

Changes in electricity use following COVID-19 stay-at-home behavior

Dylan Brewer*

November 1, 2022

Abstract

This article uses hourly electricity consumption data from the PJM Interconnection in the United States and stay-at-home metrics from cell phone location data to study the effect of the COVID-19 pandemic on electricity consumption using a difference-in-predicted-differences strategy. I show that while in the first months of the COVID-19 pandemic total electricity consumption declined by 2.7-3.8% relative to a predicted counterfactual, in June through August 2020 electricity consumption was 2.1-3.5% higher than the predicted counterfactual. Time spent at home reduces electricity consumption, and a reduction in time at home after May lead to increased electricity consumption in the summer months. In addition, higher temperatures had an increased effect on electricity consumption in 2020 relative to previous years. Nationwide monthly data on electricity consumption by load class reveals that commercial and industrial consumption was below its expected baseline from March-December 2020, while residential consumption was above its expected baseline, peaking in July. This suggests that increased demand for residential cooling offset declines in commercial and industrial demand for electricity. Estimates of the total effect of the pandemic on electricity consumption from March through December 2020 suggest that early reductions in electricity use were offset by later increases, implying that any expected “silver lining” of decreased emissions from electricity generation may be smaller than previously thought.

Keywords: COVID-19, work-from-home, social distancing, location data, electricity

*Georgia Institute of Technology, School of Economics, 221 Bobby Dodd Way, Office 224, Atlanta, Georgia 30332, brewer@gatech.edu. Declarations of interest: none. Thank you to comments by Soren Anderson, Matt Oliver, Paul Brehm, Alecia Cassidy, and from participants at the 2021 IAEE Online Conference and 2020 Southern Economics Association and Midwest Economics Association meetings.

1 Introduction

The COVID-19 pandemic and associated stay-at-home policies and behaviors have significantly damaged the health and livelihoods of millions of individuals worldwide. Moreover, many energy scholars showed that electricity demand had sharply declined during the early months of 2020 (Benatia, 2020; Benatia and Gringas, 2021; Brewer, 2020; Carroll et al., 2020; Narajewski and Ziel, 2020; Percy and Mountain, 2020), noting that electricity demand could serve as a real-time proxy for overall economic activity (e.g., Bui and Wolfers, 2020). In addition, one of the few “silver linings” expected to come out of the pandemic is a reduction in local and global pollution from a reduced demand for fossil fuels used for transportation and electricity generation (Gillingham et al., 2020; Cicala et al., 2020; Dang and Trinh, 2021). However, as stay-at-home policies expired in areas where the initial outbreak of the virus had subsided, commercial activity resumed to an extent, and individuals and firms had the opportunity to make up for lost economic activity during earlier months. In addition, as households spent more time in quarantine, they had the opportunity to adapt behavior to the home environment, potentially increasing the use of electrical appliances for work in new home offices or acquiring additional electronics for at-home leisure. Any long-run economic losses or environmental health benefits from changes in energy use during the pandemic would have to be from permanently *displaced* consumption of energy rather than *delayed* consumption of energy.

This paper provides the first empirical evidence that electricity use increased following the expiry of initial stay-at-home policies and behaviors. To evaluate the longer-term impact of staying at home on electricity consumption, I analyze data on hourly electricity use in the PJM Interconnection of the United States during 2020.¹ The analysis uses a nonparametric matching algorithm to predict electricity consumption for 2020 based on weather patterns and hourly, daily, and monthly seasonality if 2020 consumption resembled consumption from the previous five years. This predicted consumption serves as the counterfactual electricity

¹PJM serves 65 million people across 13 states plus the District of Columbia, including major population centers of Chicago, Cleveland, Philadelphia, Pittsburgh, and Washington D.C. In 2019, total billing of electricity in PJM was valued at \$39.2 billion (PJM, 2019).

consumption for a difference-in-predicted-differences estimation that compares the difference between actual electricity consumption and predicted consumption during 2020 before and after the start of the pandemic. This empirical strategy allows me to control for weather patterns, seasonality, and unobserved differences between 2020 and previous years. The results show that during the initial surge of COVID-19 cases and stay-at-home policies from March through May 2020, total electricity use was 2.7-3.8% lower than predicted consumption each month; however, after these policies expired in June 2020, total electricity use increased relative to the predicted baseline and in August was and 3.5% *higher* than predicted consumption.

