in figure 5 includes both GHGs and criteria air pollution. Damages are reported according to the state where operational data centers are located.

Figure 6: Comparing AI GED and Contribution to GDP.

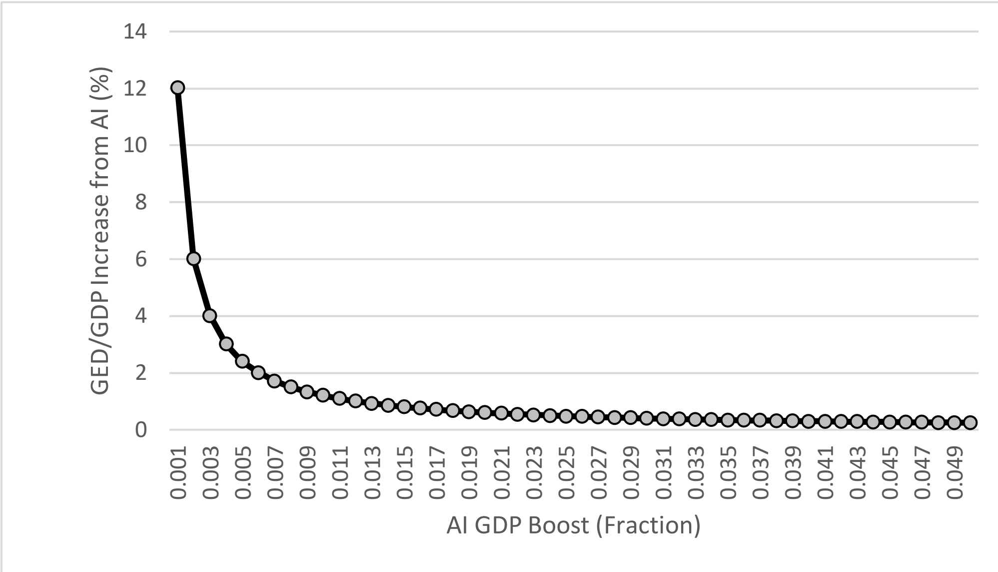

Figure 6 assumes that 15% of all data center activity powers AI training and inference (IEA, 2025). Author's calculations.

References.

1) Blunt, K., and J. Hiller. 2026. “America’s Biggest Power Grid Operator Has an AI Problem-Too Many Data Centers.” *Wall Street Journal,* Jan. 12, 2026. https://www.wsj.com/business/energy-oil/power-grid-ai-data-centers-1235f296
2) Brown, C. 2026. “Data Centers and Your Power Bill.” *New York Times*. Feb. 19, 2026. https://www.nytimes.com/2026/02/19/climate/data-centers-power-bills.html
3) Centers for Disease Control and Prevention (CDC), 2025. https://wonder.cdc.gov/
4) Datacentermap.com. 2025. https://www.datacentermap.com/datacenters/
5) Dennin, L. et al., 2025. "Modeling wildland fire smoke damages in the U.S. and unpacking impact disparities by social vulnerability." *Nature Communications: Earth and Environment*. 6(190): 1 - 19.
6) eGRID, 2026. https://www.epa.gov/egrid
7) EPRI, 2024. “Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption.” Electric Power Research Institute.
8) Gilmore, E. et al., 2018. "An Inter-Comparison of Air Quality Social Cost Estimates from Reduced-Complexity Models" *Environmental Research Letters*: 14(2019) 074016 https://doi.org/10.1088/1748-9326/ab1ab5.
9) Goldman Sachs, 2023. “Generative AI Could Raise Global GDP by 7%.”
10) Han Y., Z. Wu, P. Li, A. Wierman, S. Ren. 2024 “The Unpaid Toll: Quantifying the Public Health Impact of AI.” https://arxiv.org/pdf/2412.06288
11) Hudson, C. 2026. “Locals Are Using AI to Fight Data Centers Being Built in Their Backyards.” *Wall Street Journal*. April 8, 2026. https://www.wsj.com/pro/sustainable-business/locals-are-using-ai-to-fight-data-centers-being-built-in-their-backyards-d0642630
12) IEA, 2025. https://www.iea-4e.org/wp-content/uploads/2025/05/Data-Centre-Energy-Use-Critical-Review-of-Models-and-Results.pdf
13) Krewski D, Jerrett M, Burnett RT, Ma R, Hughes E, Shi Y, Turner MC, Pope CA 3rd, Thurston G, Calle EE, Thun MJ, Beckerman B, DeLuca P, Finkelstein N, Ito K, Moore DK, Newbold KB, Ramsay T, Ross Z, Shin H, Tempalski B. 2009. “Extended follow-up and spatial analysis of the American Cancer Society study linking particulate air pollution and mortality.” *Res Rep Health Eff Inst*. May;(140):5-114; discussion 115-36. PMID: 19627030.
14) Koomey, J.G. 2008. “Worldwide electricity used in data centers.” *Environmental Research Letters*. 3: 034008.
15) Lepeule, Johanna, Francine Laden, Douglas Dockery, and Joel Schwartz. 2012. “Chronic Exposure to Fine Particles and Mortality: An Extended Follow-up of the

Harvard Six Cities Study from 1974 to 2009." *Environmental Health Perspectives*, 120(7): 965–970. https://www.jstor.org/stable/41548777.

16) Mohan, Aniruddh, N.Z. Muller, Akshay Thyagarajan, Randall V. Martin, Melanie S. Hammer, Aaron van Donkelaar. 2024. "Measuring global monetary damages from particulate matter and carbon dioxide emissions to track sustainable growth." Nature Communications: Earth and Environment. 5, 264. https://doi.org/10.1038/s43247-024-01426-3
17) Muller, N.Z., R.O. Mendelsohn. 2007. "Measuring the Damages of Air Pollution." *Journal of Environmental Economics and Management*. 54(1): 1 – 14.
18) Muller, N.Z., R.O. Mendelsohn. 2009. "Efficient Pollution Regulation: Getting the Prices Right." *American Economic Review*. 99(5): 1714 – 1739.
19) Muller, N.Z., R.O. Mendelsohn, W.D. Nordhaus. 2011. "Environmental Accounting: Methods with and Application to the U.S. Economy." *American Economic Review*. 101(5): 1649 – 1675.
20) Mytton, D. and M. Ashtine. 2022. "Sources of data center energy estimates: A comprehensive review." *Joule*. 6: 2032 – 2056.
21) Plumer, B.2024. "Three Mile Island Plans to Reopen as Demand for Nuclear Power Grows." *New York Times*. Sept. 20, 2024. https://www.nytimes.com/2024/09/20/climate/three-mile-island-reopening.html
22) Rust, M. and W. Parker. 2026. "These Cities and States Are Taking Aim at Data Centers." *Wall Street Journal*. April 7, 2026. https://www.wsj.com/tech/ai/these-cities-and-states-are-taking-aim-at-data-centers-3b98adf1?mod=article_inline
23) S&P Global (2025) Data Center KnowledgeBase, December 2025. S&P Global Energy Horizons.
24) Tschofen, P., et al., 2019. "Fine Particulate Matter Damages and Value Added in the United States Economy," *Proceedings of the National Academy of Science*. https://doi.org/10.1073/pnas.1905030116.
25) United States Bureau of Economic Analysis (USBEA). 2026a. Gross Domestic Product. https://www.bea.gov/data/gdp/gdp-state
26) United States Bureau of Economic Analysis (USBEA). 2026b. Personal Income by County, Metro, and Other Areas.https://www.bea.gov/data/income-saving/personal-income-county-metro-and-other-areas
27) United States Department of Energy (USDOE). 2026. Form 923. https://www.eia.gov/electricity/data/eia923/
28) USEPA. 2011. The Benefits and Costs of the Clean Air Act: 1990--2020. Final Report. Office of Air and Radiation, Office of Policy, Washington, D.C.
29) USEPA, 2023. https://www.epa.gov/system/files/documents/2023-12/epa_scghg_2023_report_final.pdf

30) USEPA. 2025. https://www.epa.gov/environmental-economics/mortality-risk-valuation#whatisvsl
31) USEPA, 2026. https://www.epa.gov/air-emissions-inventories/national-emissions-inventory-nei
32) USEPA, 2026. Clean Air Markets Program Data. https://campd.epa.gov/data/custom-data-download
33) Wall Street Journal, 2026. "AI and the Data Center Backlash." https://www.wsj.com/opinion/ai-companies-data-centers-electricity-trump-administration-df665913
34) White House. 2021. "Technical Support Document: Social Cost of Carbon, Methane, and Nitrous Oxide. Interim Estimates Under Executive Order 13990." https://www.energy.gov/sites/default/files/2023-04/57.%20Social%20Cost%20of%20Carbon%202021.pdf

Appendix

Table A1: Operational Data Center Power Consumption and GED by Facility Type.

Facility Type	Count	Elec. Load[A] (Twh)	Facility Area[B] (ft^2)	Power Density (Mwh/$ft.^2$)	GED ($ mill.)	GED($)/ Mwh
Cryptocurrency	121	50,600	192,321	2.61	5,150	102
Hyperscale	432	103,000	186,549	1.36	10,700	104
Retail	1508	25,300	20,273	0.79	2,340	93
Wholesale	525	59,100	113,645	0.97	5,240	89
Other	229	13,700	53,449	0.92	1,190	87
All Facilities	2815	250,000	73,298	1.00	24,620	98

A = electricity load from S&P Global.
B = Average across all facilities of each type.
GED = authors calculations.

Table A2: Sensitivity Analysis.

Scenario	Elec. Load[A] (Gwh)	Total GED[B]	GED GHG	GED LAP	GED ($/MWh)	GED (%) IT GDP[C]
Default	251,000	24,600	20,000	4,590	101	4.7
High Run Rate	343,000	33,600	27,300	6,260	101	6.4
Low SCC	251,000	9,870	5,280	4,590	40	1.9
High Dose Response	251,000	30,700	20,000	10,700	126	5.8

A = electricity load from S&P Global.
B = All GED values in table 5 are in ($millions), expect when expressed relative to MWh.
C = GDP for Data processing, hosting, and related services. These are industries in the Data Processing, Hosting, and Related Services subsector group establishments that provide the infrastructure for hosting and/or data processing services. (USBEA, 2026)

Figures.

Figure A1: Summary Statistics for Data Centers by Type.

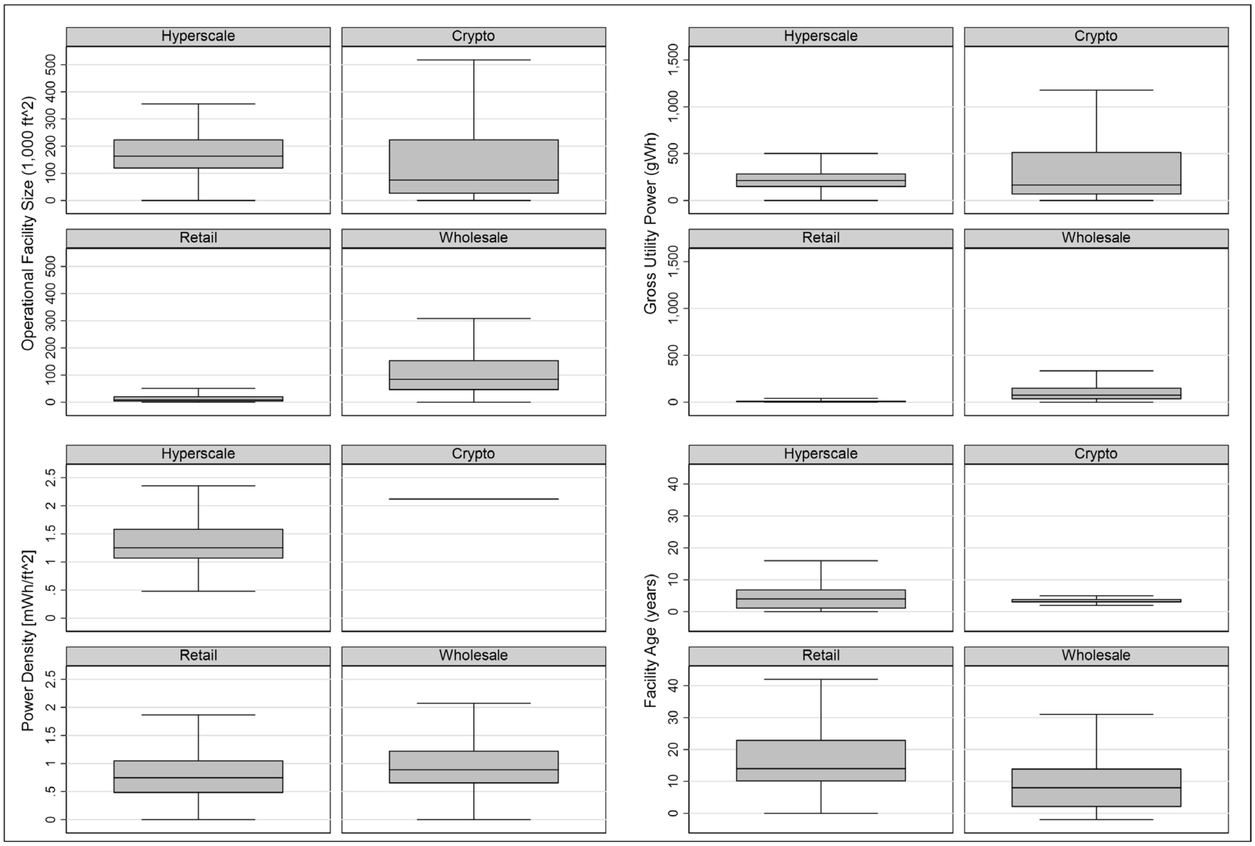

Source: S&P Global (2025) and Author's Calculations.

Figure A2: Change in $PM_{2.5}$ Concentrations from Power Consumption at Planned Data Centers.

