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.

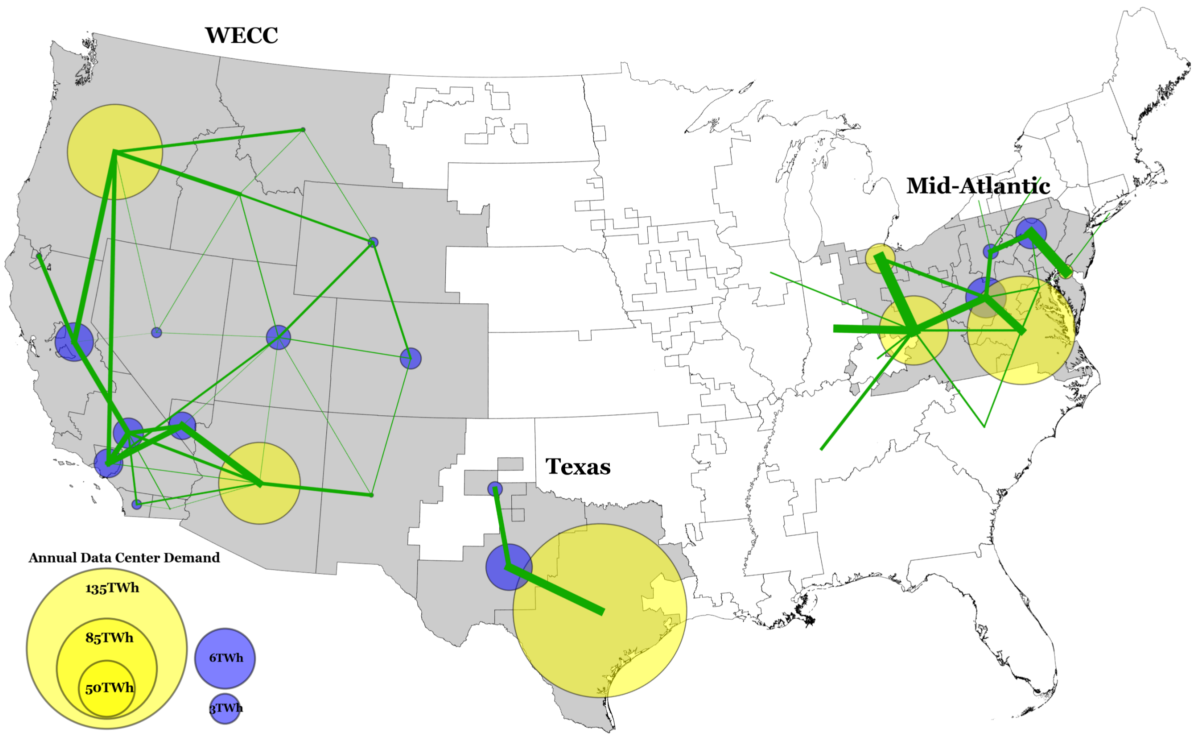

Figure 1: Model Zones with Transmission Lines and Data Center Load. Yellow and Blue circles represent annual data center demand per zone. Green line segments indicate transmission links between zones in a region where the width is scaled to the capacity of the line. Detailed information on load assumptions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

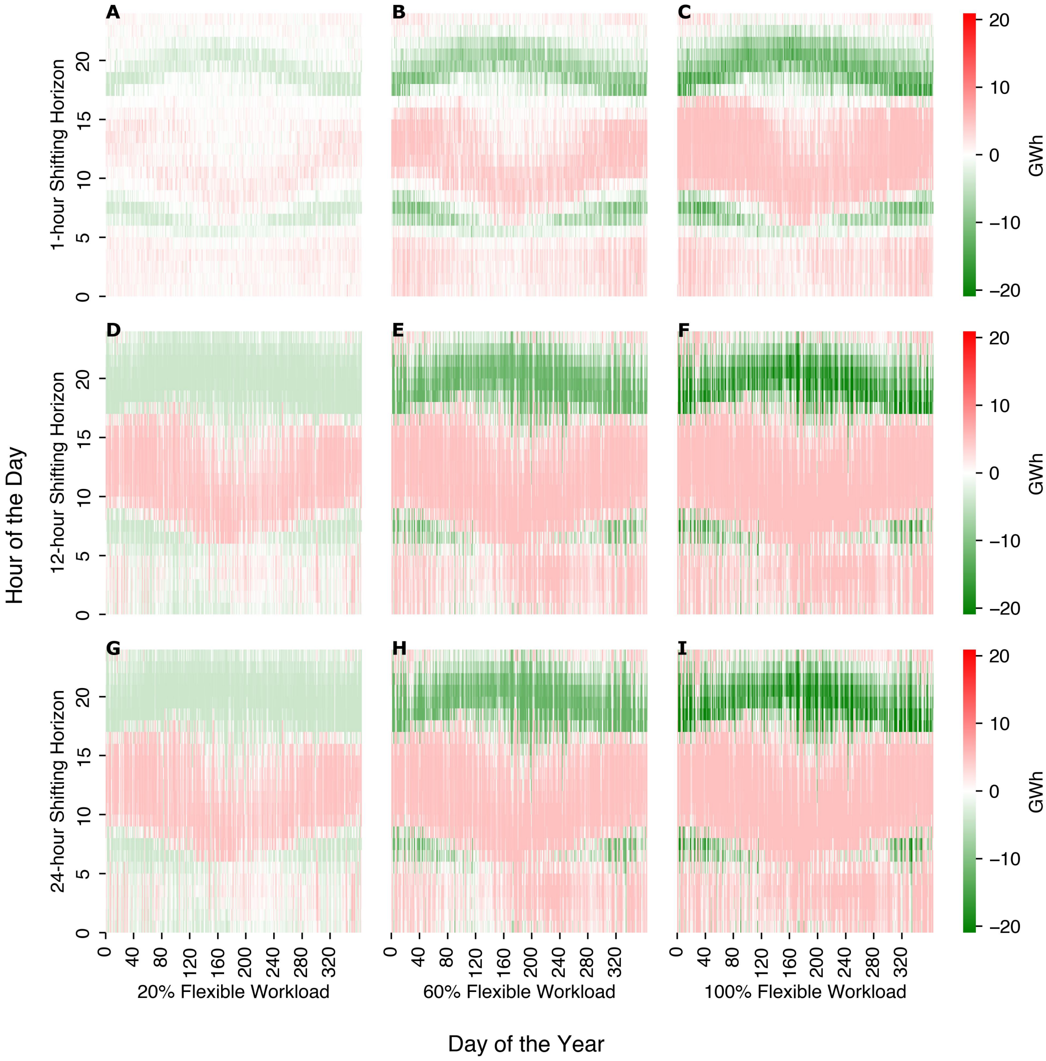

Figure 2: Data Center Load Shifting Operations for the Mid-Atlantic. Each panel displays the net hourly workload shiftted (in GWh) across an entire year, with the x-axis representing days (1–365) and the y-axis representing hours of the day (0–23). Positive values (red) indicate workload shifted into 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