There are several possible explanations for this increased electricity demand during the summer months. First, it is possible that people began to leave their homes again after the “first wave” of the pandemic ended, causing a rebound recovery. Second, it is possible that commercial and industrial consumption was delayed rather than displaced. For example, businesses may have attempted to fulfill contracts that were put on hold and operated more intensively to make up for lost time. Another explanation may be that the relationship between temperature and electricity consumption changed. This could occur if individuals working from home demanded more electricity for cooling during the summer months than in previous years while commercial and industrial buildings still needed to be cooled despite being at low occupancy. Finally, non-cooling-related residential electricity use may have increased. This may have occurred as work-from-home practices became more established and people used more office equipment from home, or if leisure activities changed. Implicit in the latter two hypotheses is a loss of scale economies in electricity use in workplaces with shared electricity-using equipment.

I analyze the relationship between stay-at-home behavior and electricity consumption using cell-phone location data provided by SafeGraph. These data include the median amount of time devices spent at home in each US census block. After matching census blocks to PJM zones, I use the amount of time spent at home as a third difference in the aforementioned

difference-in-predicted differences strategy.² I find that additional time spent at home resulted in decreased consumption of electricity. In addition, the data show that stay-at-home behavior peaked in April or May and declined in the summer months. This analysis indicates that a reduction in stay-at-home behavior contributed to the increased consumption of electricity in the summer, but this alone cannot fully explain the electricity consumption patterns. For instance, time spent at home in December 2020 was roughly on-par with that in December 2019, while electricity consumption was higher in 2020. Moreover, the effect of time at home changes from month to month, suggesting that households were spending time at home differently as the pandemic progressed.

To further examine the source of the surplus demand, I analyze monthly nationwide reported electricity consumption data by sector at the utility level. These data show that at the beginning of the pandemic in March 2020, electricity use in all sectors was lower relative to the previous five years after controlling for utility-specific unobservables, temperature, and seasonality. From April through August, residential electricity consumption was significantly higher relative to expected consumption, peaking in July. During the same period, commercial and industrial electricity demand was below normal, initially outweighing increased residential demand from April through June. In July and August, residential demand peaked while industrial and commercial demand were at the highest levels since March. This suggests that the additional load came primarily from the residential sector, which could be driven by demand for cooling, demand for work-from-home purposes, or demand for electricity as an input for leisure. In addition, there is some evidence for a partial commercial and industrial recovery during this period.

Returning to the hourly data, I then estimate the summer temperature-electricity exposure-response relationship for 2014-2019 versus 2020 to test the hypothesis that cooling more residential homes during the day is more costly than cooling workplaces. I find that in 2020, higher temperatures resulted in proportionally higher electricity consumption than in 2014-2019, though the effect is not large. Thus, it appears that the cost of cooling residential

²The triple differences compares predicted and actual electricity use, before and during the pandemic, and across areas with more stay-at-home behavior versus less stay-at-home behavior.

homes is driving some of the increased demand.

The findings in this paper are important because they suggest that some energy consumption was delayed rather than displaced by COVID-19, meaning that some of the economic and environmental health effects of the first wave of the pandemic were not permanent and may have been offset in the following months. Furthermore, it provides evidence that as individuals and businesses adapted to the pandemic, the relationship between electricity consumption and underlying drivers such as temperature and time-of-day have changed. Ultimately, the loss of economies of scale in cooling people at their places of work outweighed reductions in consumption due to new electricity-use habits.