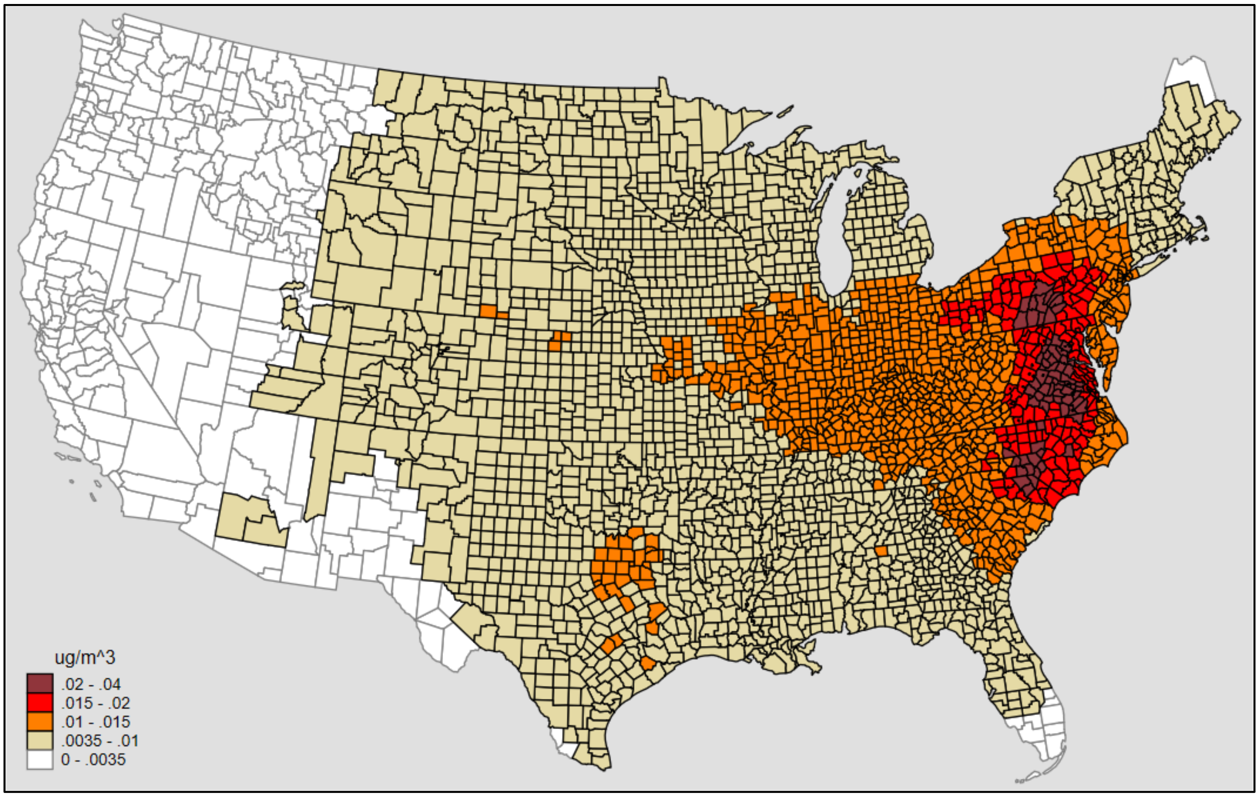

Source: Author's calculations.

Figure A3: GED from Power Consumption at Planned Data Centers.

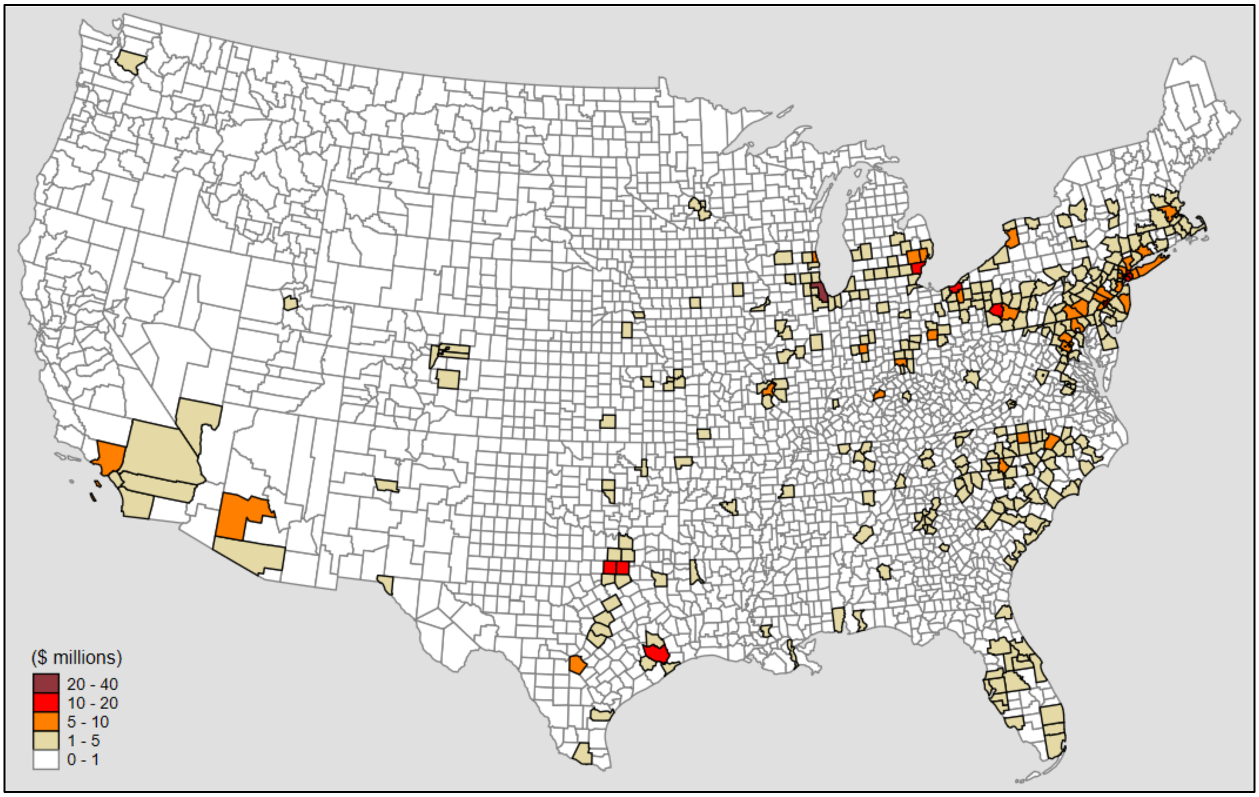

Source: Author's calculations.

NBER WORKING PAPER SERIES

DATA CENTERS AND LOCAL ECONOMIES IN THE AGE OF AI:
A SHIFT--SHARE APPROACH

Fernando E. Alvarez
David Argente
Joyce Chow
Diana Van Patten

Working Paper 35194
http://www.nber.org/papers/w35194

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2026

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w35194

Data Centers and Local Economies in the Age of AI: A Shift--Share Approach
Fernando E. Alvarez, David Argente, Joyce Chow, and Diana Van Patten
NBER Working Paper No. 35194
May 2026
JEL No. D8, O3

ABSTRACT

Data centers are the physical infrastructure behind cloud computing, artificial intelligence, and enterprise software. The rapid diffusion of artificial intelligence (AI) is intensifying demand for compute, accelerating investment in data centers, and raising concerns about the local economic and environmental footprint of these facilities. Their expansion creates a local policy tradeoff. A data center can bring capital investment, construction activity, and specialized employment, but it can also increase demand for electricity, land, and grid capacity. This paper studies these effects at the U.S. county level. We assemble a facility-level panel of global data centers with precise coordinates, scale metrics, and annualized revenue. We map facilities to U.S. counties and combine them with County Business Patterns, county-level IRS income, county-level house prices, and electricity prices. To address endogenous siting, we instrument for data center growth using two shift-share instruments, which leverage pre-existing proximity to InterTubes long-haul fiber nodes and the 1980 county share of U.S. urban college population as shares, and both Chinese and rest-of-the-world data center revenue growth as shifts. The IV estimates show positive effects on total employment, data-processing employment, construction employment, establishments, house prices, and electricity prices at different horizons after data center growth. We also find positive effects on tax returns, adjusted gross income, and wages, while annual payroll responds less robustly. The results suggest that data centers create measurable local activity, increase house prices, and affect local electricity markets through higher prices.

Fernando E. Alvarez
The University of Chicago
Department of Economics
and NBER
f-alvarez1@uchicago.edu

David Argente
Yale University
Yale School of Management
Economics
and NBER
david.argente@yale.edu

Joyce Chow
Yale University
joyce.chow@yale.edu

Diana Van Patten
Yale University
Yale School of Management
and NBER
diana.vanpatten@yale.edu

1 Introduction

Digital services are often experienced as placeless, but the infrastructure that supports them is not. Behind cloud computing, artificial intelligence, streaming, payments, enterprise software, and digital services are physical facilities that occupy land, connect to fiber networks, draw power continuously, and require cooling, security, and grid reliability. These facilities, data centers, have become one of the most visible elements of the digital economy.

The local stakes are high. A data center can look attractive to local governments because it brings a large capital project, construction activity, a property-tax base, and specialized activity in data-processing and related services. At the same time, data centers are energy-intensive and spatially concentrated. They can require expansions, transmission upgrades, and long-term power contracts. They may also use scarce land and water, while creating fewer permanent jobs than a traditional manufacturing plant with similar capital expenditure. The result is a practical policy question: what do counties actually gain, and what costs do they face, upon data centers entry?

This question has become more urgent with the rise of AI. Data centers accounted for around 1.5 percent of global electricity consumption in 2024, or about 415 TWh, and the United States accounted for the largest share of that global data center electricity use (International Energy Agency, 2025). In the United States, the Department of Energy reports that data center electricity use rose from 58 TWh in 2014 to 176 TWh in 2023. This 300% increase represented about 4.4 percent of U.S. electricity consumption, with large uncertainty but substantial projected growth through 2028 (U.S. Department of Energy, 2024; Shehabi et al., 2024). While these are national totals, impacts are highly local. Data centers are spatially clustered, so their effects on labor markets, utilities, and prices are concentrated on exposed counties.

This paper studies whether counties more exposed to data center growth experience changes in employment, establishments, payroll, income, house prices, and electricity prices. The analysis considers long differences from 1995 to years 2005, 2010, 2015, and 2020. We leverage facility-level data center activity from the 451 Research Datacenter Knowledge Base (DCKB), a proprietary S&P Global/451 Research database with site-level information on facilities, capacity, and annualized revenue. We also measure local economic outcomes from County Business Patterns (CBP), IRS Statistics of Income county files, county-level house prices, and county-level electricity prices constructed from EIA Form 861 data.

The main identification challenge is that data centers do not locate randomly. Firms choose locations with favorable power prices, grid capacity, land, tax policy, climate, disaster risk, and fiber connectivity. We address this concern with a two-instrument shift-share design. The first instrument interacts a county's pre-existing proximity to a long-haul fiber node with growth in Chinese data center revenue. The second interacts a county's 1980 share of national urban college population with rest-of-world data center revenue growth, outside the United States and China.

The first stage is strong across horizons.[1] The IV estimation then shows that data center growth has positive local effects on employment and number of establishments. The most persistent labor-market response occurs in data-processing employment, while construction employment is strongest at short horizons, consistent with requirements to build the centers themselves. We also find positive effects on tax returns, adjusted gross income, and wages, although annual payroll responds less robustly. Finally, electricity-price and house-price effects are positive and highly significant across horizons, indicating that data centers place measurable pressure on local energy and real estate markets. In sum, we document that data centers generate localized employment, increase the number of establishments, increase house prices and lead to aggregate income gains, while also raising local electricity prices.

Related literature. This paper is related to work on digital infrastructure and local labor markets. Broadband and internet infrastructure can affect wages, productivity, and firm location, but the effects depend on technology, adoption, and local absorptive capacity (Forman et al., 2012; Kolko, 2012; Akerman et al., 2015). Data centers, however, differ from household broadband access. They are spatially concentrated capital projects, and their effects operate through construction activity, specialized operations, electricity demand, tax base, and local supply chains.

The paper also relates to research on large plant openings and place-based economic development. Greenstone et al. (2010) use runner-up locations for large plant openings to study agglomeration spillovers. Moretti (2010) studies local multipliers, while Busso et al. (2013) and Glaeser and Gottlieb (2008) discuss the incidence and efficiency of place-based policies. Data centers are relevant in this context because local governments often compete for them, but their labor and power intensity differ from traditional manufacturing plants.

Finally, the paper relates to research on data center electricity use. Shehabi et al. (2016) estimate U.S. data center energy use, Masanet et al. (2020) reassess global data center energy demand in light of efficiency improvements, and Shehabi et al. (2024) report recent U.S. data center energy trends and future scenarios. We study a different margin: how the local arrival and expansion of data centers is associated with county-level employment, establishments, income, house prices, and electricity prices. Contemporaneous work by Yue and Zeng (2026) studies the local effects of data-center entry using a difference-in-differences

[1] In the total-employment baseline specifications, first-stage F-statistics range from 786.1 to 907.0.

design. Entry, however, is likely to be an endogenous outcome: data centers are sited in places with favorable infrastructure, energy access, permitting conditions, tax incentives, expected growth, and other local advantages (e.g. JLL 2026; CBRE 2026, 2025). Our instruments are based on a shift-share design, aiming to relax this concern by using historical local exposure interacted with global data-center revenue shocks. This difference in identification, as well as our focus on cumulative data-center scale rather than first entry, likely explains differences in magnitudes.

2 Data

Our analysis combines facility-level information on data centers with county-level measures of employment, business activity, income, house prices, and electricity prices. The empirical unit is a U.S. county observed over time. We use the facility data to measure where data centers operate and how their scale evolves, and we use public county-level sources to measure the local economic conditions that may respond to this expansion.

Data centers. The data center panel comes from the 451 Research Datacenter Knowledge Base (DCKB). DCKB is a proprietary facility-level database maintained by S&P Global Market Intelligence. It covers more than 12,960 data center facilities across over 2,700 providers and 131 countries, and reports detailed information on facility location, ownership, capacity, and annualized revenue (S&P Global Market Intelligence, 2026).

We map each U.S. facility to a county using its geographic coordinates and county boundary shapefiles. We then aggregate facility-level information to the county-year level. The resulting panel records whether a county has an operating data center in a given year, the number of facilities, and the total annualized revenue associated with those facilities. Our main measure of data center activity is cumulative annualized revenue. Namely, for county c and horizon h, with end year $t_h = 1995 + h$, we define

$$\Delta_h R_c = \log\left(\sum_{t \leq t_h} R_{ct}\right) - \log\left(\sum_{t \leq 1995} R_{ct}\right), \tag{1}$$

where R_{ct} is annualized data center revenue in county c and year t.[2] The cumulative measure captures the scale of data center activity that has arrived in a county by a given horizon. We also report specifications using cumulative facility counts, an indicator for whether the county has a data center by the end year, and an ever-treated indicator.

[2]For variables that can be zero, the empirical implementation uses the inverse hyperbolic sine transformation.

The same DCKB data are also used to construct the aggregate shifts in the shift-share design. We build annual measures of data center revenue for the United States, China, and the rest of the world outside the United States and China. The baseline instruments interact historical county exposure with growth in Chinese data center revenue and with revenue growth in the rest of the world outside the United States and China.