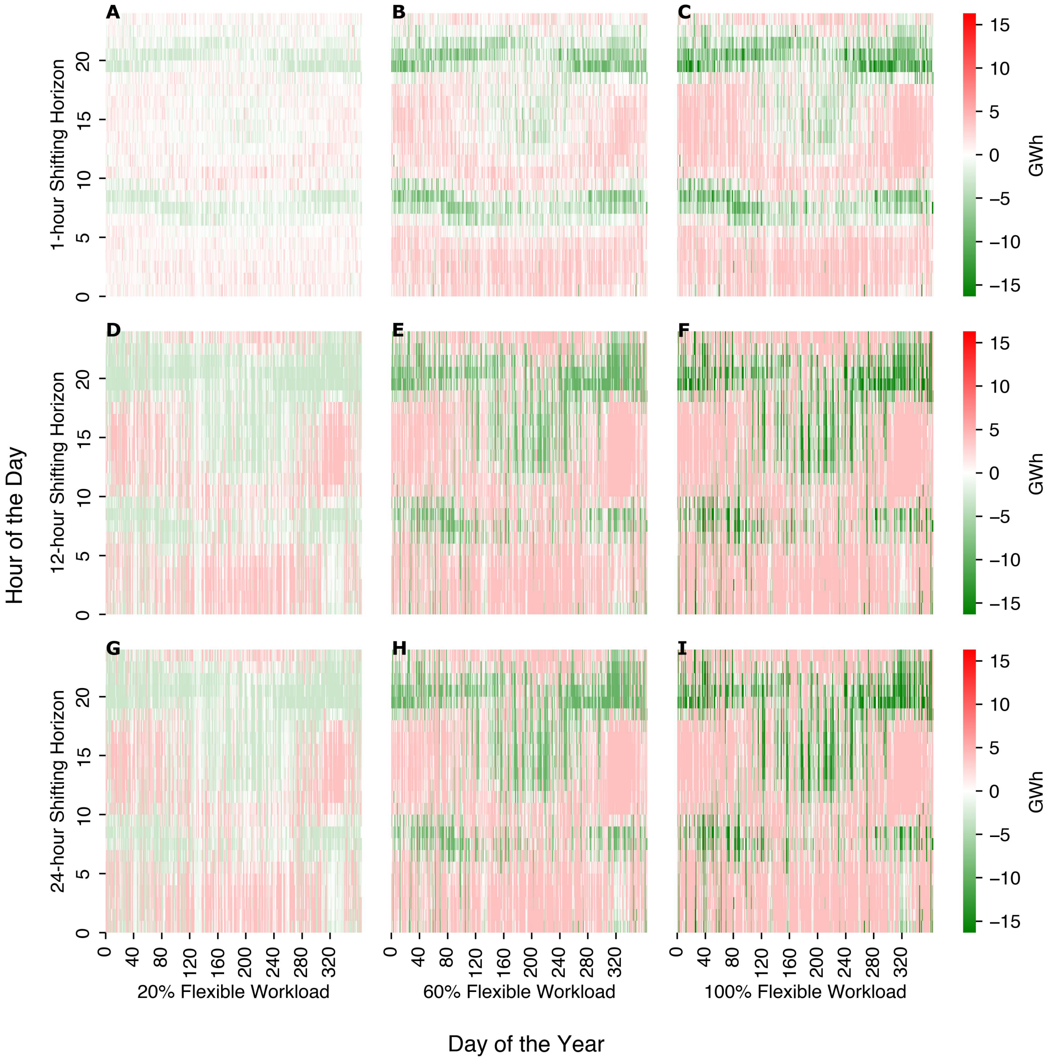

Figure 3: Data Center Load Shifting Operations for Texas. Each panel displays the net hourly workload shiftted (in GWh) across an entire year, with the x-axis representing days (1–365) and the y-axis representing hours of the day (0–23). Positive values (red) indicate workload shifted into a given hour; 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.

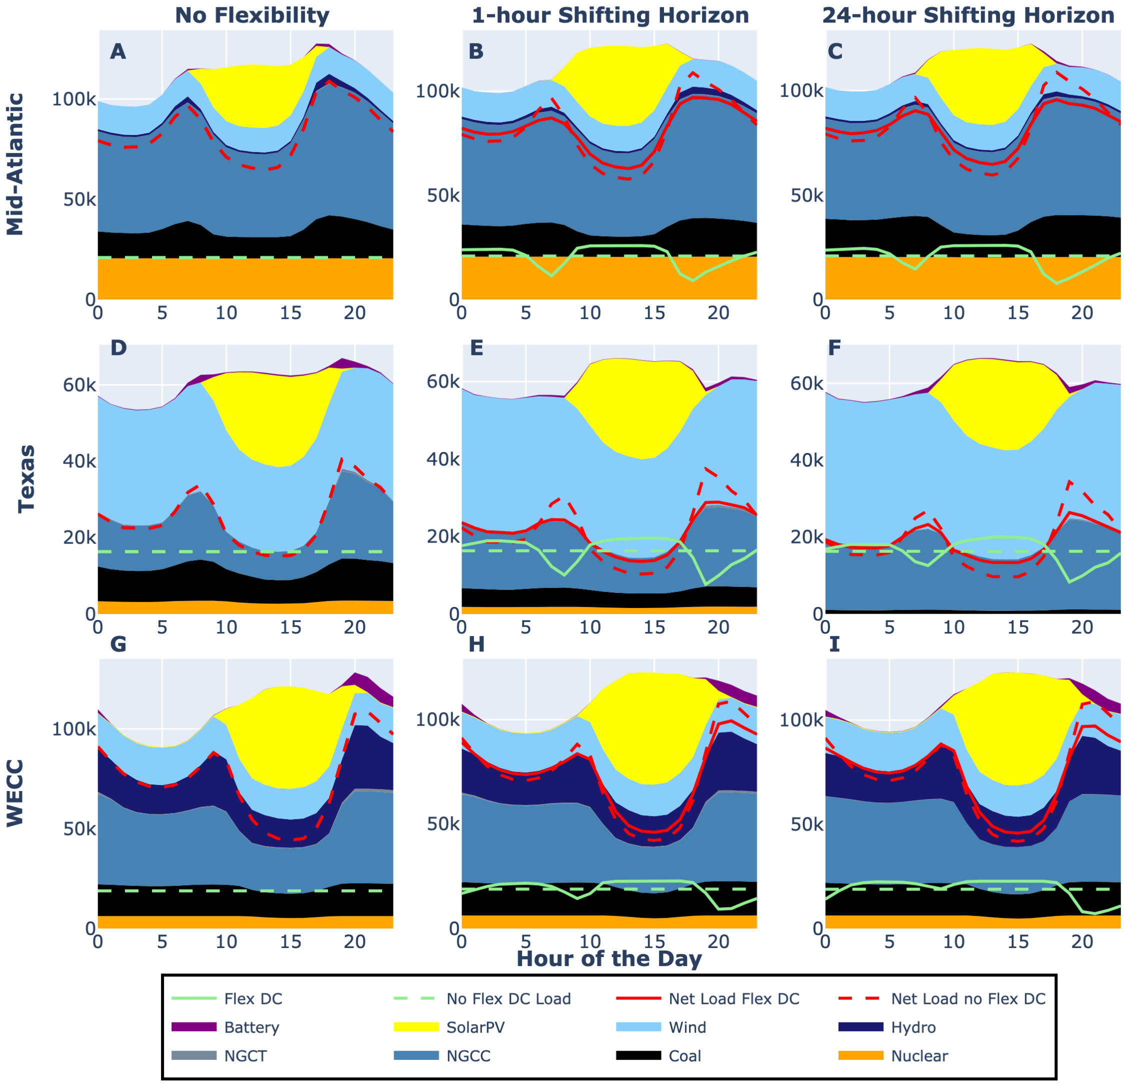

Figure 4: Average Winter Operations per Region. Each subplot shows the average generation per technology, net load, and data center load shifts in MWh per hour of the day during the Winter season. Panels correspond to three levels of operational flexibility: No Flexibility **(A, D, G)** 1-hour **(B, 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