Other new and forthcoming work has analyzed the effect of the pandemic on electricity consumption in New York (Benatia and Gringas, 2021), Texas (Cicala, 2022), France (Benatia, 2022), and Spain (Bover et al., 2022). In each of these contexts, residential consumption increased while industrial and commercial consumption decreased. Similarly, in New York, Texas, and Spain, early reductions in electricity consumption were partially made up for by the increase in residential electricity consumption during the second half of 2020, but in these regions the overall effect of the pandemic was to reduce electricity consumption. My finding is distinct in that I estimate a slight increase in electricity consumption over 2020 in PJM, which is driven by rural zones in the PJM service area. In this paper, I find that the PJM zone corresponding to Chicago shows persistent reductions over all of 2020, while zones serving rural areas display the largest summer increases in consumption. This is consistent with the findings of Benatia and Gringas (2021), whose estimates for New York City are similar to what I find in Chicago. Benatia and Gringas (2021) also focuses on the effect of the pandemic on the wholesale market, finding that forecasting errors result in reduced efficiency in the wholesale market. While I do not focus on the wholesale market in this paper, Brewer (2020) finds a similar increase in forecast errors in PJM during the early part of the pandemic. Benatia (2022) shows that demand reductions due to COVID-19 in France led to large declines in electricity prices and total revenues in the energy sector, which may reflect a future with a larger renewable energy generation mix. Finally, the use of cell phone location

data to estimate the effect of time spent at home on electricity consumption is distinct to this paper. Bover et al. (2022) find that their estimates of residential electricity consumption are correlated with time spent at home measured by Google’s Residential Mobility Index, but do not directly estimate the effect of time spent at home on consumption. Similarly, Cicala (2022) finds that residential consumption is correlated with the fraction of workers who could work from home or number of non-essential business closures.

In addition, this paper contributes to the emerging literature on the economic and environmental impacts of the COVID-19 pandemic and subsequent policy and behavioral responses. Fabra et al. (2022) use the pandemic as a natural experiment to study the effect of economic growth on carbon emissions in Spain. Brodeur et al. (2021) show that early COVID-19 stay-at-home orders reduced vehicle crashes by 20% and that particulate matter emissions decreased as well. Irwin and Livy (2021) find that during the pandemic, the perceived value of being near major roadways decreased while the value of being near open space did not change. Isphording and Pestel (2021) and Persico and Johnson (2021) study the effect of air quality on the severity of the pandemic, finding that more air pollution was associated with increased COVID-19 mortality. Other work has examined the effect of the pandemic on a variety of outcomes including labor supply and demand (Forsythe et al., 2020; Kong and Prinz, 2020; Auray and Eyquem, 2020; Couch et al., 2020; Hensvik et al., 2021), domestic violence (Leslie and Wilson, 2020; Baron et al., 2020), and crime (Abrams, 2021) among other outcomes. In addition, this paper contributes to recent work using machine-learning algorithms to estimate counterfactual consumption for a natural experiment with no control group in the fashion of Burlig et al. (2020). This machine-learning-enhanced prediction approach to study the pandemic’s effect is also used in related applications by Benatia and Gringas (2021), Graf et al. (2021), and Fabra et al. (2022).

The next section summarizes the data used in the analysis and plots 2020 hourly electricity consumption with predicted electricity consumption from 2014-2019. Section 3 discusses the empirical strategies used to identify the effect of the pandemic on electricity consumption as well as analyze the primary drivers of the effect. Section 4 describes the results, and

PJM Zones

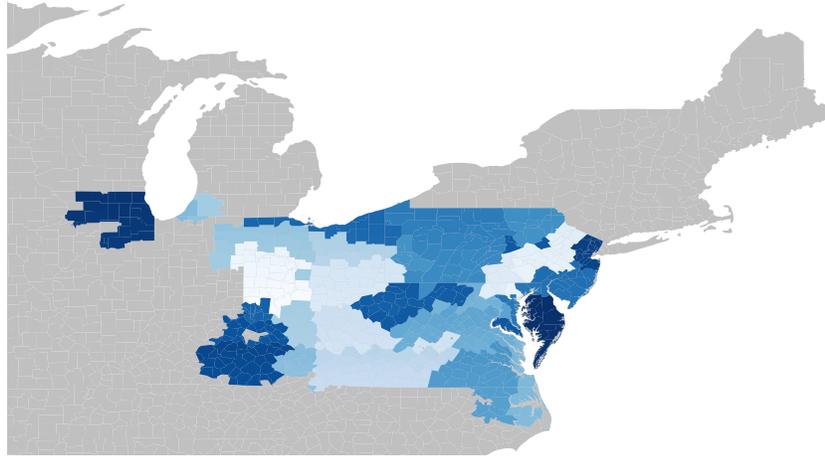


Figure 1: Each shade of blue is a PJM zone. PJM serves 13 states plus the District of Columbia, including major population centers of Chicago, Cleveland, Philadelphia, Pittsburgh, and Washington D.C. Shape files used to construct this map were provided by Jeremy Lin and Dazhi Yang.

section 5 concludes.