County Business Patterns. We measure local business activity using County Business Patterns (CBP). The main CBP outcomes are total employment, total establishments, and annual payroll. These variables provide a broad picture of county-level economic activity in the private sector. We also construct two sectoral employment measures that are especially relevant for data centers. The first is data-processing employment, which captures activity in industries directly related to hosting, processing, and related digital infrastructure services. In the SIC period, we use SIC 7374, Computer Processing and Data Preparation and Processing Services. In the NAICS period, we use NAICS 518210, Data Processing, Hosting, and Related Services. The second is construction employment, defined using SIC 15 and NAICS 236. This outcome is intended to capture the local build-out margin associated with constructing large facilities.[3]

Income and tax outcomes. We use county-level IRS Statistics of Income files to measure broader local income responses. These data provide the number of tax returns, adjusted gross income, and wages at the county-year level. These outcomes are useful because employment gains around data centers need not imply broad income gains for local residents. Some jobs may be temporary, specialized, or filled by workers outside the county. The IRS outcomes therefore allow us to ask whether data center expansion is visible in broader county-level income aggregates, beyond employment and establishments.

House prices. We measure local housing-market responses using the county-level House Price Index (HPI) from the Federal Housing Finance Agency (FHFA). Specifically, we use the FHFA annual county HPI workbook, which provides repeat-sales house price indexes at the county-year level. The HPI outcome is the percent change in the index relative to the 1995 baseline.

[3] Because the industry classification system changes from SIC to NAICS during the sample, the sector-specific series should be interpreted as approximate, not as perfectly harmonized industry measures. This caveat is less important for the aggregate CBP outcomes, which are measured consistently at the county level.

Electricity prices. We measure electricity prices using a county-level panel constructed from EIA Form 861. The underlying EIA files report utility-level revenue, electricity sales, and customer counts by sector, as well as service-territory information linking utilities to counties. The county-level price is constructed by aggregating revenue and MWh across utilities serving a county and then dividing revenue by sales. Our main electricity outcome is the total price. Conceptually, this variable captures the average price paid by electricity consumers in a county.

Historical shares The shift-share design uses two predetermined county-level exposure measures. The first is proximity to long-haul fiber infrastructure. Data centers require reliable high-capacity connectivity, and historical fiber routes created persistent geographic advantages for some locations. We measure this exposure as negative log distance to the nearest InterTubes long-haul fiber node. The node locations come from the Internet Atlas and InterTubes projects, which document long-haul fiber links, nodes, conduits, and conduit sharing among U.S.-based providers (Durairajan et al., 2013, 2015). We assign FIPS codes to node locations and use the NBER county distance database to compute county-to-node distances (Roth, 2024). The measure is predetermined relative to the post-1995 expansion studied in the paper. In constructing it, we exclude six nodes for which we were unable to verify clear pre-2000 tube or conduit evidence. The second exposure measure is the county's 1980 share of the national urban college population. This variable is constructed from historical Census county data using total population, urban population, and college population in 1980 (U.S. Census Bureau, 1998; Schroeder et al., 2025). The motivation is that counties with larger concentrations of urban college-educated population in 1980 were more exposed to the early human-capital and urban-demand conditions that shaped the geography of digital infrastructure. Because the share is measured fifteen years before the 1995 baseline, it is predetermined with respect to the subsequent growth in data center revenue that we use as the shift.

3 Industry background

A fast-growing, concentrated industry Data centers house servers, storage, and networking equipment, together with the power and cooling systems needed to keep that equipment running. Their growth reflects several demand forces: the migration of enterprise computing to the cloud, software-as-a-service, streaming and digital public services, and now AI training and inference. To better understand this industry's dynamics, we leverage data center facility-level data. Namely, Figure 1 considers the total number of data centers globally and across time. The figure shows a sharp acceleration after the mid-2000s, with especially rapid growth in Asia-Pacific, the United States and Canada, and Europe. The U.S. market is also highly concentrated. Figure 2 maps state-level data center counts and shows the outsized role of Virginia and Texas.

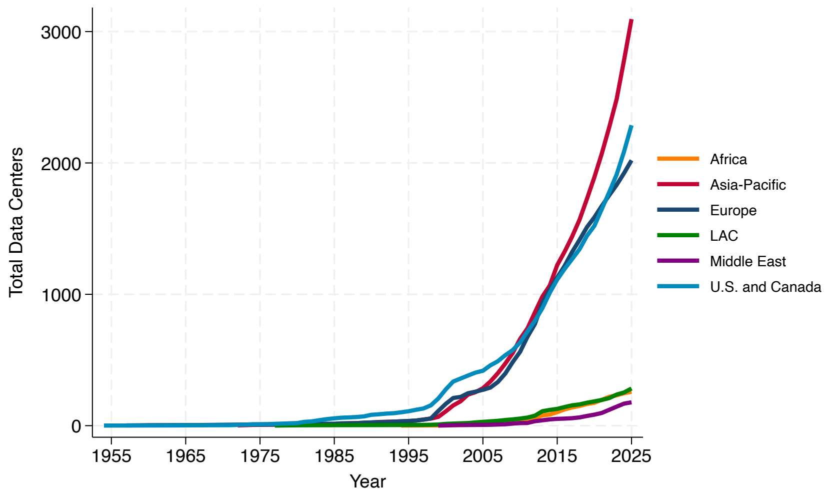

Figure 1: Growth in data centers by subregion

Notes: The figure plots cumulative data center counts by subregion across time.

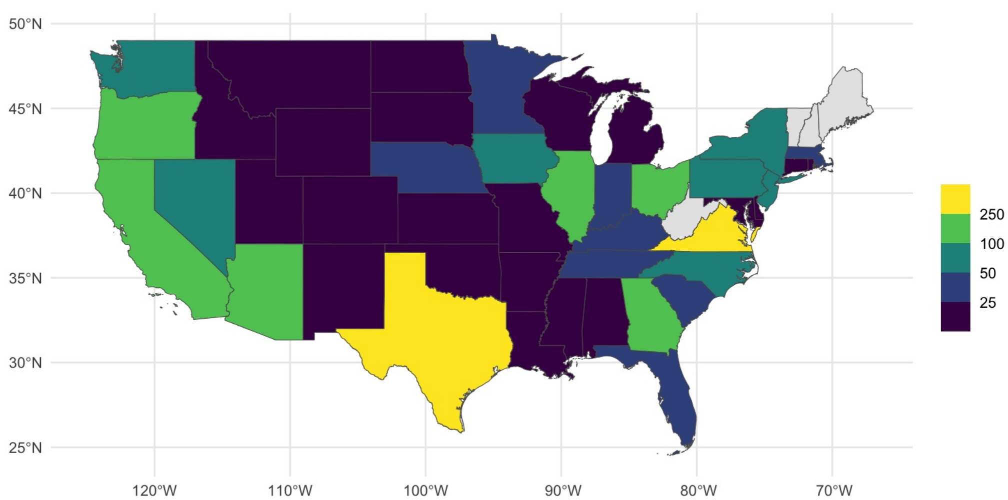

Figure 2: Number of data centers by state

Notes: The figure reports state-level facility counts in the contiguous United States.

Revenue, capacity and construction costs A useful way to think about a data center is as a factory for computation. Its output is compute, storage, and network services. Its inputs are land, structures, power equipment, cooling, servers, specialized labor, and connectivity. Industry capacity is therefore measured in several ways. Figure 3 shows the growth in revenue from 1995 onwards. Revenue in China is zero prior to 1995, and very low elsewhere.

Figure 3: Cumulative total revenue

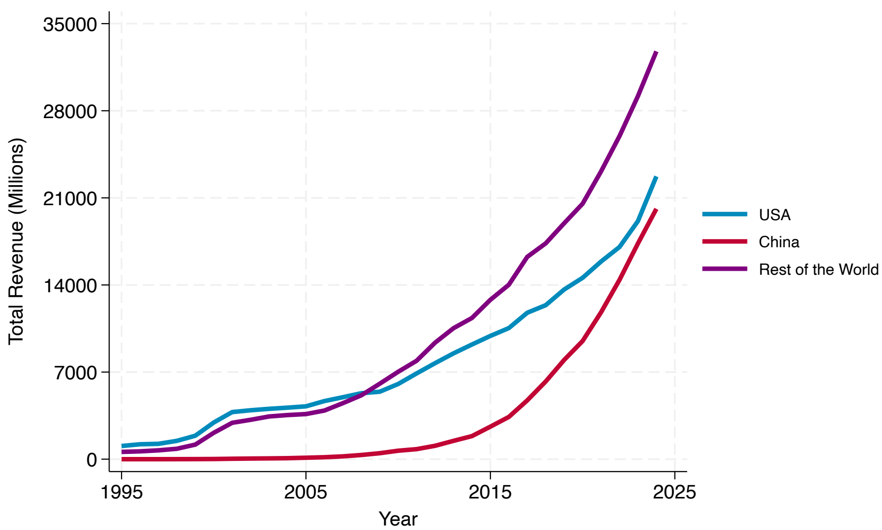

Notes: The figure shows the total cumulative revenue (in millions of US dollars) over time.

Figure 4 shows that capacity measures grew sharply after 1995, especially after 2015. Namely, Panels (a) descrives the total operational space required by the data centers. Panel considers the total count of racks; a measure of the number of equipment cabinets, and Panel (c) relies on an Uninterruptible Power Supply (UPS) measure, which depedns on the electricity capacity available to IT equipment and is one of the most common industry measures of data center scale.

Critical capacity, the electrical and mechanical infrastructure that supports IT load, accounts for 53% of construction costs (Cushman & Wakefield, 2025). This helps explain why data centers can generate large construction demand and large utility-system requirements even when permanent employment is modest. Land is another source of geographic differentiation. There are large differences in land prices across selected data center markets. Established markets such as Chicago, Phoenix, and Virginia command large land premia relative to several emerging markets, like Iowa, Indiana, and Minneapolis (Cushman & Wakefield, 2025). Those premia reflect the same factors that make data center location endogenous: fiber, power availability, permitting, proximity to demand, and the ability to

Figure 4: Cumulative capacity growth of data centers

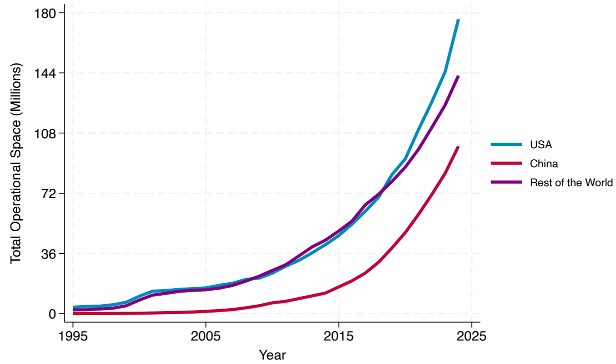

(a) Cumulative operational space

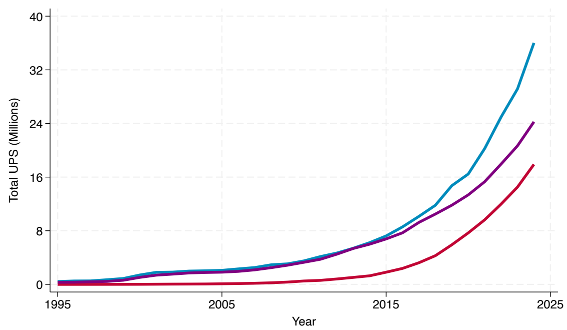

(b) Cumulative UPS

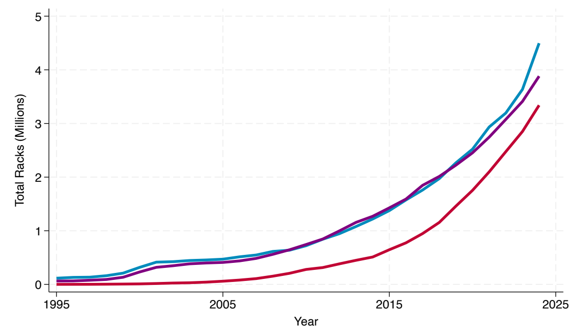

(c) Cumulative total racks

Notes: The figure shows the cumulative dynamics of data center outcomes. Panel (a) considers total operational space (in millions of square feet), Panel (b) considers the cumulative UPS consumed, and Panel (c) displays the cumulative number of racks.

expand.

Electricity, cooling, and efficiency Electricity is central to data center operations. In a typical facility, power reaches IT equipment through switch gear, uninterruptible power supplies, and power distribution units. Cooling systems and other support loads sit alongside the servers, storage devices, and network equipment. This architecture explains why the same facility can affect both local employment and local electricity prices.

Location choice Data center location depends on power prices, grid capacity, fiber access, latency, land, climate, water, permitting, taxes, and disaster risk. These forces make naive comparisons of data center and non-data center counties hard to interpret. They also motivate the shift-share design proposed below to account for endogenous site choice.

4 Empirical strategy

The baseline second-stage equation for horizon h is

$$\Delta_h Y_c = \beta_h \widehat{\Delta_h R_c} + \varepsilon_{ch}, \tag{2}$$

where $\Delta_h Y_c$ is the change in the log outcome from 1995 to year $1995 + h$, and $\Delta_h R_c$ is the change in log cumulative data center annualized revenue. For the HPI, $\Delta_h Y_c$ is the percentage change in the index relative to 1995. The first stage is

$$\Delta_h R_c = \pi_{1h} Z^F_{ch} + \pi_{2h} Z^U_{ch} + v_{ch}. \tag{3}$$

The two excluded instruments are

$$Z^F_{ch} = [-\log(\text{distance to nearest fiber node}_c)] \times \Delta_h R^{China}, \tag{4}$$

$$Z^U_{ch} = (\text{1980 urban college share}_c) \times \Delta_h R^{ROW}, \tag{5}$$

where $\Delta_h R^{China}$ is the change in log data center annualized revenue in China. Finally, $\Delta_h R^{ROW}$ is the corresponding rest-of-world change, excluding the United States and China. In the data center time series, both China and the rest-of-the-world have virtually no data center revenue before the 1995 baseline, which makes these shocks post-baseline expansions.

The empirical design follows the logic of shift-share or Bartik instruments (Bartik, 1991), where identification can be framed as relying on the exogeneity of predetermined shares (Goldsmith-Pinkham et al., 2020) *or* on the quasi-randomness of shocks (Borusyak et al., 2022). While our shares are predetermined, our shocks are arguably quasi-random from individual U.S. counties' perspective, and in the spirit of the *China shock* literature (Autor et al., 2013).

5 Descriptive statistics for main outcomes

Table 1 reports summary statistics for the 1995 to 2020 long differences. Counties are highly heterogeneous. The mean change in data-processing employment is large relative to the median county, but many counties have no data center exposure. The data center treatment is sparse: by 2020, 7.6 percent of counties have positive data center presence over the horizon, while the mean change in log cumulative data center revenue is 0.190.