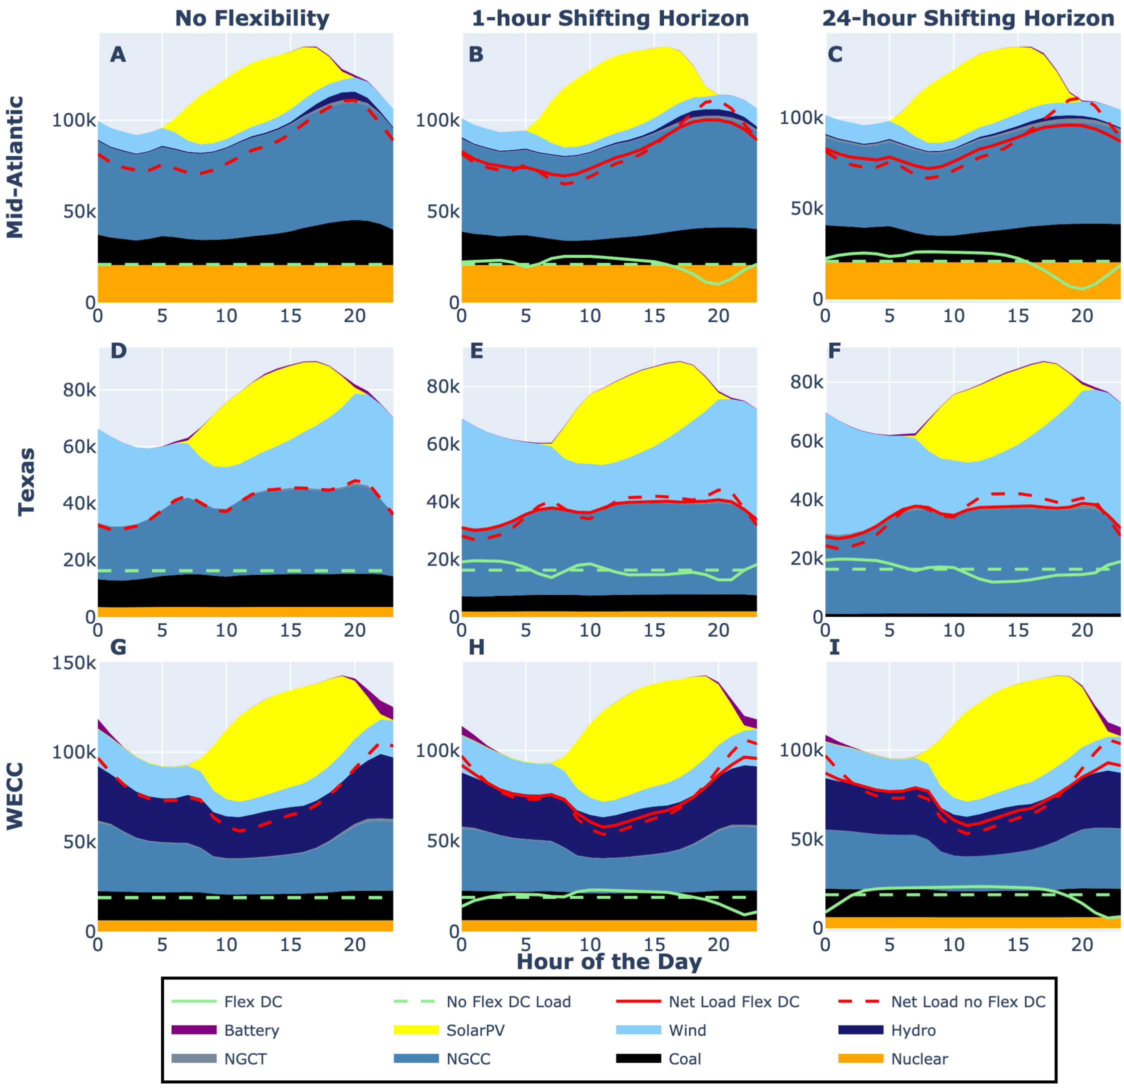

Figure 5: Average Summer Operations per Region. Each subplot shows the average generation per technology, net load, and data center load shifts in MWh per hour of the day during the Summer season. Panels correspond to three levels of operational flexibility: No Flexibility **(A, D, G)** 1-hour **(B, 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

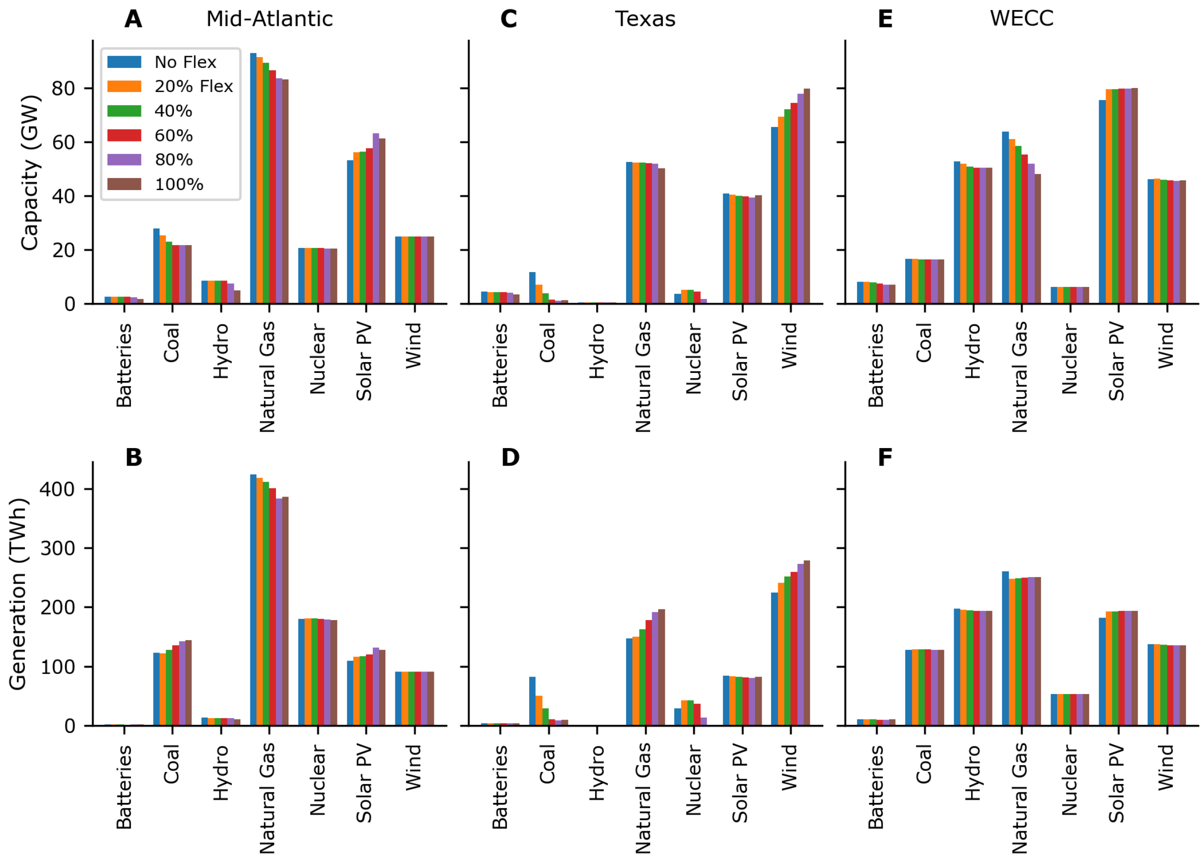

Figure 6: Capacity and Generation per Region with 24-hour shifting horizon. Top row **(A, C, E)** shows total installed capacity by technology, accounting for new investments and retirements. Bottom row **(B, D, F)** shows corresponding total generation by technology. Results are shown across increasing 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

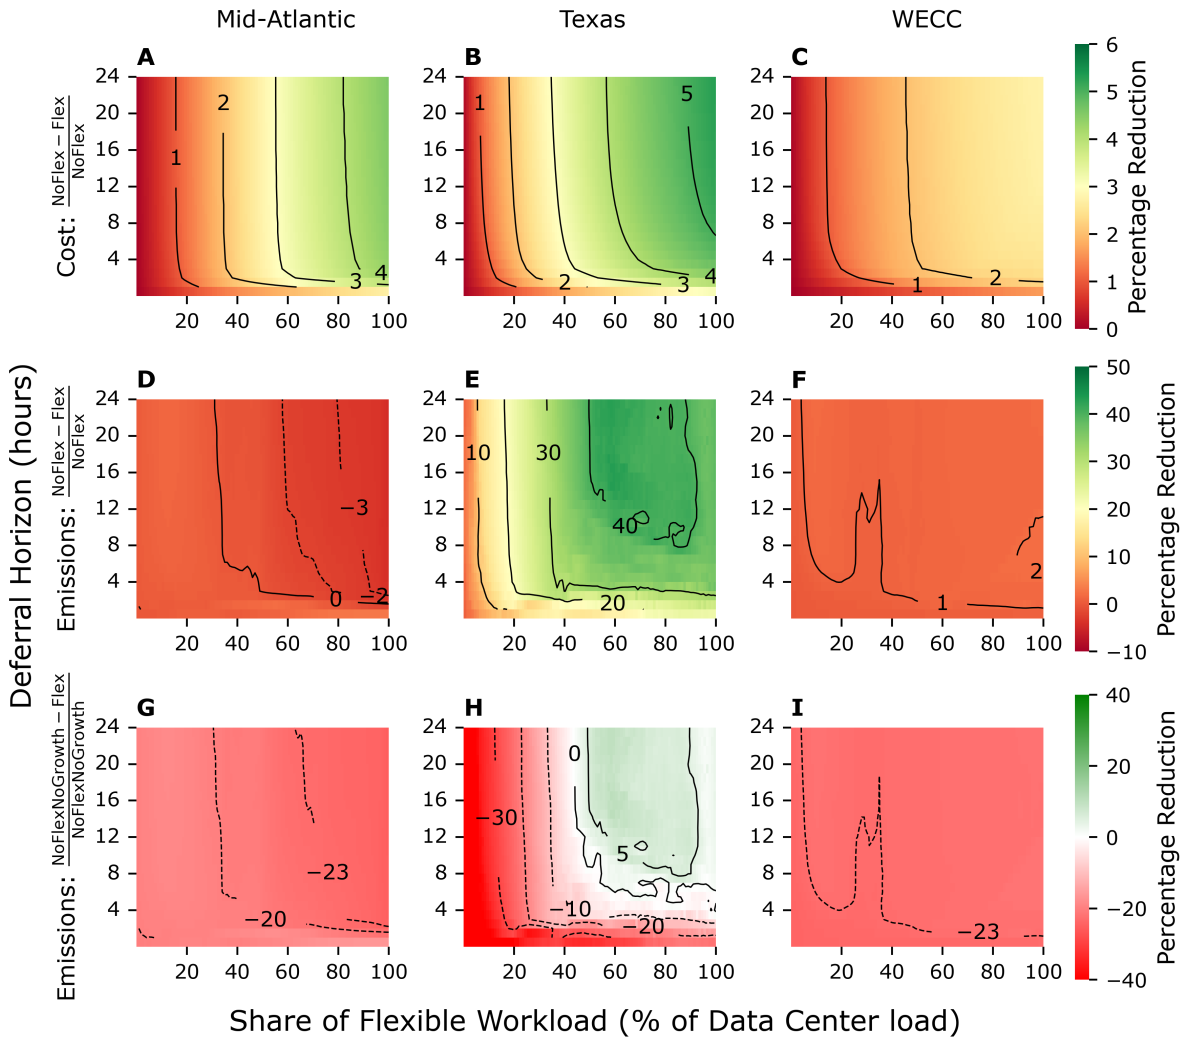

Figure 7: Cost and Emissions Reduction. Top row **(A, B, C)**: Heatmaps show the percentage reduction in total system cost from introducing flexible data centers, relative to a system without flexibility, across combinations of shifting horizon and share of flexible workload. Middle row **(D, E, F)**: 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.