2 Data

To analyze the impact of COVID-19 on electricity consumption, I assemble data on metered hourly electricity consumption in the PJM Interconnection, one of nine Independent System Operators (ISOs) managing electricity distribution in the United States (PJM, 2020). PJM coordinates electricity distribution to individuals and firms in Delaware, Illinois, Indiana, Kentucky, Maryland, Michigan, New Jersey, North Carolina, Ohio, Pennsylvania, Tennessee, Virginia, West Virginia and the District of Columbia. My sample period runs from January 1, 2014 through December 31, 2020.³ Figure 1 maps the PJM zones, which

³Data on electricity consumption for the hour beginning at 2 AM on March 8, 2020 is missing online, so electricity consumption is interpolated from the surrounding hours. The data are likely missing due to the

cover the major population centers of Chicago, Cleveland, Philadelphia, Pittsburgh, and Washington D.C. In 2019, total billing of electricity in PJM was valued at \$39.2 billion and there were 65 million people living in its service area (PJM, 2019).

Hourly weather data comes from the Integrated Surface Database maintained by the United States National Centers for Environmental Information (2020a). PJM is divided into subregions called “zones.” Each zone’s weather is defined as a weighted average of one or several nearest weather stations following definitions published by PJM (2018). I construct all zone-level temperatures and precipitation for use as predictors of electricity consumption.⁴

To investigate the relationship between stay-at-home behavior and electricity consumption, I use data on social distancing provided by SafeGraph. These data use cell-phone GPS records to create a measure of the median amount of time spent at home each day for devices in each census block in the United States. I match census blocks to PJM zones to construct a measure of the average median amount of time spent at home daily and the total number of devices with a home location in each PJM zone daily.

Finally, to determine which sectors are responsible for changes in load patterns, I analyze a countrywide panel of monthly electricity demand at the utility-by-state level collected on the EIA form 861-M (United States Energy Information Administration, 2020b).⁵ These data describe monthly electricity consumption by residential, commercial, industrial, and transportation users. I match each utility to the counties it serves via EIA Form 861-S (United States Energy Information Administration, 2020a), which allows me to control for monthly weather patterns using data from the United States National Oceanic and Atmospheric Administration’s nClimDiv county-level database (National Centers for Environmental Information, 2020b). For each utility, the monthly weather includes the average, high, and low

beginning of daylight savings time at this hour and day in 2020. This interpolation does not affect any of the findings.

⁴0.1% of the station-hour observations are missing temperature data, and 0.5% of the station-hour observations are missing precipitation data. I interpolate the missing temperature data using the observed temperatures in surrounding hours using a simple linear interpolation method. Missing precipitation data are imputed as having no precipitation (92% of all non-missing station hours indicated that there was no precipitation). Inclusion of the interpolated and imputed data does not substantially change the results.

⁵The unit of observation is monthly by utility by state. If a utility serves customers in two states in the same month, the utility will have two entries, one under each state.

monthly temperatures and average monthly precipitation across the counties it serves.⁶ Due to gaps in coverage of the nClimDiv weather data for Hawaii and Alaska, I restrict analysis to electricity consumption in the contiguous United States.

3 Empirical strategy and suggestive evidence

To measure the causal effect of COVID-19 on electricity consumption, one would ideally compare electricity load during COVID-19 2020 to a counterfactual 2020 not affected by the pandemic. Because this counterfactual 2020 does not exist, the main empirical challenge is estimating what electricity consumption would have been without COVID-19. Electricity consumption varies hourly and based on day of the week due to sleep and work schedules. In addition, weather patterns affect electricity consumption by changing the demand for heating and cooling energy services. Thus, I use weather patterns and hourly, daily, and monthly seasonality to predict counterfactual electricity consumption for 2020 using observed electricity consumption from 2014-2019.