Table 2 shows how treatment exposure grows across horizons. The average data center revenue and facility-count exposure rises from 2005 to 2020. The share of counties with

Table 1: Summary statistics for 1995 to 2020 long differences

Variable	Obs.	Mean	Std. dev.
% Δ employment	3,173	0.212	0.858
% Δ data-processing employment	3,173	0.525	1.627
% Δ construction employment	3,173	-1.430	1.530
% Δ annual payroll	3,173	0.969	1.039
% Δ establishments	3,173	0.091	0.488
% Δ tax returns	3,130	0.338	0.297
% Δ adjusted gross income	3,130	0.958	0.358
% Δ wages	3,130	0.861	0.340
% Δ house price index	2,046	1.093	0.441
% Δ electricity price	2,954	0.038	0.022
% Δ cumulative data center revenue	3,188	0.190	0.869
% Δ cumulative data center facilities	3,188	0.130	0.524
Data center present over horizon	3,188	0.076	0.264

Notes: Variables are county-level long differences from 1995 to 2020. Rows labeled as percent changes are inverse-hyperbolic-sine differences in the implementation, which approximate log differences for positive values and retain observations with zeros. The house price index row is the percentage change in the HPI relative to 1995.

positive data center presence during the horizon rises from 3.0 percent in 2005 to 7.6 percent in 2020.

Table 2: Mean treatment exposure by horizon

	2005	2010	2015	2020
Cumulative data center revenue	0.080	0.115	0.172	0.190
	(0.554)	(0.658)	(0.815)	(0.869)
Cumulative data center facilities	0.043	0.066	0.106	0.130
	(0.274)	(0.346)	(0.459)	(0.524)
Data center present over horizon	0.030	0.044	0.065	0.076
	(0.170)	(0.206)	(0.246)	(0.264)

Notes: The table reprts mean treatment by horizon. Standard deviations are in parenthesis. The cumulative outcomes report inverse-hyperbolic-sine long differences in cumulative data center revenue and facility counts. The deta center presence outcome reports an indicator equal to one when a county has positive data center presence during the horizon.

6 First stage

Table 3 reports the first stage for the main endogenous treatment. Both instruments are highly predictive. The first-stage F-statistics range from 786.1 to 907.0 across horizons in the total-employment sample. The coefficient on the 1980 urban college share interacted with rest-of-world growth is negative in the long difference with respect to 2005 because the rest-of-world revenue shock is negative over 1995 to 2005 in the constructed data center aggregate. In later horizons, however, it is positive.

Table 3: First stage for cumulative data center annualized revenue

End year	$-$ log distance × China shock		1980 urban college share × ROW shock		First-stage F	R^2	Obs.
	Coef.	SE	Coef.	SE			
2005	0.010***	(0.003)	-569.681***	(14.811)	786.11	0.335	3,130
2010	0.012***	(0.002)	118.561***	(3.071)	811.24	0.342	3,131
2015	0.016***	(0.002)	126.021***	(3.131)	904.22	0.366	3,131
2020	0.016***	(0.002)	130.279***	(3.238)	906.95	0.367	3,129

Notes: The dependent variable is the horizon-specific change in log cumulative data center annualized revenue. Robust standard errors are in parentheses. The China and ROW shocks are long differences in data center annualized revenue outside the United States, with ROW excluding both the United States and China. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7 Main IV results

Employment and establishments Table 4 reports the IV estimates related to employment and number of establishments. Total employment increases at most horizons, with statistically significant effects in 2005, 2010, and 2020. The 2020 coefficient is 0.039, implying that a one-unit increase in cumulative data center revenue is associated with a 3.9 percent increase in total employment between 1995 and 2020. This effect is not mechanically driven by data center jobs themselves. When we instead use total county employment excluding data-processing employment, the coefficient is 0.038. This similarity is expected, since data-processing employment accounts for less than one percent of total county employment.

The more direct labor-market response appears in data-processing employment. The estimates imply that a one-unit increase in cumulative data center revenue raises data-processing employment by roughly 27 percent in 2015 and 29 percent in 2020. Construction employment follows a different time profile. The largest coefficients appear in 2005 and 2010, consistent with a build-out phase in which the local labor demand created by data centers is concentrated around construction and installation. Establishments also rise at every horizon, with coefficients between 0.026 and 0.062.

House and electricity prices House prices increase as well. As shown in Table 4, HPI coefficients are positive and statistically significant in every horizon, with a 2020 coefficient of 0.177. Electricity prices also increase in all horizons. The 2020 coefficient is 0.009, and the electricity-price estimates are statistically significant at the 1 percent level in all baseline horizons. The table therefore points to both benefits and costs. Data centers generate measurable local activity, especially in sectors directly connected to their construction and operation, and they also place upward pressure on house prices and electricity prices.

Table 4: Main IV estimates: employment, establishments, house prices, and electricity prices

	End year			
Outcome	2005	2010	2015	2020
Employment	0.022*	0.063***	0.014	0.039***
	(0.012)	(0.020)	(0.010)	(0.011)
Data-processing employment	0.092	-0.119	0.265***	0.294***
	(0.067)	(0.073)	(0.048)	(0.053)
Construction employment	0.600***	0.473***	0.052***	0.071***
	(0.117)	(0.075)	(0.019)	(0.019)
Establishments	0.026***	0.034***	0.046***	0.062***
	(0.007)	(0.008)	(0.007)	(0.008)
House price index	0.239***	0.081***	0.124***	0.177***
	(0.039)	(0.019)	(0.021)	(0.028)
Electricity price	0.007***	0.008***	0.008***	0.009***
	(0.002)	(0.002)	(0.001)	(0.001)
Observations, employment sample	3,130	3,131	3,131	3,129
First-stage F, employment sample	786.11	811.24	904.22	906.95

Notes: Each cell reports a separate 2SLS coefficient of the listed long-difference outcome on the horizon-specific change in log cumulative data center annualized revenue. Robust standard errors are in parentheses. All specifications use the two excluded instruments reported in Table 3. Rows are log-change approximations, with inverse-hyperbolic-sine differences used for variables that can be zero. The HPI outcome is the percentage change in the house price index relative to 1995. Electricity-price regressions use 2,986, 2,997, 2,975, and 2,954 observations across the four horizons; HPI regressions use 2,036 observations in each horizon. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Payroll, income, and wages Table 5 reports outcomes for annual payroll, tax returns, adjusted gross income, and wages. The estimates show positive effects on several broad income aggregates. Tax returns, adjusted gross income, and wages are positive and statistically significant in every horizon. The 2020 coefficients imply increases of about 5.9 percent in tax returns, 8.2 percent in adjusted gross income, and 5.9 percent in wages. Annual payroll is positive and significant in 2005 and 2010, but not in the longer horizons.

Table 5: Main IV estimates: payroll, income, and wages

	End year			
Outcome	2005	2010	2015	2020
Annual payroll	0.046***	0.049**	-0.005	0.017
	(0.018)	(0.021)	(0.013)	(0.013)
Tax returns	0.021***	0.050***	0.058***	0.059***
	(0.007)	(0.010)	(0.009)	(0.010)
Adjusted gross income	0.021***	0.055***	0.065***	0.082***
	(0.008)	(0.013)	(0.009)	(0.011)
Wages	0.020***	0.040***	0.054***	0.059***
	(0.008)	(0.010)	(0.009)	(0.009)
Observations, payroll sample	3,130	3,131	3,131	3,129
First-stage F, payroll sample	786.11	811.24	904.22	906.95

Notes: Each cell reports a separate 2SLS coefficient of the listed long-difference outcome on the horizon-specific change in log cumulative data center annualized revenue. Robust standard errors are in parentheses. All specifications use the two excluded instruments reported in Table 3. Rows are log-change approximations, with inverse-hyperbolic-sine differences used for variables that can be zero. Observation counts differ slightly across IRS and average-pay outcomes because of data availability. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

8 Alternative treatment definitions

Table 6 reports 1995 to 2020 estimates using alternative measures of data center exposure. The first column is the baseline continuous revenue treatment. The second uses cumulative data center facilities. The third uses an indicator for whether the county has a data center over the 25-year horizon. The fourth uses an ever-data center indicator. The latter is included as a diagnostic because it is defined over the full available panel rather than only by the horizon.

The qualitative pattern is stable across treatment definitions. Data-processing employment, construction employment, establishments, house prices, and electricity prices are positive in all columns. The binary specifications have larger coefficients because the treatment changes from a continuous log revenue measure to an indicator. These estimates should therefore be interpreted as effects of moving from no data center presence to data center presence, not as elasticities.

Table 6: Alternative endogenous treatment measures, 1995 to 2020

Outcome	Log DC revenue	Log DC facilities	Ever DC
Employment	0.039***	0.065***	0.175***
	(0.011)	(0.019)	(0.049)
Data-processing employment	0.294***	0.490***	1.327***
	(0.053)	(0.091)	(0.178)
Construction employment	0.071***	0.130***	0.278***
	(0.019)	(0.032)	(0.085)
Establishments	0.062***	0.102***	0.288***
	(0.008)	(0.015)	(0.029)
Electricity price	0.009***	0.015***	0.039***
	(0.001)	(0.002)	(0.005)
House price index	0.177***	0.314***	0.772***
	(0.028)	(0.053)	(0.124)
Annual payroll	0.017	0.031	0.066
	(0.013)	(0.021)	(0.062)
Tax returns	0.059***	0.094***	0.278***
	(0.010)	(0.016)	(0.032)
Adjusted gross income	0.082***	0.136***	0.369***
	(0.011)	(0.019)	(0.037)
Wages	0.059***	0.098***	0.263***
	(0.009)	(0.017)	(0.034)
Observations, CBP outcomes	3,129	3,129	3,129
Observations, IRS outcomes	3,130	3,130	3,130
Observations, electricity	2,954	2,954	2,954
Observations, HPI	2,036	2,036	2,036
First-stage F, CBP sample	906.95	776.48	437.91
First-stage F, IRS sample	907.29	776.77	438.09

Notes: Each cell reports a separate 2SLS coefficient for the 1995 to 2020 outcome. The first two columns use horizon-specific treatments. The ever-data-center column uses treatment status over the full available panel and is included as a diagnostic specification. Rows are log-change approximations, with inverse-hyperbolic-sine differences used for variables that can be zero. The house price index row is the proportional change in the FHFA county-level HPI relative to 1995. Robust standard errors are in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

9 Robustness

Table 7 reports 1995 to 2020 robustness checks for the main continuous revenue treatment. The first column repeats the baseline. The second absorbs state fixed effects and clusters by state. With state fixed effects, the positive effects on data-processing employment, establishments, and house prices remain statistically significant. The total-employment coefficient falls from 0.039 to 0.026 and is no longer statistically significant. Construction and electricity-price estimates are also no longer statistically significant with state fixed effects.

Table 7: Robustness for 1995 to 2020 estimates

Outcome	Baseline	State fixed effects
Employment	0.039***	0.026
	(0.011)	(0.016)
Data-processing employment	0.294***	0.246***
	(0.053)	(0.073)
Construction employment	0.071***	0.003
	(0.019)	(0.033)
Establishments	0.062***	0.063***
	(0.008)	(0.008)
Electricity price	0.009***	0.001
	(0.001)	(0.002)
House price index	0.177***	0.136***
	(0.028)	(0.034)
Annual payroll	0.017	0.020
	(0.013)	(0.020)
Tax returns	0.059***	0.042***
	(0.010)	(0.012)
Adjusted gross income	0.082***	0.067***
	(0.011)	(0.011)
Wages	0.059***	0.056***
	(0.009)	(0.011)
Observations, CBP outcomes	3,129	3,129
Observations, IRS outcomes	3,130	3,130
Observations, electricity	2,954	2,954
Observations, HPI	2,036	2,035
First-stage statistic, CBP sample	906.95	46.25
First-stage statistic, IRS sample	907.29	46.27

Notes: Each cell reports a separate 2SLS coefficient for the 1995 to 2020 outcome on the change in log cumulative data-center annualized revenue. The state fixed effects column absorbs state indicators and clusters standard errors by state; the first-stage statistic in that column is the Kleibergen-Paap Wald statistic reported in the log. Rows are log-change approximations, with inverse-hyperbolic-sine differences used for variables that can be zero. The house price index row is the proportional change in the FHFA county-level HPI relative to 1995. Robust standard errors are in parentheses in the baseline column; state-clustered standard errors are in parentheses in the state fixed effects column. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

10 Conclusion

This paper studies the local effects of data center expansion using county-level long differences and two historical shift-share instruments. The instruments combine pre-1995 digital infrastructure and 1980 human-capital geography with global data center growth shocks. The estimates show positive effects on total employment, data-processing employment, establishments, construction employment during early build-out horizons, house prices, and electricity prices. We also find positive effects on tax returns, adjusted gross income, and wages. Annual payroll responds less robustly.

The policy implication is therefore mixed. Data centers create economic activity, especially in directly related sectors and during construction, and they are associated with larger county-level income aggregates. They also raise electricity prices and are associated with higher house prices, which may benefit property owners while increasing costs for renters and prospective homebuyers. Counties considering data center projects should therefore weigh employment, establishment, and aggregate income gains against pressures on electricity markets and local housing affordability.

References

Akerman, A., Gaarder, I., Mogstad, M., 2015. The skill complementarity of broadband internet. Quarterly Journal of Economics 130, 1781–1824.

Autor, D.H., Dorn, D., Hanson, G.H., 2013. The china syndrome: Local labor market effects of import competition in the united states. American Economic Review 103, 2121–68. URL: https://www.aeaweb.org/articles?id=10.1257/aer.103.6.2121, doi:10.1257/aer.103.6.2121.

Bartik, T.J., 1991. Who Benefits from State and Local Economic Development Policies? W.E. Upjohn Institute for Employment Research, Kalamazoo, MI.

Borusyak, K., Hull, P., Jaravel, X., 2022. Quasi-experimental shift-share research designs. Review of Economic Studies 89, 181–213.

Busso, M., Gregory, J., Kline, P., 2013. Assessing the incidence and efficiency of a prominent place based policy. American Economic Review 103, 897–947.

CBRE, 2025. North america data center trends h2 2025. URL: https://www.cbre.com/insights/books/north-america-data-center-trends-h2-2025. cBRE Research.

CBRE, 2026. Key factors to consider for effective data centre site selection. URL: https://www.cbre.co.uk/insights/articles/key-factors-to-consider-for-effective-data-centre-site-selection. cBRE Research.

Cushman & Wakefield, 2025. U.S. data center development cost guide. Industry report. Data-center development cost and land cost benchmarks.