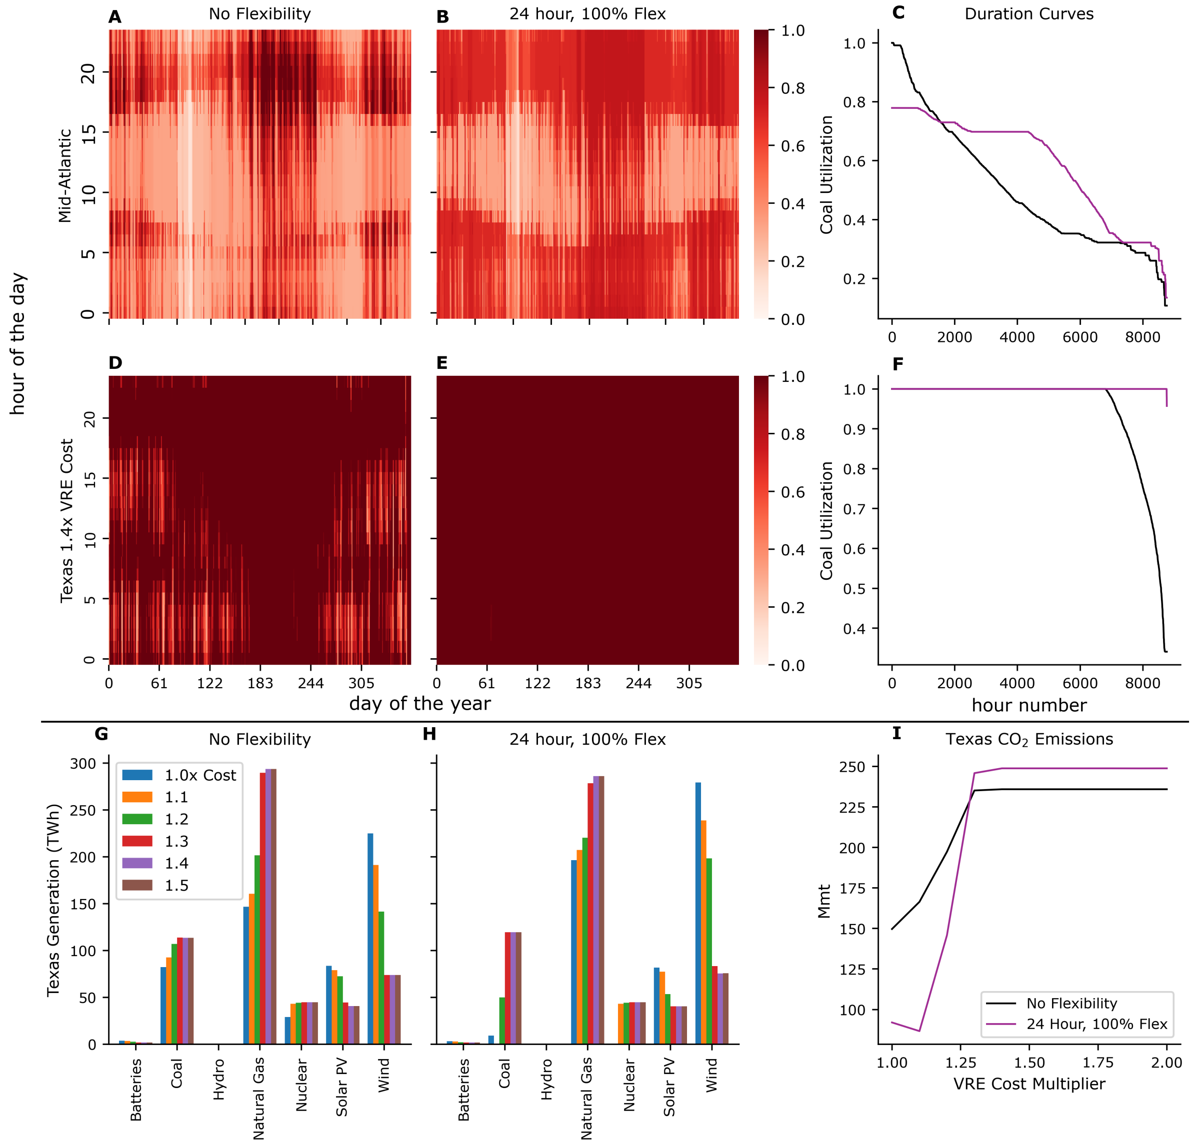

Figure 8: Coal utilization and emissions under varying VRE costs and data center flexibility. Panels **(A-B)** show heatmaps of coal utilization in the Mid-Atlantic under no flexibility and with a 24-hour shifting horizon with 100% flexible workload. Panels **(D-E)** show analogous results for Texas 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	Capital Recovery Period	WACC	Investment Cost	Fixed O&M	Variable O&M
	($/MW)	(years)		($/MW-yr)	($/MW-yr)	($/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