To generate a counterfactual 2020 electricity consumption pattern, I use a matching algorithm that matches hours in 2020 to the most similar hour in the last five years within the same zone. For each zone in each hour of 2020, the algorithm searches among all hours within the past five years exactly within the same month-of-year, hour-of-day, and day-of-week, and finds the hour with the most similar temperature and precipitation as measured using the Mahalanobis metric.⁷ Thus, for the hour of 10:00-11:00 AM of Sunday 16 August 2020 in the COMED zone serving Chicago, the algorithm searches among all August Sunday 10:00-11:00 AM hours in COMED between 2014 and 2019 for the hour with the most similar weather.⁸ This approach has precedent: ISOs such as PJM use similar matching algorithms to forecast day-ahead load alongside other prediction algorithms

⁶For some utilities, no counties are listed in the service territory data from the EIA. Utilities with missing service territory represents less than 1.5% of generation from 2014-2020. These utilities are excluded from the analysis.

⁷Appendix A provides a more precise description of the matching algorithms.

⁸I resolve ties by averaging the load within all the matched hours when applicable. Matching on relative humidity did not significantly change the results.

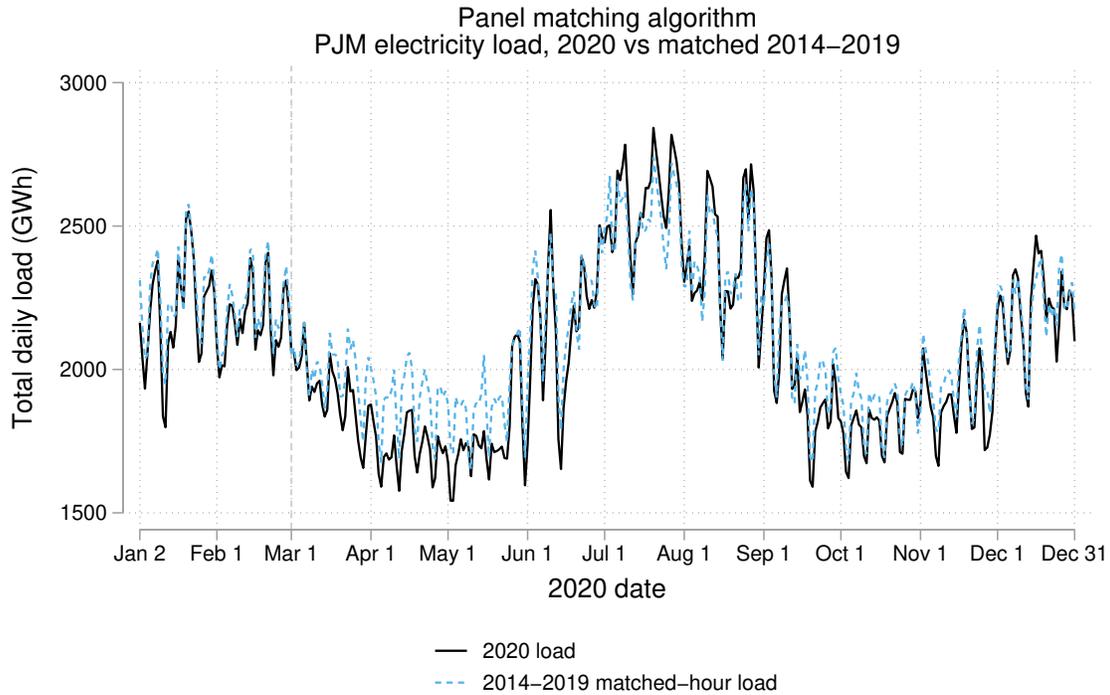


Figure 2: 2020 Electricity load matched to the most-similar weather hour from the previous five years, matching within zone exactly on the month-of-year, hour-of-day, and day-of-week. Figure displays the series in units of mean daily load for ease of interpretation.

(Anastasio, 2020). In addition, I perform the same prediction procedure to predict 2019 electricity consumption using observed electricity consumption from 2014-2018, which allows me to correct for potential prediction bias correlated with monthly seasonality or zone fixed effects.