Durairajan, R., Barford, P., Sommers, J., Willinger, W., 2015. Intertubes: A study of the US long-haul fiber-optic infrastructure, in: Proceedings of ACM SIGCOMM, pp. 565–578.

Durairajan, R., Ghosh, S., Tang, X., Barford, P., Eriksson, B., 2013. Internet atlas: A geographic database of the internet, in: Proceedings of ACM HotPlanet.

Forman, C., Goldfarb, A., Greenstein, S., 2012. The internet and local wages: A puzzle. American Economic Review 102, 556–575.

Glaeser, E.L., Gottlieb, J.D., 2008. The economics of place-making policies. Brookings Papers on Economic Activity , 155–253.

Goldsmith-Pinkham, P., Sorkin, I., Swift, H., 2020. Bartik instruments: What, when, why, and how. American Economic Review 110, 2586–2624.

Greenstone, M., Hornbeck, R., Moretti, E., 2010. Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. Journal of Political Economy 118, 536–598.

International Energy Agency, 2025. Energy and AI. IEA report. Analysis of electricity demand from data centres and artificial intelligence.

JLL, 2026. 2026 global data center outlook. URL: https://www.jll.com/en-us/insights/market-outlook/data-center-outlook. jones Lang LaSalle Research.

Kolko, J., 2012. Broadband and local growth. Journal of Urban Economics 71, 100–113.

Masanet, E., Shehabi, A., Lei, N., Smith, S., Koomey, J., 2020. Recalibrating global data center energy-use estimates. Science 367, 984–986.

Moretti, E., 2010. Local multipliers. American Economic Review: Papers and Proceedings 100, 373–377.

Roth, J., 2024. County distance database. National Bureau of Economic Research data file.

Schroeder, J.P., Van Riper, D., Manson, S., Knowles, K., Kugler, T., Roberts, F., Ruggles, S., 2025. IPUMS national historical geographic information system: Version 20.0 [dataset].

Shehabi, A., Smith, S.J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M.A.B., Holecek, B., Koomey, J.G., Masanet, E.R., Sartor, D.A., 2024. 2024 United States Data Center Energy Usage Report. Technical Report LBNL-2001637. Lawrence Berkeley National Laboratory.

Shehabi, A., Smith, S.J., Sartor, D.A., Brown, R.E., Herrlin, M., Koomey, J.G., Masanet, E.R., Horner, N., Azevedo, I.L., Lintner, W., 2016. United States Data Center Energy Usage Report. Technical Report LBNL-1005775. Lawrence Berkeley National Laboratory.

S&P Global Market Intelligence, 2026. 451 research datacenter knowledgebase. Product documentation and database accessed through Yale subscription.

U.S. Census Bureau, 1998. Usa counties: Historical county data files. Historical county data used for 1980 population, urban population, and education tabulations.

U.S. Department of Energy, 2024. DOE releases new report evaluating increase in electricity demand from data centers. Press release summarizing the 2024 United States Data Center Energy Usage Report.

Yue, D., Zeng, Y., 2026. The local economic effects of data center entry. Working paper, SSRN.

OCTOBER 2025 | VOL. 42, NO. 10

APA

ZONING PRACTICE

Unique Insights | Innovative Approaches | Practical Solutions

The Physical Footprint of Artificial Intelligence

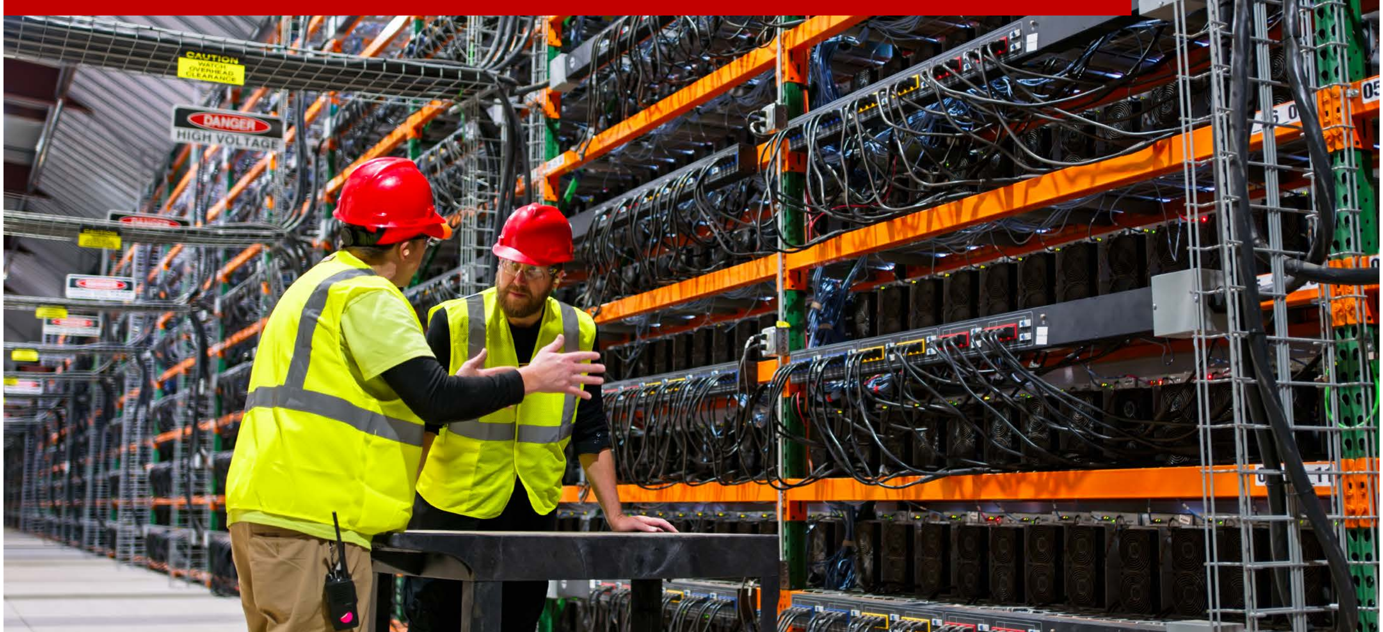

The Physical Footprint of Artificial Intelligence

By Charlie Nichols, AICP

Every time you ask ChatGPT, Gemini, or Claude a question, you are tapping into a sprawling, power-hungry network of machines. Somewhere, a data center's processors are whirring, fans are spinning, and megawatts of electricity are flowing.

Artificial intelligence (AI) may feel virtual, but its footprint is intensely physical. Behind every chatbot interaction, predictive algorithm, or autonomous system lies a vast network of data centers, power generators, and electricity transmission and distribution infrastructure. As vast as it is now, the demand for computing power is growing at an exponential rate, and local zoning is on the front lines.

This issue of *Zoning Practice* explores the physical effects of AI deployment and highlights core considerations for local planning and zoning. It begins with a summary of the land use characteristics of the system of data centers that host and serve contemporary AI models before highlighting noteworthy regulatory approaches and areas of opportunity for zoning updates and land use decision-making processes.

Data center infrastructure in the United States, 2025 (Credit: NREL)

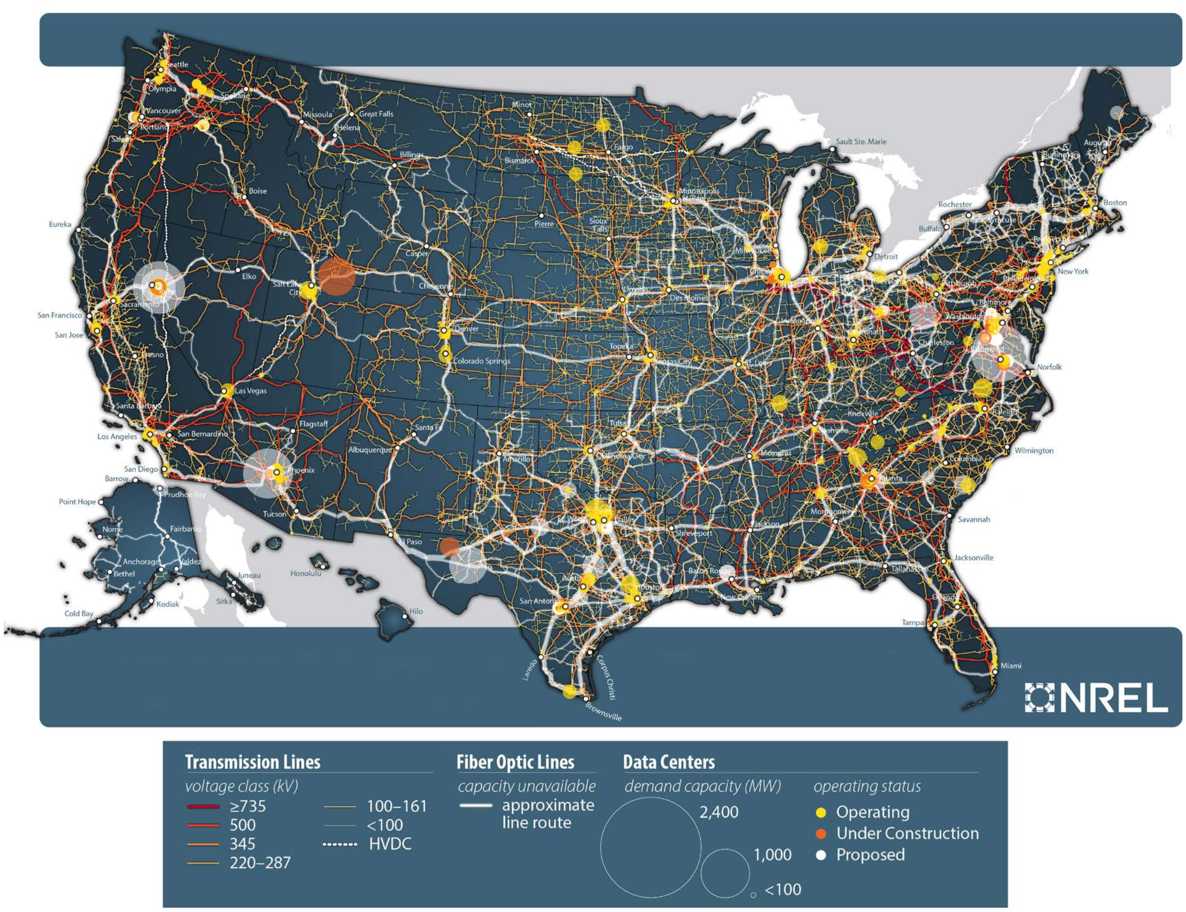

What Are the Physical Needs of AI?

When we think about artificial intelligence, we often imagine abstract ideas or algorithms, software, or maybe a chat assistant or a robot. But AI is deeply physical. It runs on powerful hardware that lives in large buildings, draws enormous amounts of electricity, and requires robust infrastructure to keep it cool and operational. These needs are shaping land use decisions in ways many communities have never dealt with before.

AI Lives in Data Centers

The primary home of AI is the data center. These are large, sometimes windowless, buildings filled with servers, networking equipment, and backup systems. While some are sleek and high-tech, many look like simple warehouses. But inside, the technology is anything but simple.

AI workloads require far more computational power than traditional cloud computing. That means more servers packed with graphics processing units (GPUs), which are optimized for machine learning tasks. These GPUs are energy-intensive and generate a significant amount of heat (Shehabi et al. 2024; Casey 2025).

This is why the design, location, and infrastructure of data centers have become such a big deal. For example, Meta's Altoona, Iowa, data-center campus has more than five million square feet of space and is still growing (Miller 2022).

Data centers themselves fall into several distinct categories. Edge or micro facilities are the smallest, often modular container-sized enclosures ranging from a few hundred to a few thousand square feet. Enterprise data centers, typically operated by corporations or universities, can range from about 5,000 to 50,000 square feet, sometimes larger. Colocation facilities lease space to multiple tenants and often fall between 50,000 and 600,000 square feet, with many averaging around 150,000 square feet. At the largest scale are hyperscale data centers, typically built by major cloud or AI providers, which can easily reach hundreds of thousands of square feet per building and exceed one million square feet across a campus (Zhang 2023).

While many forecasts focus on power demand rather than square footage, it is possible to translate one into the other. Deloitte estimates that AI-driven data centers could require up to 123 gigawatts (GW) of capacity in the U.S. by 2035, compared to roughly 4 GW today (Stansbury et al. 2025). Real-world projects suggest that every megawatt of IT load requires between 5,000 and 12,000 square feet of total building area. Applying that ratio to 123 GW implies a national buildout of 615 million to 1.48 billion square feet of data center space, equivalent to about 22 to 53 square miles. Land use estimates point in a similar direction, with recent projects averaging 0.5 to 1.5 acres per MW, which would translate to roughly 96 to 288 square miles of U.S. land devoted to AI-related data center campuses by 2035 (Stansbury et al. 2025).

A proposed 612-acre hyperscale data center campus in Cedar Rapids, Iowa (Credit: QTS)

AI Needs Lots of Electricity

Power demand is one of the most critical limiting factors in scaling AI. The U.S. Department of Energy's Secretary of Energy Advisory Board notes that legacy hyperscale data centers have typically connected at 20–50 megawatts (MW), but utilities are now receiving AI-driven connection requests for single campuses of 300–1,000 MW (2024). To put the low end of that new range in context, a 300 MW facility running around the clock would consume about 2.6 terawatt-hours a year—roughly the annual electricity use of 250,000 U.S. homes (calculated with the U.S. EIA average of 10,500 kWh per household). These unprecedented loads are forcing planners, utilities, and regulators to rethink siting, transmission capacity, and community-impact mitigation.

This demand is driving data centers to locate near existing transmission infrastructure, substations, or power plants. In some cases, new substations or transmission lines are being proposed just to support AI infrastructure. Local planners are being asked to approve not just buildings, but energy projects with regional impacts.

There is also growing concern about the climate impacts of AI. Researchers estimate that the cumulative carbon emissions from AI models could reach 3.66 to 8.72 million tons in the U.S. alone—the equivalent of driving an average gasoline-powered car nine to 22 billion miles (Ding et al. 2025; USEPA 2024). This has led to pressure for data centers to run on renewable energy, adding another layer of land use complexity as solar or wind farms are proposed nearby or colocated together with data centers.

AI Needs Water and Cooling

All that power generates heat, and that heat has to go somewhere. Most data centers use a combination of air- and watercooling systems. Some of the largest AIfocused facilities can consume hundreds of thousands of gallons of water per day for evaporative cooling (Lei et al. 2025; Shehabi et al. 2024; Selsky 2022). That's raising concerns in water-scarce regions or places where water infrastructure is already stretched thin.