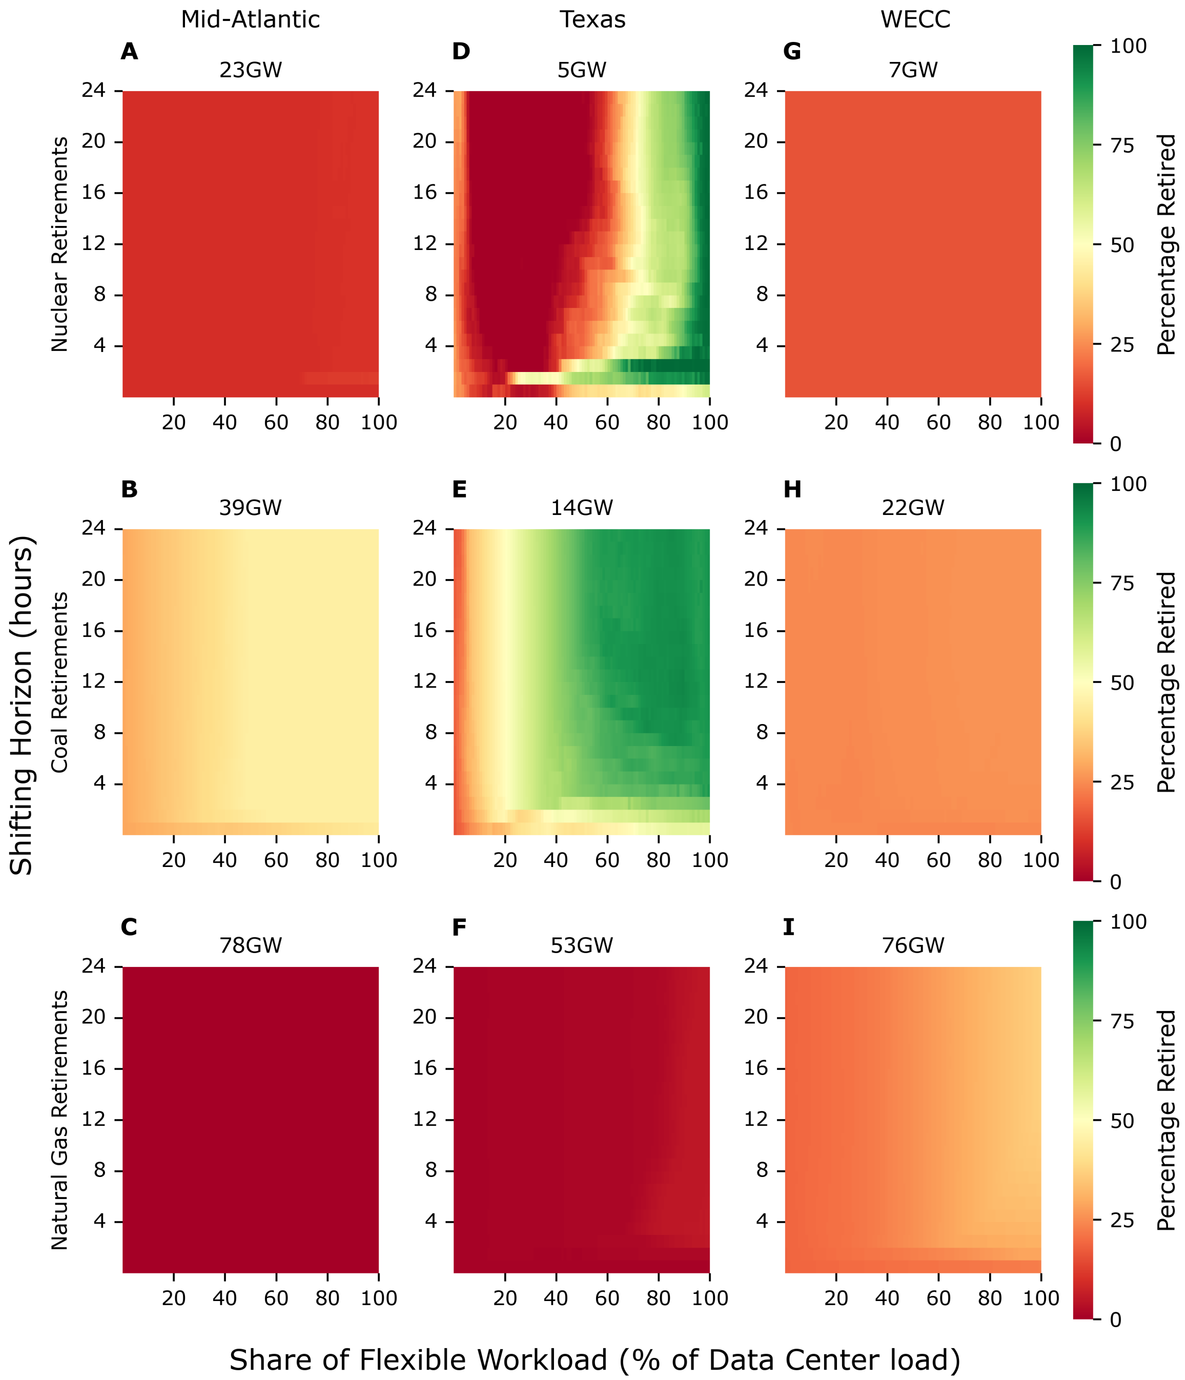

Figure S1: Percentage Retirement per Technology for each Region. Panels show the percentage of existing capacity retired for nuclear **(A, D, G)**, coal **(B, E, H)**, and natural gas **(C, F, I)** across combinations of data center shifting horizon (1 to 24 hours) and flexible workload share (1% to 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.

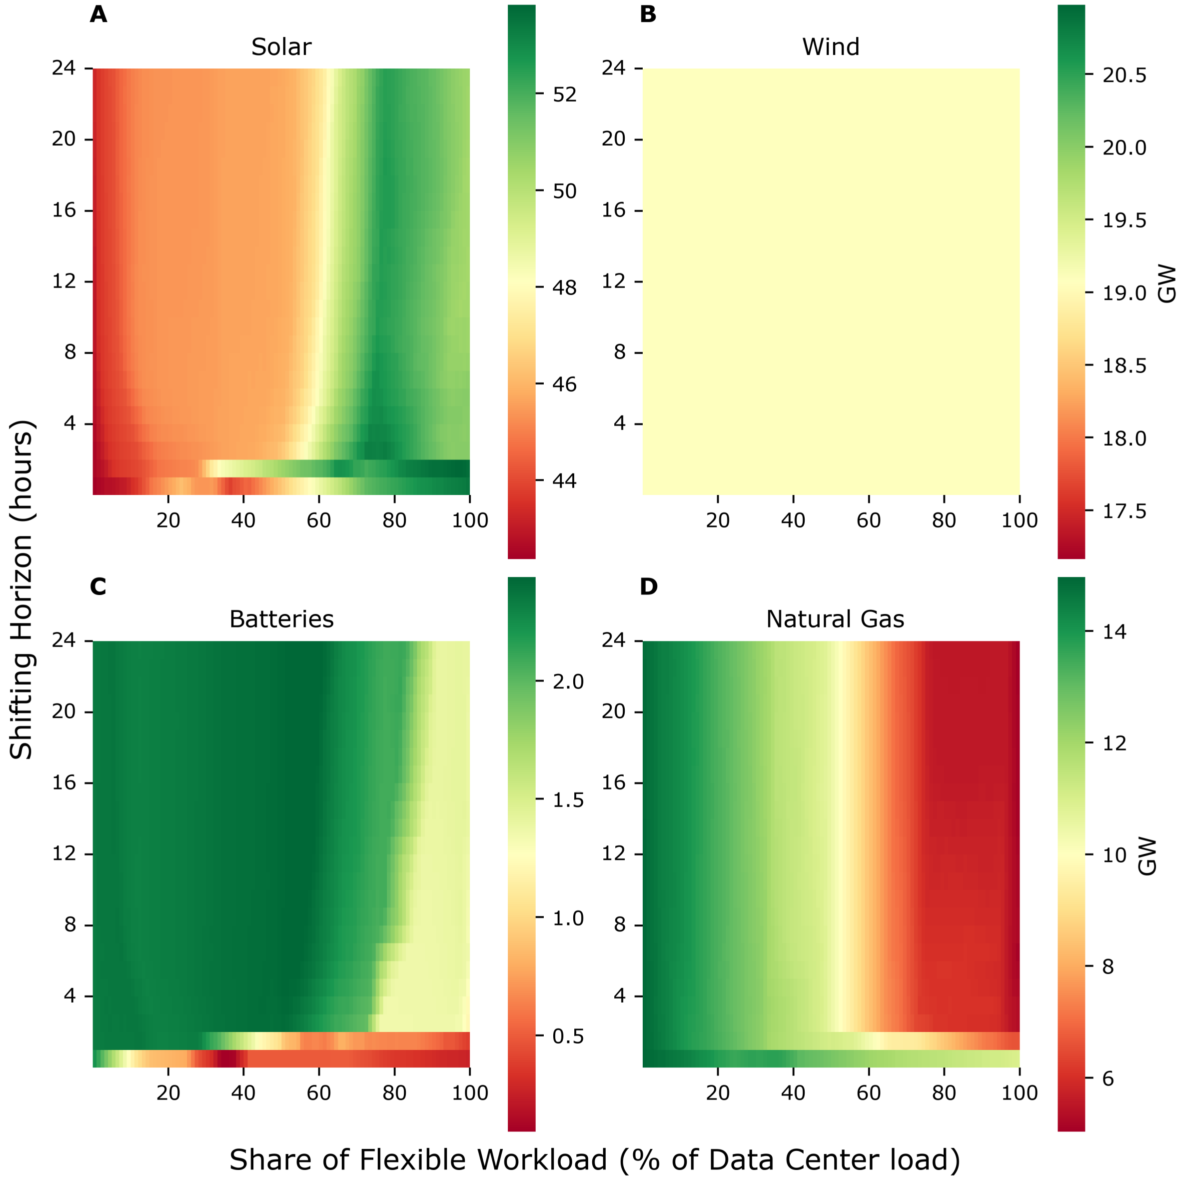

Figure S2: Capacity Investments per Technology in the Mid-Atlantic. Panels show new capacity additions for solar **(A)**, wind **(B)**, batteries **(C)**), and natural gas **(D)** across combinations of data center shifting horizon (1 to 24 hours) and flexible workload share (1% to 100%). Note that color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