After matching, the result is a panel of counterfactual hourly zone-level electricity consumption from 2019-2020. I aggregate this up to the PJM level by day and plot 2020's consumption and counterfactual consumption in figure 2. First, it is important to note that prior to COVID-19 (around March 1), the matched electricity load looks to be a close predictor of actual electricity load, which provides confidence that the algorithm provides good estimates of what electricity consumption would have been in a counterfactual 2020 without the virus. In appendix section A, I further investigate the goodness of fit by predicting 2019 zone-level electricity load using load from 2014 - 2018. I compare this to an approach that

forecasts load using linear regression as well as an approach that predicts aggregate load for all of PJM using a similar matching process. I find that the zone-level matching approach performs best in a root-mean-squared error sense and in percent prediction error (RMSE = 71.3, percent error = 2.5%), followed by the regression approach (RMSE = 81.4, percent error = 2.9%), and finally the PJM-level matching approach (RMSE = 170.7, percent error = 6.3%). Given the disparity in performance, I prefer modeling counterfactual electricity demand using the zone-level matching approach, which I use throughout the paper, though I replicate the results using the regression approach in appendix B. Mean daily load in 2019 was 2,155 GWh (s.d. = 287), so the prediction errors are relatively small. For reference, the typical day-ahead forecast error published by the PJM market operator is typically off by around 2% (Brewer, 2020), though this is not a perfect comparison given the year-ahead prediction approach taken here.

Looking at the matched load and realized load in figure 2, one can clearly see the effect of stay-at-home behavior between March 1st and June 1st. After June 1st, however, electricity consumption begins to track much more closely with the predicted consumption. By the beginning of July and through August, electricity consumption is much closer to expected consumption. Finally, it is interesting to note that the gap between consumption and predicted consumption widens again beginning in September, perhaps due to the rising number of COVID-19 cases and subsequent behavioral responses.

Denoting Y_{zt} electricity consumption in zone z during hour t and \hat{Y}_{zt} the matched electricity consumption, $\frac{Y_{zt} - \hat{Y}_{zt}}{\hat{Y}_{zt}}$, is the percent difference between predicted and actual electricity consumption. When $t \in 2019$, this fraction is the percent prediction error, and when $t \in 2020$, this fraction is the percent difference between COVID-19 2020 and counterfactual 2020 that would be expected if electricity consumption looked similar to the last five years. This difference cannot be solely attributed to COVID-19 because even without COVID-19, electricity demand would be likely to differ from previous years due to differences in economic activity or differences in energy efficiency, which is likely given year-over-year declines in electricity consumption in the United States. To adjust for this possibility, I estimate

the following difference-in-predicted-differences model on the data from all of 2019 and 2020 using a least-squares regression:

$$\frac{Y_{zt} - \hat{Y}_{zt}}{\hat{Y}_{zt}} = \delta_{zymdh} + \sum_{m=3}^{12} \beta_m post_t 1(month_t = m) + \varepsilon_{zt}, \quad (1)$$

where $post_t$ is an indicator variable equal to one for all hours after the beginning of the COVID-19 period, and β_m varies by month. The term δ_{zymdh} represents controls for fixed effects by zone z , indicators for year y , month of year m , and day of week times hour of day dh , and an intercept. Estimates of β_m measure the heterogeneous effect of COVID-19 using a bias-correction approach to estimate the difference in predicted differences before and after COVID-19 stay-at-home behavior began. It is unclear what the exact treatment date was: most state governments issued the first advisory or mandatory stay-at-home orders between March 7th and March 15th, but voluntary and private-sector-lead stay at home behavior likely began earlier. Furthermore, news announcements of one state’s stay-at-home policies may have influenced private behavior in other states, so it is not clear that coding each state’s treatment date heterogeneously is a better approach. Instead, I consider separate specifications with treatment dates of March 1st, 8th, and 15th as well as “doughnut specifications” that exclude all hours from March 1-7 or March 1-15 from the estimating sample. My preferred specification is that which excludes March 1 - 15 from the sample and thus avoids measurement error in the treatment variable. In addition, I also estimate the effect separately for zones corresponding to major cities to characterize regional heterogeneity.