For example, in The Dalles, Oregon, a dispute between Google and the city over water use became national news when the city council approved a water agreement to support Google's data center expansion, despite local concerns about long-term water availability (Selsky 2022).

Water and cooling infrastructure also raise siting questions. Should data centers be allowed in areas with limited water supply? What happens when a tech company becomes one of the largest users of municipal water? These questions are starting to reach planning commissions and city councils.

Top-five U.S. Google data centers by annual water withdrawals, 2024 (Credit: [Google's 2025 Environmental Report](#))

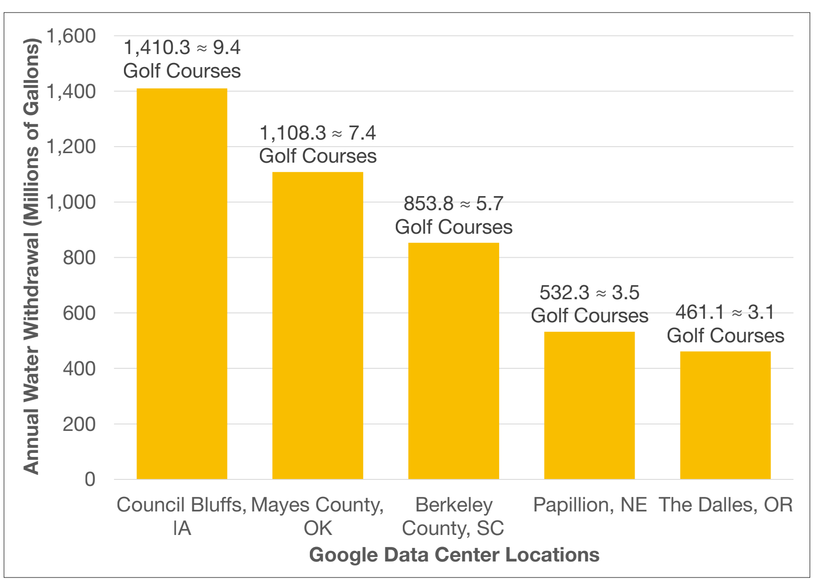

AI Needs Fiber and Connectivity

Finally, AI infrastructure depends on high-speed fiberoptic connections. Training models and delivering AI services both require fast, reliable data transmission. This can drive the need for new fiber lines, telecom infrastructure, or even small-cell installations in rural or suburban areas (RVA LLC 2025; Walker 2024).

It's not just big cities seeing these investments. Some rural areas are gaining interest from AI developers because they offer space, lower land costs, and cooperative local governments—provided they can offer fiber access and a willing utility partner.

How Is AI Infrastructure Regulated (or Not)?

If your city or county does not already have a data center, just wait. The odds are increasing that a tech company, or the utility that serves them, will soon come knocking. Yet most local governments are not fully prepared to regulate AI infrastructure. In many places, the regulatory framework is either nonexistent or built for a different era of technology.

Zoning Codes Rarely Mention AI or Data Centers

Many zoning codes still make no explicit reference to "artificial intelligence" or even to "data centers." Where definitions are absent, planners may choose to slot these facilities into broad buckets such as warehousing, light-industrial, or public-utility uses, even though the buildings may be packed wall-to-wall with servers instead of pallets.

Yet these facilities behave very differently from the categories they're often shoehorned into, and there are many reasons why local governments may want to specifically define data center uses (Morley 2022). Their continuous operation demands megawatts of electricity and, in many climates, hundreds of thousands of gallons of cooling water per day; the equipment generates heat and noise; and the employment footprint is minimal. When such impacts are overlooked, communities can be blindsided—as happened in Prince William County, Virginia, where approval of a massive datacenter corridor sparked backlash over noise, power delivery, and land use compatibility.

Recognizing this mismatch, an increasing number of jurisdictions have begun to write data-center-specific rules. Loudoun County, Virginia, imposes façade, screening, lighting, and pedestrian-connectivity standards on by-right data centers to blunt visual impacts while leveraging their tax base (§4.06.02). Prince William County uses a Data Center Opportunity Zone Overlay to funnel projects to infrastructure-served parcels and require design review (§32-509). Missoula County, Montana, offers a different model. The county's ordinance, crafted for cryptocurrency mines, confines those operations to industrial zones and requires them to offset 100 percent of their electricity use with renewable energy (§5.10). Because cryptocurrency mines and large-scale data centers both run continuously, draw high-density power, and employ few on-site workers, planners can adapt the same toolkit—clear land use definitions, targeted overlay districts, and energy-focused performance standards—to data centers when communities want comparable safeguards.

Looking ahead, AI training clusters dwarf the loads discussed in 2022, with utilities now fielding single-campus interconnection requests of 300 MW and more. The zoning fundamentals remain the same, but the stakes are higher. Without proactive definitions, locational criteria, and impact standards, local governments risk conceding critical decisions about land, water, and grid capacity to developers' timetables rather than community goals.

Data Center Alley in Loudoun County, Virginia (Credit: Gerville/iStock/Getty Images Plus)

Many AI Facilities Are Allowed by Right

In areas that do allow data centers by right, local officials often have little authority to influence their design or siting (Morley 2022). Developers may be able to build massive facilities with only administrative approval. If the project complies with the basic zoning and building code, it can move forward, even if it brings significant impacts to neighboring properties or the local infrastructure system.

This hands-off, by-right approach can leave neighbors in the dark when a campus that draws 100 MW or more of power is permitted the same way a warehouse is. Such facilities may also require hundreds of thousands of gallons of cooling water per day and generate continuous low-frequency noise from chillers, pumps, and backup generators (Van Geet and Sickinger 2024). Without a public-hearing trigger, residents may not learn what is coming until the bulldozers roll.

That said, relying on discretionary use permits alone is not a perfect fix. Case-by-case approvals can introduce uncertainty, increase timelines, and duplicate reviews that utilities already perform when they decide whether to supply the necessary electricity and water. A more balanced strategy is to embed objective, use-specific standards (e.g., caps on sound at the property line, requirements for renewable-energy procurement, and water-recycling targets) directly into the zoning code. Guidance from the Urban Land Institute shows how clear definitions, overlay districts, and measurable performance thresholds can give developers predictability while still protecting community interests (Miet 2024). By pairing these standards with early coordination among planners, utilities, and residents, communities can address local impacts without resorting to duplicative or open-ended discretionary reviews.

Infrastructure Approvals May Be Handled Separately

Adding to the complexity, the infrastructure needed to support AI such as transmission lines, substations, power generation facilities, battery energy storage. and fiber installations is often regulated under different frameworks. Utilities may have their own review and siting authority at the state level, which can bypass local land use processes entirely.

Large solar or wind projects, for example, are pre-empted from local control in more than 20 U.S. states, leaving local governments to vet the data-center building, while the power generation facility that feeds it is debated elsewhere (Gomez and Morley 2023; Morley 2025). Fragmented approvals make it hard for planners to tally cumulative effects such as substations, access roads, or groundwater withdrawals.

Battery-energy-storage systems (BESS) create another layer of complexity, and a clear trend of data centers colocating BESS on-site is accelerating (ZincFive 2024). Some states exempt utility-scale BESS that are colocated with generation assets, while others treat them as industrial equipment needing only an electrical permit. Where local authority does apply, recent guidance recommends clear definitions, district regulations, and objective safety standards, thermal-runaway monitoring, minimum setbacks, and emergency-response plans to avoid ad-hoc hearings (Ross and Vadali 2024).

Developers are now bundling data centers with on-site renewables and storage in microgrid “energy parks,” aiming to bypass long interconnection queues and control energy costs. Recent projects in Texas and Virginia pair hundreds of megawatts of generation and storage with adjacent server halls, creating hybrid campuses that straddle state energy-facility review, regional transmission rules, and local zoning (DiGangi 2025). To keep pace, planners can identify jurisdictional triggers early, embed measurable performance standards (e.g., noise caps, screening, or renewable-energy sourcing) in their codes,

The Eland Solar-plus-Storage Center in Kern County, California (Credit: The Desert Photo/iStock/Getty Images Plus)

and coordinate with utilities so local and state reviews proceed on aligned timelines.

Environmental Review Is Inconsistent

Environmental review of AI infrastructure also varies widely. In states that require environmental impact statements (EIS), large-scale data centers may undergo detailed scrutiny. But in states without EIS laws, or for smaller projects, there may be minimal analysis of water use, energy consumption, or greenhouse gas emissions (Morris 2024).

Even where review is required, the focus may be on the building itself, rather than the full ecosystem of impacts. For example, if a local code does not require review of off-site power infrastructure or supporting utility upgrades, critical issues related to energy delivery, environmental impact, or long-term capacity may fall through the cracks.

Local Governments Are Starting to Catch Up

Local governments are no longer standing still while hyperscale campuses spring up at the edge of town. Since 2023, a wave of city councils, county boards, and planning commissions have begun moving data centers out of catch-all industrial categories and into their own, better-defined regulatory boxes. Some jurisdictions, such as Atlanta, now require special-use permits tied to energy, water, and noise studies (Ordinance 25-O-1063). Others, such as Cedar Rapids, Iowa, leverage community-benefit agreements to ensure local reinvestment when a project wins approval (Pratt 2025).

Approaches vary, but the trend is unmistakable: Communities are adopting objective, use-specific standards rather than relying solely on ad-hoc discretionary permits. Some ordinances steer projects into infrastructure-served corridors, others set caps on sound and water use, and a growing number link approvals to renewable energy procurement or on-site battery storage. Table 1 highlights seven recent examples illustrating the breadth of new zoning language, overlay districts, and design guidelines that together show local governments are indeed catching up.

Table 1. Examples of Recent Local Regulatory Updates for Data Centers

Jurisdiction	How it regulates data-center impacts
Atlanta, GA	Requires a special-use permit for every new data center and empowers the city council to review water-consumption, energy-efficiency, and noise-mitigation plans (Ordinance 25O1063, 2024)
Brainerd, MN	Prohibits data centers unless the planning commission approves a conditional-use permit that addresses cooling noise and utility demand (Ordinance No. 1581, 2025)
Chandler, AZ	Adds a data center use category; limits the use to Planned Area Development zones and sets size, generator-testing and water-recycling standards (Ordinance No. 5033, 2022)
Tempe, AZ	Requires a water use plan and enhanced setbacks next to homes and schools, and “innovation hubs” (Ordinance No. O2025-23, 2025)
Phoenix, AZ	Defines “data center,” restricts locations, and introduces design standards such as façade articulation and noise studies (Ordinance G-7396, 2025)
Sugar Grove, IL	Creates a dedicated district with height limits, façade screening, and a master-utility-plan requirement (Ordinance No. 2022-1206B, 2022)
Frederick County, MD	Establishes an overlay zone that limits where data centers can be built (Bill No. 25-05, 2025)

What Should Planners Be Thinking About?

Artificial intelligence may sound futuristic, but the decisions that shape its physical footprint are being made today. Local governments that wait too long to prepare may find themselves reacting to projects rather than guiding them. So what should planners be thinking about now?

Think About Scale

AI infrastructure often hides in plain sight until its true footprint emerges. What looks like a single “warehouse” can blossom into a portfolio buildout—multiple server halls, two substations, a battery yard, and a 30-inch water main, all staged over a decade (USDOE SEAB 2024). To avoid approving these megaprojects one slice at a time, some jurisdictions now demand a phased master plan up front. For example, Loudoun County, Virginia, requires every data-center rezoning to include a "Data Center Development Plan" showing the full buildout of power feeds, cooling infrastructure, and utility corridors before the first site plan is approved (2025).

Regional utilities are following suit by running scenario-based load models to test whether transmission and groundwater supplies can keep up. A 2024 white paper by Energy + Environmental Economics describes how such models informed Portland (Oregon) General Electric's latest integrated-resource plan and helped local planners identify future right-of-way corridors for two new 230-kV lines (Riu et al. 2024). By asking for phased utility exhibits and participating in utility load-growth scenarios, planners can make sure each new server hall fits into a system-wide picture rather than becoming an isolated surprise.

Many comprehensive plans still treat "technology infrastructure" as an afterthought, yet data center proposals are now shaping decisions on land supply, energy policy, water allocation, and broadband.

Think About Alignment With Your Plans

Many comprehensive plans still treat "technology infrastructure" as an afterthought, yet data-center proposals are now shaping decisions on land supply, energy policy, water allocation, and broadband. Start by inventorying where AI-related facilities touch existing plan elements—utilities, environmental stewardship, economic development—and flag the gaps.

One emerging best practice is to link data-center approvals directly to community climate goals. Embedding such benchmarks in comprehensive plans or codes gives planners clear decision criteria and ensures that new AI infrastructure advances, rather than conflicts with, local resiliency objectives.

Plans can also weave data-center growth into broadband and workforce strategies. The U.S. Department of Energy's 2024 report on AI infrastructure recommends that local governments coordinate land-use designations with state broadband-expansion maps so that fiber corridors serving data centers double as backbone routes for underserved neighborhoods (USDOE SEAB 2024). Aligning these layers up front helps planners negotiate public-benefit clauses—such as dark-fiber setasides or training programs, rather than scrambling for concessions late in the process.

Updating your plan first and then adopting measurable standards that flow from it gives applicants clarity, while ensuring projects advance the community's long-term vision.

Think About Infrastructure Capacity

AI campuses can overwhelm local utilities faster than many other land uses. Virginia's Joint Legislative Audit and Review Commission estimates that data centers will require 11 gigawatts (GW) of new electric generation and transmission in that state alone by 2035, roughly one-third of Dominion Energy's entire current system (VJLARC 2024). National modeling by Energy + Environmental Economics shows a similar surge, with some balancingareas seeing load grow 25 percent in a single decade under an "AI-high" scenario (Riu et al. 2024).

Water systems face parallel stress. At Google's complex in The Dalles, Oregon, public records show cooling demand could top one-quarter of the city's current supply, prompting a 2023 agreement that pauses future phases unless new wells come online (Selsky 2022). Quincy, Washington, responded to similar pressures by creating a special water rate class and meter fee for data centers to fund infrastructure upgrades (2025). These examples point to tools planners can

The Three Mile Island nuclear power plant in Middleton, Pennsylvania, which is coming back online to power Microsoft data centers (Credit: gsheldon/ iStock Editorial/ Getty Images Plus)

adopt: cumulative-demand studies embedded in utility master plans, tiered rate structures that recover capital costs, and permit conditions that link new construction to confirmed water-capacity projects.

Electric and water systems are only part of the picture. Broadband providers may need additional conduit banks, and public works departments often discover that construction traffic surpasses road-design volumes. Objective, use-specific standards, such as requiring a utilityinfrastructure plan that maps ultimate substations, mains, and fiber routes, plus haul-route and pavement-repair agreements, give planners leverage without duplicating state or utility reviews.