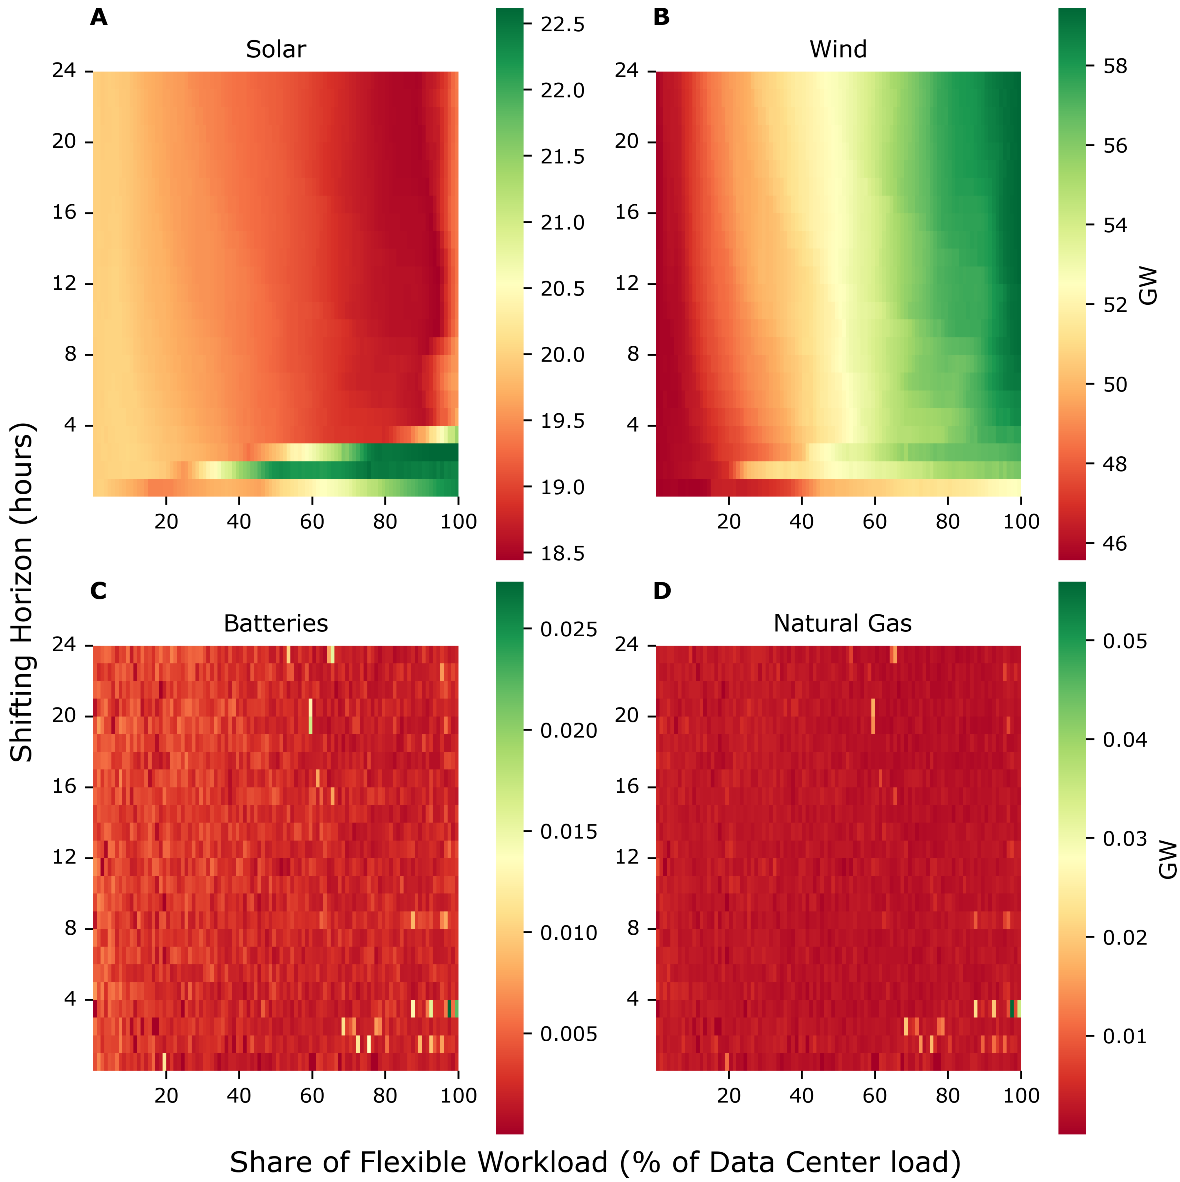

Figure S3: Capacity Investments per Technology in Texas. Panels show new capacity additions for solar **(A)**, wind **(B)**, batteries **(C)**), and natural gas **(D)** across combinations of data center shifting horizon (1 to 24 hours) and flexible workload share (1% to 100%). Note that color scales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

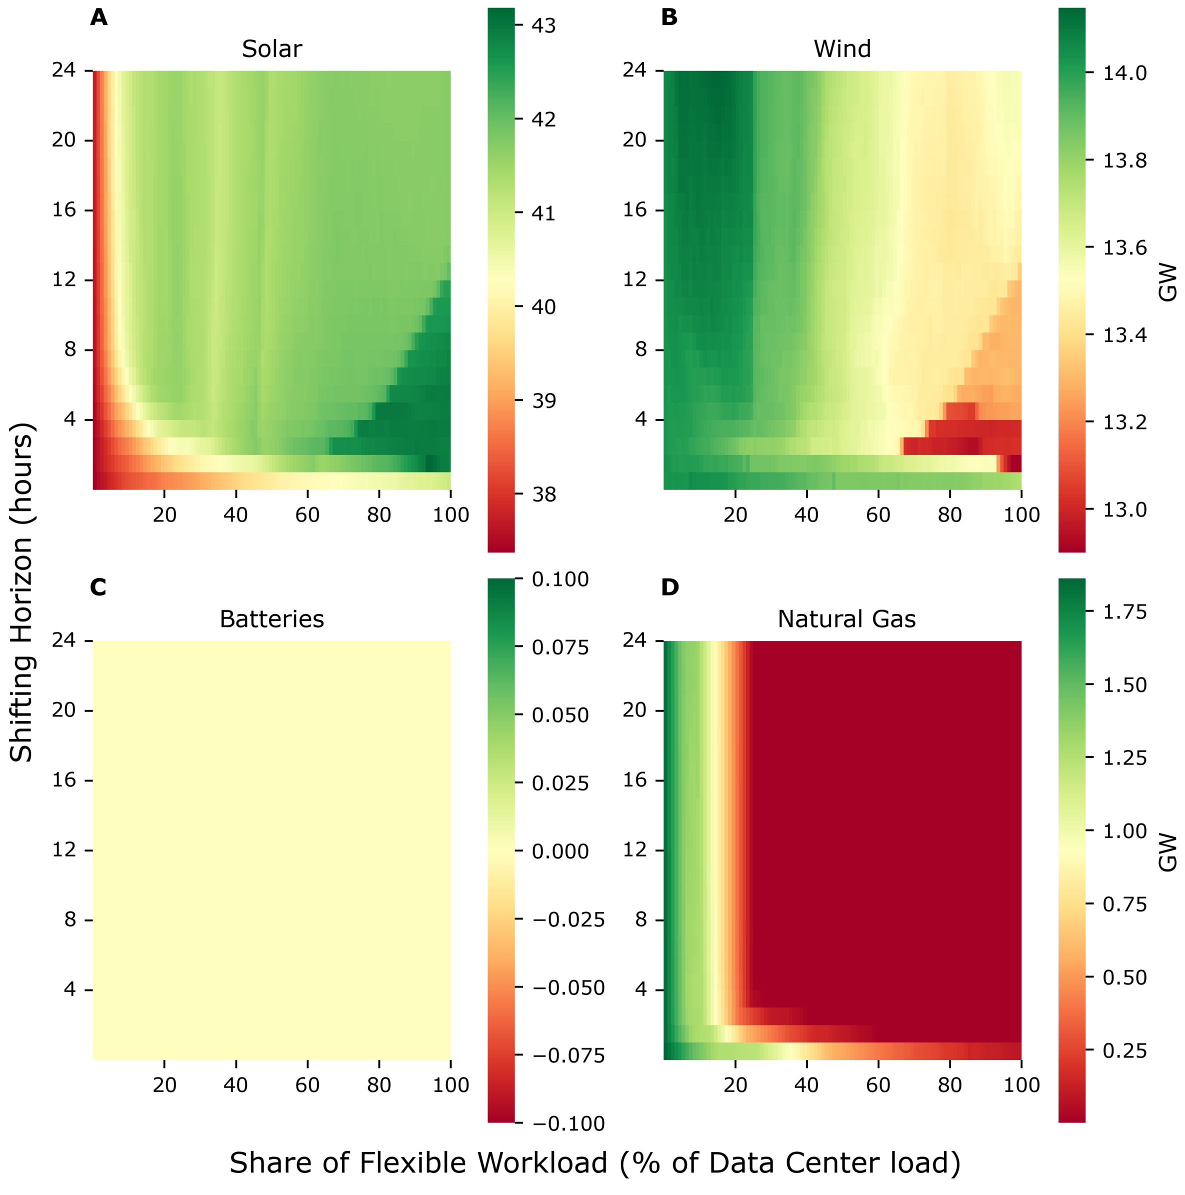

Figure S4: Capacity Investments per Technology in WECC. Panels show new capacity additions for solar **(A)**, wind **(B)**, batteries **(C)**), and natural gas **(D)** across combinations of data center shifting horizon (1 to 24 hours) and flexible workload share (1% to 100%). Note that color scales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

12 WECC Data Center Shifting Operations

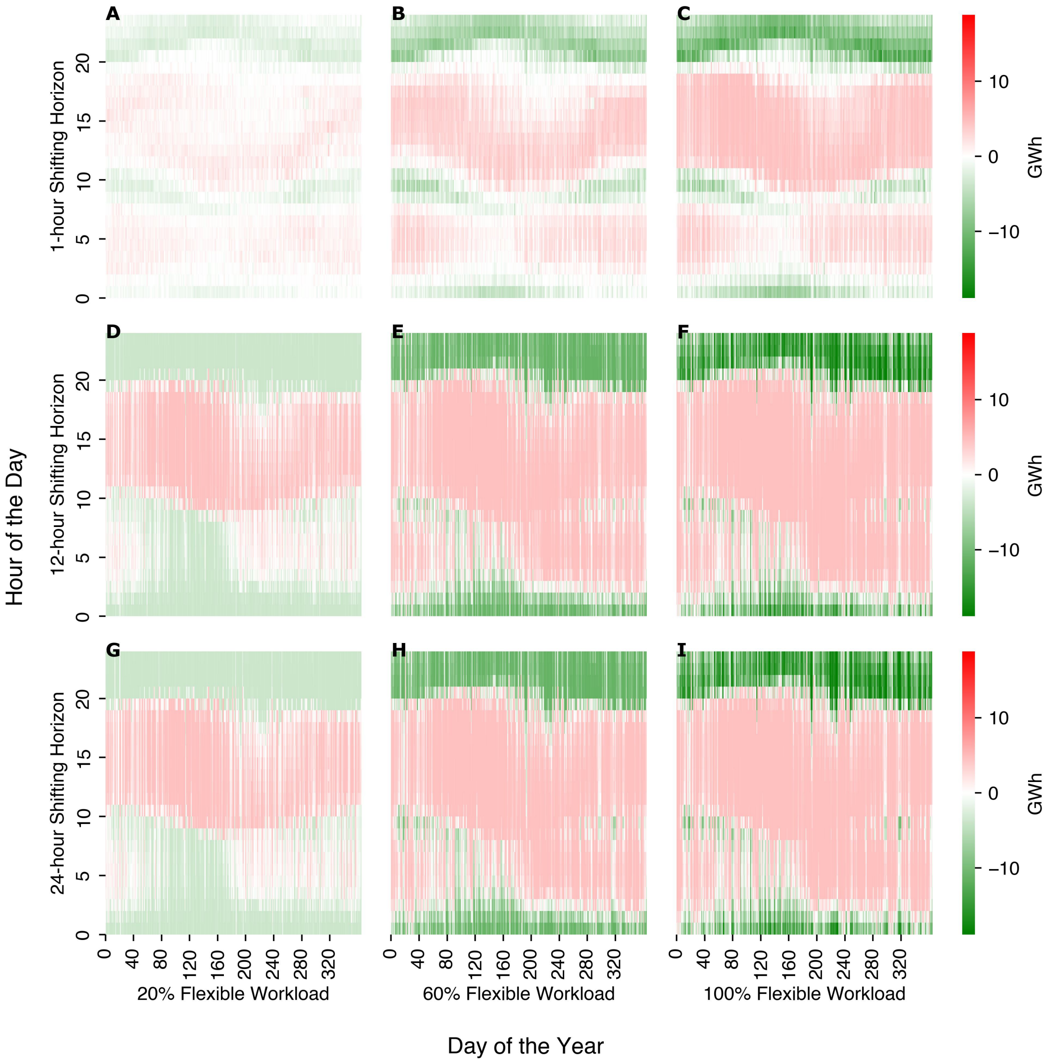

Figure S5: Data Center Load Shifting for WECC. Each panel displays the net hourly workload shifted (in GWh) across an entire year, with the x-axis representing days (1–365) and the y-axis representing hours of the day (0–23). Positive values (red) indicate workload shifted into a given hour; negative 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

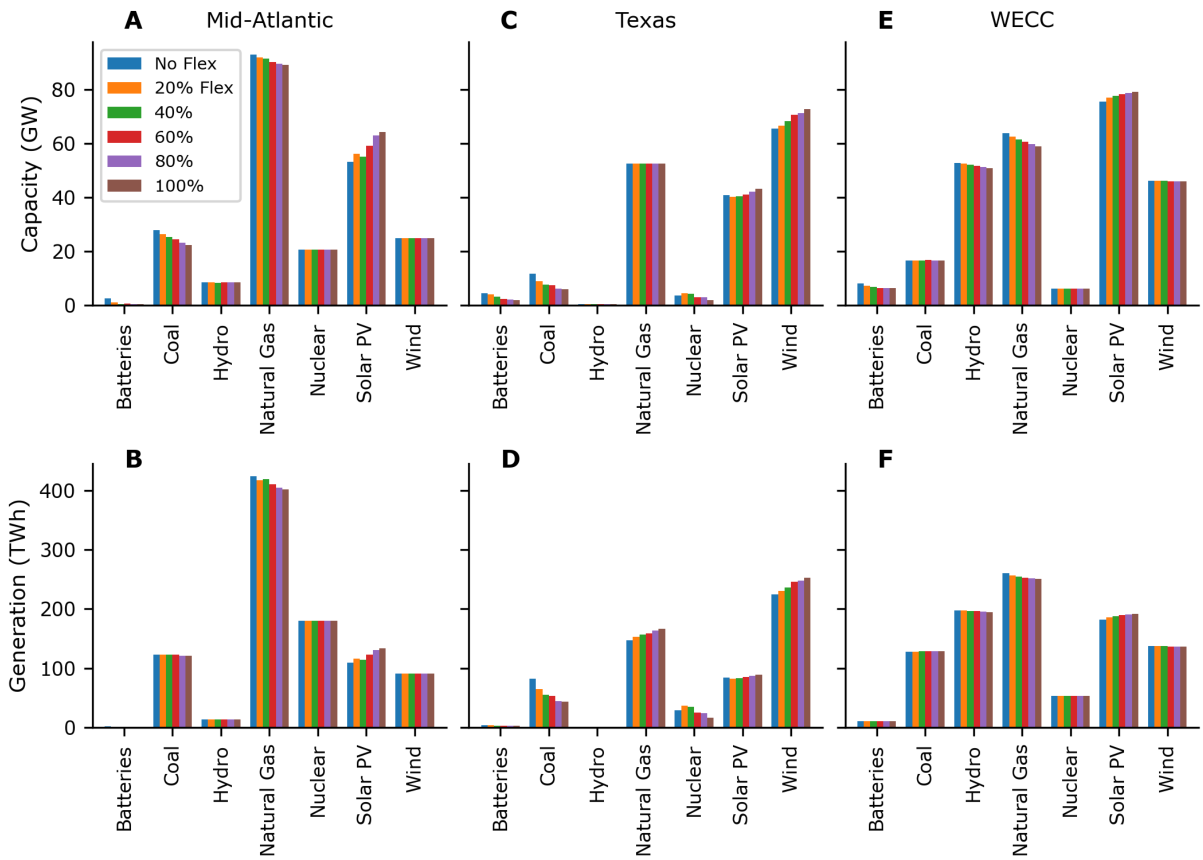

Figure S6: Capacity and Generation per Region with a 1-hour shifting horizon. Top row **(A, C, E)** shows total installed capacity by technology, net of new investments and retirements, for each region. Bottom row **(B, D, F)** presents corresponding total generation by technology. Results are shown 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

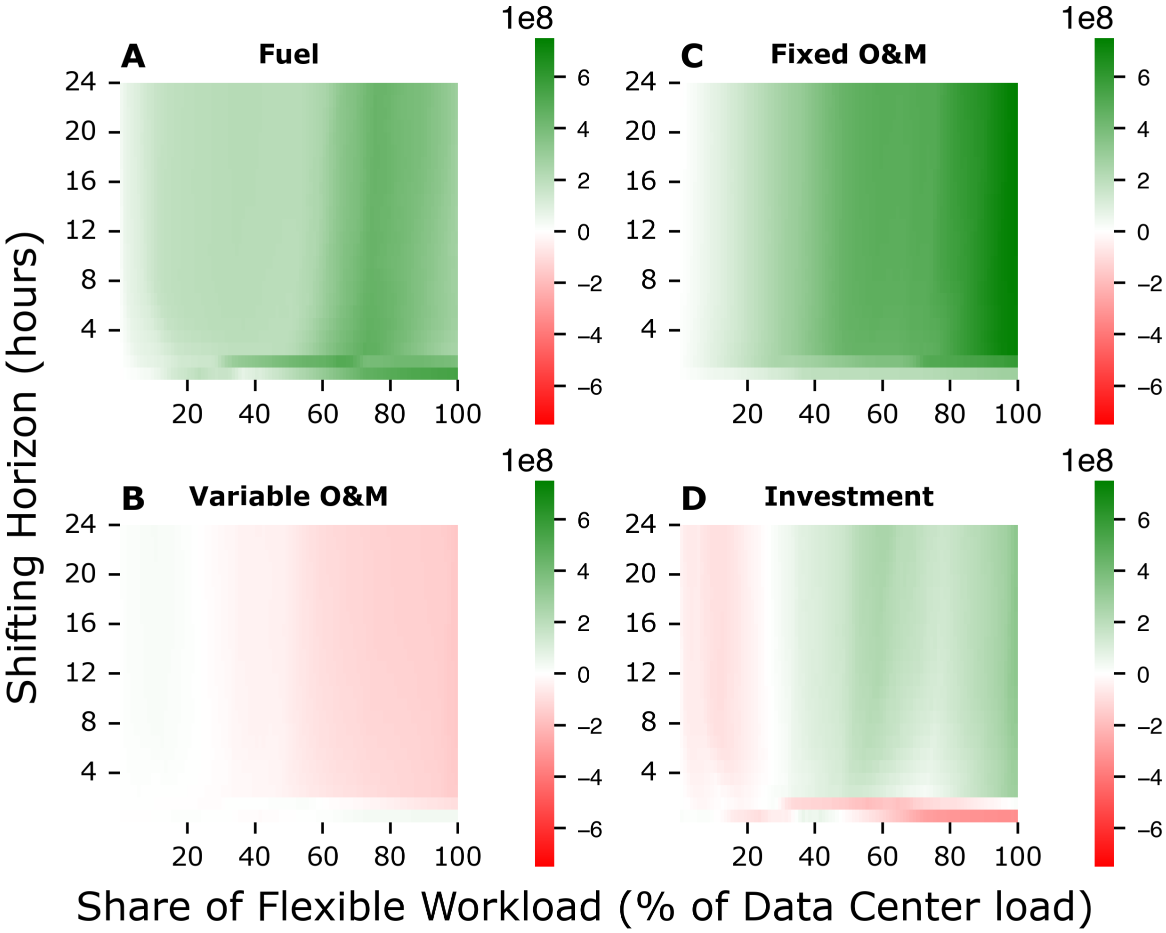

Figure S7: Cost Difference per Component for the Mid-Atlantic. Panels show the change in system costs between scenarios with and without data center flexibility for fuel **(A)**, fixed O&M **(B)**, variable O&M **(C)**, and generation investment **(D)**, across combinations of shifting horizon and flexible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