Think About Cumulative Impacts

A single 30 MW data center can feel benign, yet clusters of 10 or more along one corridor may push peak electric load past a gigawatt, double truck traffic during construction, and raise ambient sound by up to 10 dBA at nearby homes (VJLARC 2024). Project-by-project review often misses these system-level effects, so several jurisdictions now require applicants to look beyond their parcel lines.

Clustering can also amplify benefits if managed deliberately. Developers in Texas and Virginia now pair multiple server halls with a shared microgrid that combines on-site solar, wind, and battery storage—an "energy-park" model that eases interconnection delays and helps regions meet renewable-energy goals (DiGangi 2025). By mapping preferred corridors for both data centers and their supporting infrastructure, planners can steer growth to areas where capacity, compatibility, and community returns align.

Think About Equity and Community Benefits

Data-center projects promise major capital investment but generate few long-term jobs and can offload noise, truck traffic, and resource use onto nearby neighborhoods. Additionally, new cost analyses show that ordinary ratepayers are already footing most of the bill for AI's voracious appetite for electricity.

Monitoring Analytics, the independent market monitor for PJM Interconnection, the largest regional transmission organization in the U.S., calculated that between 2024 and 2025 data-center electricity demand added about $25 to the typical household's monthly bill (Biryukov 2025). PJM now projects that AI and data-center demand will double the region's energy use by 2033, whereas growth would have been only 15 percent by 2040 without new campuses (JLARC 2024).

In response to this and other similar projections of effects on ratepayers, lawmakers in New Jersey (AB 5466), Oregon (HB 3546), and other states have introduced bills or tariffs to place data centers in a separate rate class or require them to "bring their own clean power," so everyday customers are not forced to subsidize the electricity needs of trillion-dollar tech companies (Levy 2025). More communities are also moving to tie approvals to arrangements that deliver measurable local benefits.

For example, Cedar Rapids, Iowa, required QTS to sign a community benefits agreement (CBA) that will return about $18 million over 20 years for workforce training, broadband expansion, and green-infrastructure projects (Pratt 2025). Legal guidance stresses clear milestones, third-party verification, and enforcement clauses to keep such agreements credible (Eisenson 2023).

Meanwhile, Quincy, Washington, created a special water rate class for data centers in 2024, adding higher volumetric charges and meter fees earmarked for new wells and main upgrades. Targeted surcharges turn one user's high demand into system-wide resilience.

By weaving CBAs and host-community fees into zoning approvals or development agreements, planners can ensure that AI infrastructure acts as a catalyst for broader community gain rather than an enclave of private benefit.

A North Dakota data center using nonconductive fluid to cool servers rather than air or water cooling systems (Credit: halbergman/E+)

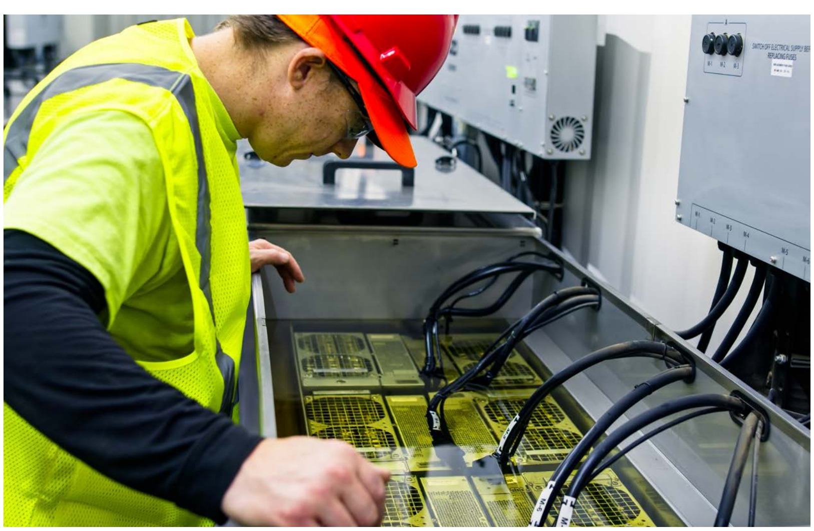

Where Can Planners Learn More?

As artificial intelligence infrastructure expands, planners have a growing need to stay informed about what these facilities are, how they function, and how to plan for them thoughtfully. The good news is that several helpful resources already exist, and more are emerging every year.

Follow the Energy

Many AI-related land use challenges stem from energy demand. That means energy planning organizations are a good place to start. Resources from the U.S. Department of Energy, National Renewable Energy Laboratory, and Lawrence Berkeley National Laboratory offer insights into data center energy use, grid impacts, and cooling technologies (Shehabi et al. 2024; USDOE SEAB 2024; Van Geet and Sickinger 2024).

State and regional energy offices are also useful partners. They can help planners understand energy trends, forecasted demand, and opportunities to align AI-related development with state energy goals.

Watch the Water

Water use is another key issue, especially in places facing drought or groundwater depletion. Reports from the U.S. Environmental Protection Agency, as well as local water utilities and watershed management agencies, can help assess water-related impacts of AI infrastructure.

Planners can also look to academic and journalistic research on water use in cooling systems, which varies significantly based on the type of cooling and climate zone (Berreby 2024).

Track Technology and Land Use Trends

For a broad view of how technology affects land use, the Lincoln Institute of Land Policy and the Urban Land Institute have both published helpful materials. These organizations explore how emerging technologies from AI to autonomous vehicles are reshaping cities, infrastructure, and land markets.

Local case studies can also be instructive. Some jurisdictions have started sharing lessons learned from planning for large-scale data centers or tech campuses. For example, Loudoun (2024; 2025) and Fairfax (2024) Counties in Virginia offer planning documents and staff reports that shed light on real-world challenges and solutions.

Build Cross-Sector Relationships

Planning for AI infrastructure requires collaboration. It touches on land use, utilities, economic development, and environmental protection. Building relationships with energy providers, water utilities, economic development groups, and regional planning agencies can help planners spot opportunities and anticipate challenges.

Conferences like the American Planning Association's National Planning Conference, Grid Forward, or Smart Cities Connect often include sessions on technology infrastructure. These events are a great way to hear from peers and industry experts.

AI infrastructure is no longer a far-off idea; it's already shaping land use decisions in communities across the country. For planners, this presents both challenges and opportunities. By understanding what AI infrastructure is, what it requires, and how it fits into broader planning goals, local governments can prepare for development that is sustainable, equitable, and forward-looking.

As with many emerging trends, the best path forward is to stay curious, build partnerships, and think holistically. AI may be powered by algorithms, but the future it creates will depend on human decisions, including the choices planners make today.

References and Resources

Berreby, David. 2024. "As Use of AI Soars, So Does the Energy and Water It Requires." *Yale Environment 360*, February 6.

Biryukov, Nikita. 2025. "Power Companies Warn Lawmakers About Their Plans to Tackle Rising Bills." *New Jersey Monitor*, April 25.

Casey, Evan. 2025. "Microsoft Built Five Data Center Campuses in This Iowa City. Here's What Wisconsin Can Expect." *Wisconsin Public Radio*, May 15.

DiGangi, Diana. 2025. "Microgrid 'Energy Parks' Could Ease Strain from Rising Power Demand, Report Says." *Utility Dive*, July 23.

Ding, Zhaohao, Jianxiao Wang, Yiyang Song, Xiaokang Zheng, Guannan He, Xiupeng Chen, Xiupeng Chen, Tiance Zhang, Wei-Jen Lee, and Jie Song. 2025. "Tracking the Carbon Footprint of Global Generative Artificial Intelligence." *The Innovation* 6(5): 100866.

Eisenson, Matthew. 2023. "Experts Identify Best Practices for Negotiating and Drafting Community Benefits Agreements." *Climate Law*, September 27.

Fairfax (Virginia) Department of Planning Development, County of. 2024. ***Data Centers Report and Recommendations***.

Gomez, Alexsandra, and David Morley. 2023. ***Solar@Scale: A Local Government Guidebook for Improving Large-Scale Solar Development Outcomes***. Chicago: American Planning Association; Washington, DC: International City/County Management Association.

Lei, Nuoa, Jun Lu, Arman Shehabi, and Eric Masanet. 2025. ***The Water Use of Data Center Workloads: A Review and Assessment of Key Determinants***. Berkeley, CA: Lawrence Berkeley National Laboratory.

Levy, Marc. 2025. "As Electric Bills Rise, Evidence Mounts That Data Centers Share Blare. States Feel Pressure to Act." *Associated Press,* August 8.

Loudoun (Virginia), County of. 2024. "Data Center Growth and Energy Constraints." Board of Supervisors Transportation and Land Use Committee Information Item, June 20.

Loudoun (Virginia), County of. 2025. Data Center Standards & Locations.

Miet, Hannah. 2024. "Local Guidelines for Data Center Development." Washington, D.C.: Urban Land Institute.

Miller, Rich. 2022. "The New MegaCampuses: The World's Largest Data Center Projects." *Data Center Frontier*, November 1.

About the Author

Charlie Nichols, AICP, is the Director of Planning and Development for Linn County, Iowa. He leads a 15-person department and has written pioneering zoning ordinances for utility-scale solar, nuclear energy, and hyperscale data centers. Nichols received his master's degree in urban and regional planning from the University of Iowa and has been working in the field of planning for over 10 years. Outside of work, he enjoys tending to his backyard chickens and working on home renovation projects with his wife and three children.

Morley, David. 2022. "Zoning for Data Centers and Cryptocurrency Mining." *Zoning Practice*, June. American Planning Association.

Morley, David. 2025. "Wait, Who Approves Large-Scale Solar Siting?" *APA Blog*, March 3.

Morris, Jackson. 2024. "Data Centers Gobbling Up Existing Nukes Threatens Grid Decarb Goals." *National Resources Defense Council Expert Blog*, July 11.

Pratt, Richard. 2025. "Cedar Rapids Council Approves Development Agreement for QTS Data Center Project." *Corridor Business Journal*, January 30.

Prince William (Virginia) Planning Office, County of. 2022. DPA2021-0020: Data Center Opportunity Zone Overlay District Comprehensive Review.

Quincy (Washington), City of. 2025. Rate Resolution 25-709.

Riu, Isabelle, Dieter Smiley, Stephen Bessasparis, and Kushal Patel. 2024. ***Load Growth Is Here to Stay, But Are Data Centers?*** San Francisco: Energy + Environmental Economics.

Ross, Brian, and Monika Vadali. 2024. "Battery Energy Storage Systems." *Zoning Practice*, March.

RVA LLC. 2025. "The Underappreciated Need to Enable AI and Data Center Growth." Washington, DC: Fiber Broadband Association.

Selsky, Andrew. 2022. "Oregon City Drops Fight to Keep Google Water Use Private." *Associated Press*, December 15.

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

Stansbury, Martin, Kelly Marchese, Kate Hardin, and Carolyn Amon. 2025. "Can U.S. infrastructure Keep Up With the AI economy?" Deloitte Insights, June 24.

U.S. Department of Energy, Secretary of Energy Advisory Board (USDOE SEAB). 2024. "Recommendations on Powering Artificial Intelligence and Data Center Infrastructure."

U.S. Energy Information Administration. 2024. Use of Energy Explained: Electricity Use in Homes.

U.S. Environmental Protection Agency (USEPA). 2024. Greenhouse Gas Equivalencies Calculator.

Virginia Joint Legislative Audit and Review Commission (VJLARC). 2024. ***Virginia Data Center Study: Electric Infrastructure and Customer Rate Impacts***.

Van Geet, Otto, and David Sickinger. 2024. ***Best Practices for Energy-Efficient Data Center Design***. Washington, D.C.: U.S. Department of Energy Federal Energy Management Program.

Walker, Willy. 2024. "The Future of Real Estate Is Digital: How Data Centers and 5G Are Shaping the Next Generation of Infrastructure." *Walker & Dunlop Market Trends*, October 30.

Zhang, Mary. 2023. "Types of Data Centers: Enterprise, Colocation, Hyperscale." *Dgtl Infra*, November 26.

ZincFive. 2024. "Data Center Energy Storage Industry Insights Report 2024."

American Planning Association

Creating Great Communities for All

ZONING PRACTICE OCTOBER 2025 | VOL. 42, NO. 10. Zoning Practice (ISSN 1548–0135) is a monthly publication of the American Planning Association. Joel Albizo, FASAE, CAE, Chief Executive Officer; Petra Hurtado, PHD, Chief Foresight and Knowledge Officer; David Morley, AICP, Editor. Subscriptions are available for $65 (individuals) and $120 (organizations).

BROOKINGS

RESEARCH

How Advanced Industrial Zones can help states thrive in the electro-industrial economy

Jake Higdon and Joseph Parilla

July 9, 2026

In an era of rising energy demand ↗, global industrial competition ↗, and growing risks of climate change, the United States needs to manufacture, generate, transmit, and deploy abundant clean electricity at scale and speed. Yet this imperative is running headfirst into a series of workforce, permitting, and governance bottlenecks that, as a recent Brookings report argues, will only be overcome when U.S. state governments better align their economic development and energy policies. Doing so will not only enable electrification, but also help states capture ↗ the economic benefits associated with what energy experts ↗, investors ↗, and economic analysts ↗ have termed the "electro-industrial economy."

By "electro-industrial economy," we refer to the increasingly integrated system of clean energy, electricity infrastructure, semiconductors, and data centers, whose growth is becoming mutually reinforcing. Advanced industries have always powered the U.S. innovation and export economy, but the country is entering a new phase. The demands of the electro-industrial era put greater value on the physical assets of places: their electricity grid, land, water, and other infrastructure, in addition to their base of workers. This strain, along with broader public anxiety ↗ about artificial intelligence, has expressed itself in the recent backlash ↗ to the data center buildout. Americans want investment and jobs, especially in places experiencing economic distress—but they are clearly saying, "Not like this."

As we approach a major midterm election, U.S. states are seeing unprecedented investment ↗ that, if channeled properly, could reduce costs for consumers ↗, mitigate supply chain risks ↗, decarbonize ↗ our energy system, and deliver more jobs and economic opportunity ↗ to more places. But elected officials can only channel investments if they have clarity on what sort of positive-sum model their voters can support. This report explores how a new tool—Advanced Industrial Zones—can help policymakers plan and govern this big electro-industrial buildout to secure private sector investment, grow good jobs, and create pathways to economic mobility in local communities.