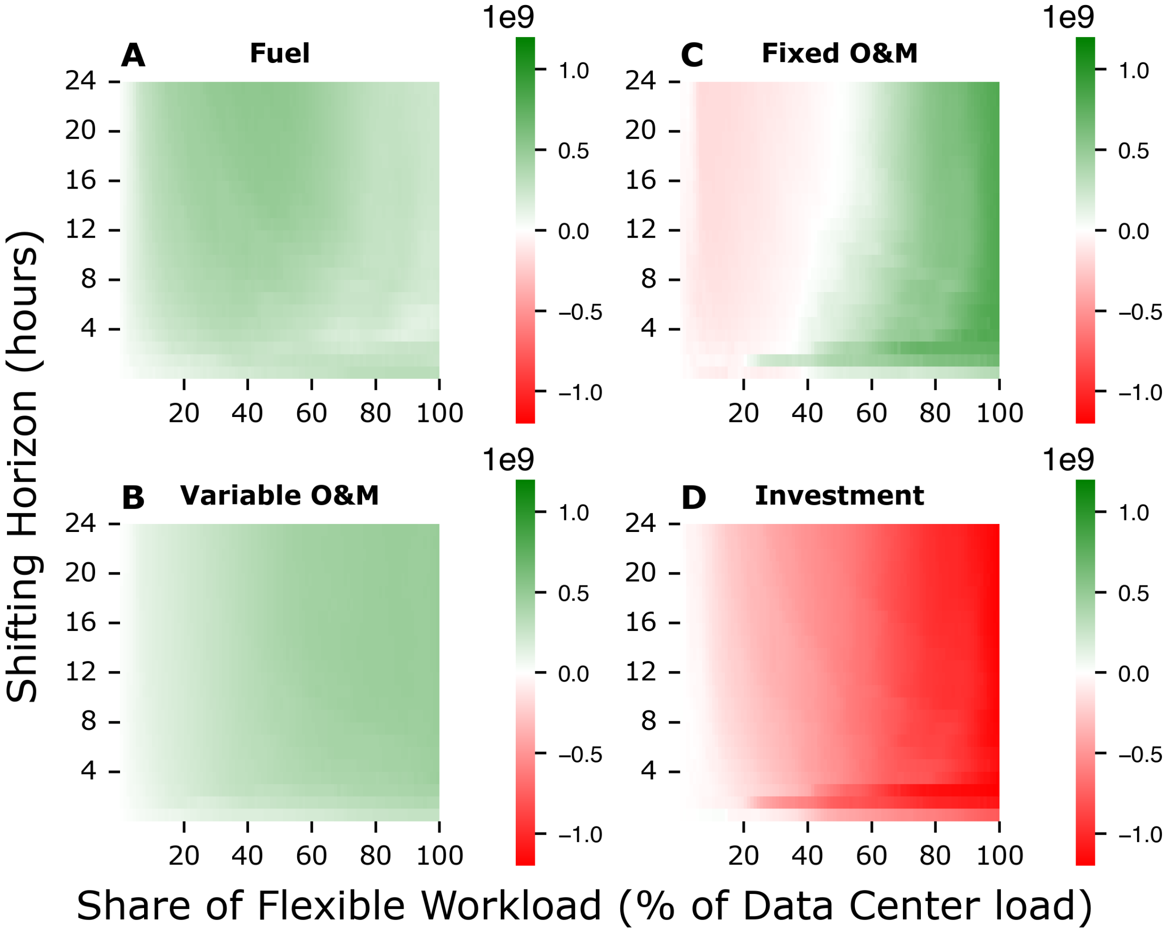

Figure S8: Cost Difference per Component for Texas. Panels show the change in system costs between scenarios with and without data center flexibility for fuel **(A)**, fixed O&M **(B)**, variable O&M **(C)**, and generation investment **(D)**, across combinations of shifting horizon and flexible workload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

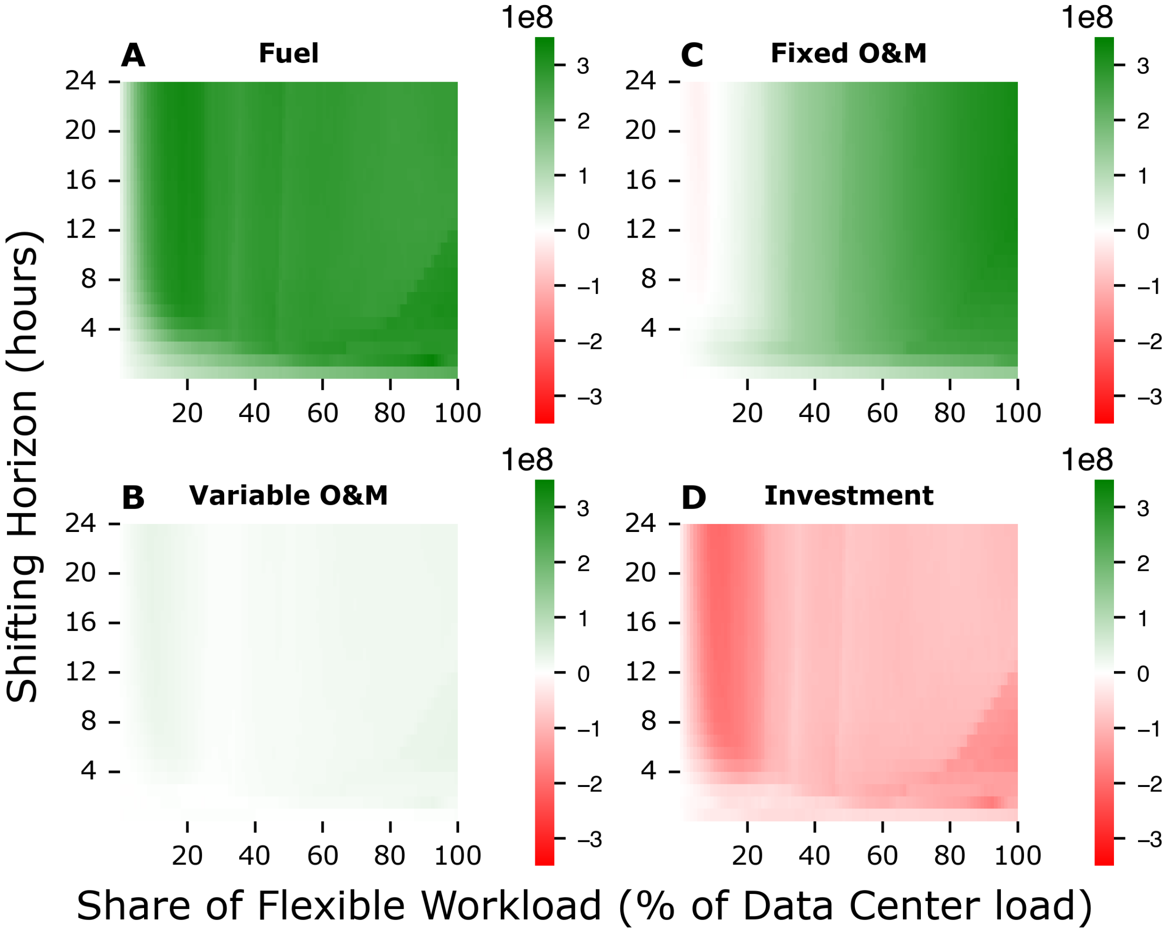

Figure S9: Cost Difference per Component for WECC. Panels show the change in system costs between scenarios with and without data center flexibility for fuel **(A)**, fixed O&M **(B)**, variable O&M **(C)**, and generation investment **(D)**, across combinations of shifting horizon and flexible workload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