Why the electro-industrial economy matters and how it is playing out on the ground

Brookings has long argued that advanced industries are the primary source of U.S. innovation, productivity, and export growth. Today, advanced industries are entering a new phase. In the 2010s, experts were portending a "fourth industrial revolution" ↗ revolving around breakthroughs in digitization and robotics. In the early 2020s, the central driver of electro-industrial investments was the shift to clean energies such as batteries, electric vehicles, heat pumps, and solar panels, which drove $275 billion in domestic investment in 2025 ↗ and are increasingly cheaper than fossil-fuel-based alternatives. Now, hulking artificial intelligence and cloud data centers are the avatars of U.S. industry. By 2030, tech giants are poised to invest something on the order of $2 trillion in U.S. compute infrastructure ↗.

The electro-industrial era is global in scope, but it will be won through local advantages. While companies compete in different market segments and organize around different global value chains, their competitiveness ultimately rests on localized ecosystems of innovation, talent, land, power, and logistics infrastructure. No single region possesses each of these assets in equal measure, nor will every region compete in the same segments of the electro-industrial economy. Different combinations of regional strengths create different opportunities. As a result, the electro-industrial transition is likely to unfold through a variety of place-based development pathways. We outline three below:

- **Innovation-production corridors** build on co-located concentrations of research institutions, startup ecosystems, and specialized talent and expertise to move technologies from discovery to production. These regions leverage national labs, universities, and innovation districts to attract manufacturing tied to frontier technologies. In the electro-industrial era, notable examples include Fremont, Calif. ↗, home to Tesla and a variety of climate tech manufacturers benefiting from Bay Area proximity; and the Illinois Quantum and Microelectronics Park ↗ (IQMP), which builds on innovation assets in the Greater Chicago region, such as Argonne National Laboratory's quantum computing expertise. This powerful innovation-production link is why Fervo Energy built its first commercial operation next door to the Department of Energy's FORGE geothermal testbed ↗ in Utah, and why the Trump administration is now looking to co-locate AI infrastructure on national lab property ↗.

- **Reindustrialization anchors** emerge in regions with a strong legacy of manufacturing, plentiful land, and a skilled production workforce. These places are repositioning themselves as sites for large-scale production in sectors such as semiconductors, batteries, and clean energy manufacturing. For instance, Weirton, W.Va., had been redeveloping its old steel mill site ↗ since 2017—years before federal funding and tax credits made it possible to build the now famous Form Energy factory. Megaprojects such as Micron's $100 billion semiconductor investment in Upstate New York or Hanwha QCells' $2.5 billion solar supply chain ↗ in and around Dalton, Ga., also exemplify this trend.

- Sustainable energy-compute hubs leverage rapid investment in compute infrastructure such as data centers to achieve low-carbon or regenerative infrastructure goals. As our Brookings colleagues have shown, these projects can directly add clean energy capacity to the system, improve water quality, cover the costs of grid infrastructure such as transmission lines, and accelerate development of new energy technologies. Examples of this include Google's partnership with Xcel Energy to leverage Form Energy's iron-air batteries at its Minnesota data center, or Google's completed project to upgrade water infrastructure ↗ in The Dalles, Ore. This archetype is primarily emerging around data centers, but can also be applied across other large industrial loads, such as Antora's new 5 gigawatt per hour thermal battery deployment to deliver 24/7 power to a bioprocessing facility in South Dakota. Sustainable energy-compute hubs offer to not only decarbonize local infrastructure, but also drive down clean technology costs and solidify domestic supply chains.

This simple typology illustrates that the electro-industrial era is not producing a single geography of growth, but a differentiated map shaped by how regions align physical infrastructure, energy systems, and innovation capacity.

Rapid electro-industrial buildout is scrambling the economic development playbook

Traditionally, regions have been able to make these industry bets over long time horizons. Local economic development practitioners think of building infrastructure, attracting manufacturers, and developing the local workforce in terms of years and decades. These processes provide time for proactive community organizing, infrastructure upgrades, and more. The types of projects that were being announced just a few years ago, such as semiconductor and solar manufacturing hubs, benefitted from this playbook.

Now, electro-industrial investments, especially data centers, are unfolding in weeks and months. A single data center campus may offer an order of magnitude more investment than some communities have seen in decades. This offers a tantalizing opportunity: Economic developers could bypass the years of effort and cost necessary to even have a chance at landing a factory.

Yet there is a structural misalignment between the scale of investment and the absorptive capacity of the communities on the other side of the table. Utilities are scrambling to adapt, public opinion ↗ is turning on data centers, and communities are establishing anti-development veto points ↗ with potential ripple effects across the economy. As local officials and economic developers evaluate the landscape of potential electro-industrial investments, it’s hard to make sense of what represents a generational reinvestment opportunity—and what’s a devil’s bargain.

We observe four major governance challenges arising from the electro-industrial buildout:

- **Uncertain economic development value.** Not all electro-industrial investments create the same degree of local economic value. Production facilities—e.g., battery plants, semiconductor fabs, electric vehicle facilities—can anchor larger localized supply chains that create more jobs and longer-term advantages for regions. For these reasons, a recent analysis ↗ showed that the majority of tax incentive deals for clean energy projects delivered positive returns to the surrounding areas. By contrast, the long-term local value of data centers is much less settled. Data centers do create jobs, but the numbers can be overstated and are complicated by the fact that data centers come in different forms (e.g., cloud versus AI training loads). Not only are data centers different from manufacturing ↗, but they also do not interface with the local economy in the same ways that more traditional infrastructure, such as rail or broadband, enables local industry. Recent walk-backs of tax incentives ↗ for data centers in places such as Virginia suggest a recalibration of the cost-benefit analysis for economic developers.

- **Information and power imbalance.** Decisionmakers grappling with electro-industrial investments are operating in an asymmetric and rapidly shifting information environment. There is a structural mismatch between the scale, sophistication, and capacity of large tech and industrial firms and local governments in the places they are looking to site their facilities; firms often use this to their advantage by making claims about their technology that are hard to verify, pitting jurisdictions against each other to secure better incentives, or sweet-talking officials. To make matters worse, tech giants often use third parties or special-purpose vehicles to manage data center development, and have leveraged nondisclosure agreements ↗ to restrict public knowledge of these projects. This undermines trust between firms, officials, and the public, and hampers future accountability.

- **Public backlash.** A deep, bipartisan backlash to data center development has emerged across the country, with implications not just for data centers but for the broader electro-industrial market. For those concerned about the macroeconomic and social implications of artificial intelligence controlled by a handful of powerful firms, data centers are one of the few places where the average citizen can throw sand in the gears. Others are concerned about local quality-of-life issues, including noise and light pollution, electricity prices, and water quality. This backlash is creating immense pressure on local officials and even prompting recall petitions ↗. Policy changes driven by public pressure, deepening not-in-my-backyard (NIMBY) sentiment, and risk aversion on the part of public officials will make it more difficult for all types of electro-industrial projects to move forward.

- **Competition over scarce resources.** Almost all electro-industrial investments require the same basic inputs: electricity, land, permits, and workforce. Local supply of these real assets is relatively inelastic. A single factory or data center will often require a new substation, grid interconnection, and additional power capacity on that grid. Industrial lands are in short supply, particularly after the investment boom of the past few years. These dynamics pit investments against one another; stretch local administrators, economic developers, and educational institutions; and incentivize jurisdictions to cut taxes (or corners) in a race to the bottom.

Given these challenges, one clear governance option is refusal, which is understandable: Data centers, factories, and energy infrastructure come with tradeoffs. Some states and regions may calculate that the costs are not worth bearing or are too opaque.

However, sidestepping electro-industrial investments also comes with a major opportunity cost—at a local scale, across regional economies, and for national resilience, competitiveness, and decarbonization. For those interested in balancing these tradeoffs and finding a path forward, it's important to define the positive-sum model for planning the electro-industrial transition. Public authorities, community and labor stakeholders, and local businesses need a vision for where they want to go and a clear sense of place. Only then can they negotiate for projects that deliver aligned economic development outcomes, enhance broader regional infrastructure, and hold projects accountable to their goals.

A proactive planning model for the buildout: Advanced Industrial Zones

One way to plan and organize these electro-industrial investments is through the designation of Advanced Industrial Zones. Many of the obstacles to building new industries in America today—from a convoluted permitting process to inadequate grid infrastructure—require a localized solution. Many of the risks—from environmental degradation to noise pollution—are due to unplanned, Wild West-style development. Designated zones provide a governance framework for navigating these challenges in specific places.

This type of planning is not new in the world of economic development, with past examples ranging from traditional industrial districts to federal tax incentives such as

Opportunity Zones. However, the electro-industrial moment requires a new type of zone—one that focuses on the strategic sectors of the modern, electrified industrial economy; works to agglomerate the real assets that they require; and delivers the types of jobs and benefits communities demand. Fortunately, we have many existing models across the U.S. ↗ and abroad that offer lessons for the designation and implementation of these zones.

Designation and governance

Advanced Industrial Zones should be community-led and state-enabled. The most effective zone policies will establish a sort of long-term governance compact between local and state authorities.

Who should take the lead in identifying the right geographies for these zones? It will be important for the state to establish a framework and process for zone designation that involves community approval. In many cases, the state may have already identified one or more initial zones of strategic importance, such as an old military base, port community, or known innovation cluster. The state may also wish to establish a bottom-up nomination process for additional zones. Either way, as locations are being considered for special designation as an Advanced Industrial Zone, there is a unique opportunity for proactive planning, education, and community engagement that is not always present in industrial development. Local economic development organizations are well positioned to facilitate this conversation given their role providing the connective tissue between community stakeholders, municipal and state governments, workforce institutions, and industry.

In the designation process, states and localities should consider existing assets and land use, the types of electro-industrial activities that are geographically suitable and desirable, and the expectations for community and worker benefits. These expectations could be codified in an "Advanced Industrial Zone Development Plan" that state and local officials approve. This is a better model for community engagement than the reactive, project-specific haggling that comes after a project has been announced. There's a widespread but faulty impression that zones bypass local preferences and process; however, with the right design, they can be a tool for community organizing.

Yet after those conditions are agreed upon and the zone is designated, officials must be able to move with independence and clarity of purpose to develop sites, deliver timely permits, and offer certainty to employers—effectively creating a landing pad for rapid investment. In successful projects such as South Korea's Ulsan Eco-Industrial Park ↗ or the Devens Regional Enterprise Zone in Massachusetts ↗, state or national governments delegated a certain set of streamlined planning, permitting, and investment tools to specialized commissions or public development corporations that administer the zone. These entities would still be subject to ongoing public oversight and responsive to the guardrails set forth in the initial designation and planning process.

Strategic land use planning

It may be counterintuitive, but for those legitimately concerned about harms from industrial activity, Advanced Industrial Zones can help regions manage and mitigate the risks of industrial development by promoting efficient and thoughtful land use. Investments in clean energy generation and battery manufacturing are incredibly important tools for broader pollution reduction and regional economic activity, but many of their would-be supporters remain torn over hyperlocal environmental and social impacts. By clustering large electricity loads and shared infrastructure, zones enable states, communities, utilities, and conservationists to mitigate sprawl and optimize service delivery—and that's even before accounting for the benefits of clustering for innovation in manufacturing.

State and utility infrastructure investment

In return for this proactive designation and planning process, states or other entities should offer preferential infrastructure investment within the zones to attract manufacturers. This can include site readiness or brownfield remediation grants, prioritized grid upgrades, and onsite power generation. These investments focus public dollars on the core physical inputs to electro-industrial projects, allowing communities to invest in a diversified industrial ecosystem that is not tied to any specific firm.

By contrast, approaches that tend to provide bespoke cash incentives for business attraction can be effective, but also highly risky. Such incentives—and the industrial infrastructure they pay for—all become the property of a single project. Recent analysis from the Upjohn Institute suggests that state and local incentives for clean energy and manufacturing projects have a generally positive return on investment ↗, but noncash services such as infrastructure upgrades tend to outperform cash incentive packages.

Targeted financing

Advanced Industrial Zones also offer a way to break free of the budgetary doom loop for rural or deindustrialized communities. The perennial challenge for these communities is that they need to invest in infrastructure, workforce, and state capacity to attract new investments, but do not have the tax base or administrative capacity to support those activities. This is true in legacy industrial regions—such as Pennsylvania's Steel Valley, Southeastern textile hubs, and Pacific Northwest timber towns—but also in rapidly growing exurbs looking to diversify their economies via greenfield industrial development. Resolving this constraint often requires fiscal transfers from other regions or a mechanism to borrow against future growth.

Designated zones with clear economic development plans, governance structures, and community buy-in help provide confidence around innovative financing mechanisms such as tax increment financing (which is what Utah's Inland Port Authority uses to upgrade its industrial districts) and revolving loan funds. It also helps state and federal entities channel financing into specific places with clear, coordinated industrial planning.

Supercharging private capital

Finally, Advanced Industrial Zones can allow communities to strategically tap into a new source of capital: hyperscalers, or the technology companies hungry for land and power to rapidly roll-out billions in investments for new in AI and cloud computing infrastructure. It also simply makes sense to site data centers within industrial parks rather than in greenfield developments.

One can imagine a scenario in which a zone courts a data center developer who is asked to directly pay for or underwrite the larger set of investments necessary for a regional manufacturing hub, such as transmission upgrades, brownfield remediation, site prep, water infrastructure, and workforce development. Regions could effectively leverage hyperscalers' infrastructure investments as a "slingshot" for broader industrialization or energy transition goals. To pull this off, economic developers will need to be clear-eyed and strategic about the types of industries they want to grow and the infrastructure they will need in specific places. Advanced Industrial Zones can organize that vision.

Building new opportunities from old place-based assets

The various agendas of hyperscalers, clean energy advocates, geopolitical and national security hawks, and economic developers have coalesced around the same physical imperative: to build vast amounts of new electricity and manufacturing capacity across the country. Like waves of different frequencies overlapping, these trends have fueled an infrastructure super-cycle. Whether this generates constructive or destructive waves remains to be seen; so far, it has run headfirst into the same set of physical constraints across the American landscape.

The good news is that this electro-industrial transition is revealing the value of place-based assets—including power infrastructure, industrial land, and a local blue collar workforce—that had fallen out of favor in the 21st century software economy. This could mean new economic opportunity for a broader swath of communities and workers, as well as a once-in-a-generation window to revitalize and clean up our antiquated energy system. States and communities that move quickly to adopt new planning and governance structures—such as Advanced Industrial Zones that organize community priorities, streamline outdated procedures, and channel public sector investments in specific geographies—will be the ones who maximize the potential of this moment.

AUTHORS

Jake Higdon Energy Transition Fellow - California Forward

Joseph Parilla Senior Fellow & Director of Applied Research - Brookings Metro in joe-parilla-14bba69

The Brookings Institution is committed to quality, independence, and impact.
We are supported by a diverse array of funders. In line with our values and policies, each Brookings publication represents the sole views of its author(s).