This difference-in-predicted-difference estimator is similar to the estimator implemented in Burlig et al. (2020). The estimated coefficients reveal whether earlier reductions in electricity consumption were offset by later increases when individuals began returning to work, controlling for differences between 2020 and prior years, weather, and hourly, daily, and monthly seasonality. The included fixed effects by zone, indicators for year, month of year, and day of week times hour of day control for the possibility of prediction error correlated with these variables. To corroborate the difference-in-predicted differences estimates,

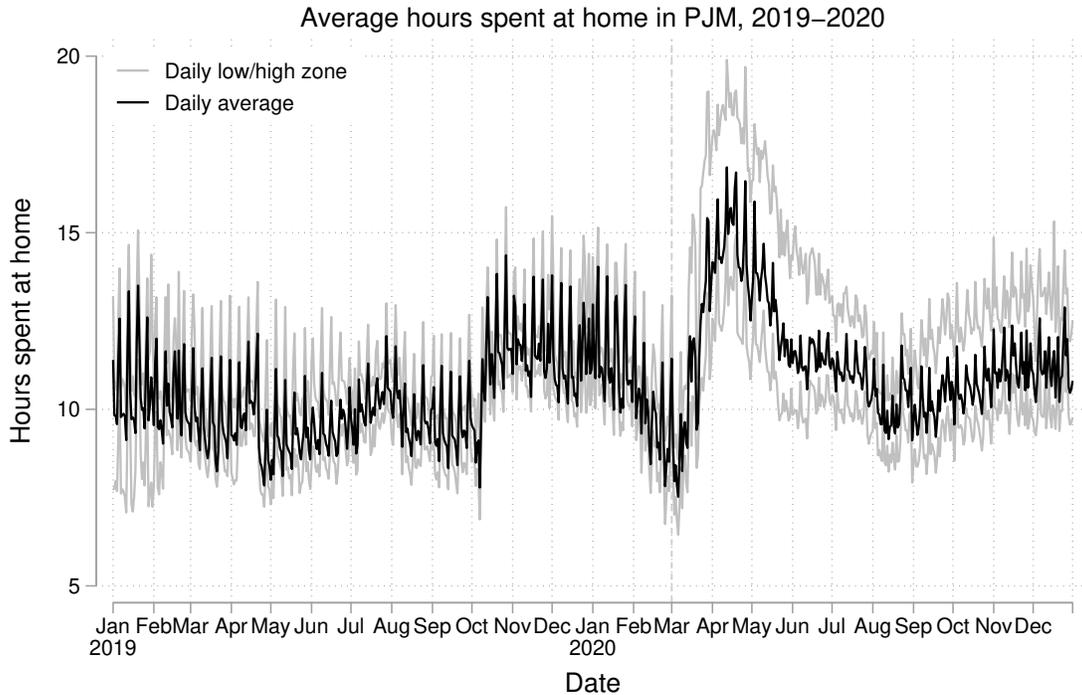


Figure 3: Average median time spent at home in PJM for 2019 through 2020 using cell phone location data. Each observation is the daily average of each census block’s time spent at home, weighted by the number of devices in each census block. Gray lines are the average median time spent at home for the highest and lowest zone.

I provide the results from a linear regression analysis in appendix B. The linear regression estimates are ultimately similar to the estimates in the main analysis, though I prefer the results from the difference-in-predicted differences procedure given the matching algorithm’s improvement in estimating electricity consumption as shown in Appendix section A.

Next, I analyze the role of stay-at-home behavior on electricity consumption using the cell phone location data. Figure 3 displays the variation in stay-at-home behavior over time and zone within PJM—after an initial increase in stay-at-home behavior, people began again to leave home. One can clearly see that substantial stay-at-home behavior did not begin until mid-to-late March and that it peaked in May. Furthermore, heterogeneity in time spent at home was substantially greater after March 2020 than in previous years. I estimate the relationship between stay-at-home behavior and electricity consumption in a triple differences specification where the first difference is between predicted and observed

consumption, the second difference is pre- and post-March 1, and the third difference is the continuous measure of time spent at home:

$$\frac{Y_{zt} - \hat{Y}_{zt}}{\hat{Y}_{zt}} = \delta_{zymdh} + \sum_{m=3}^{12} \beta_m post_t 1(month_t = m) + \eta_m home_{zt} \cdot post_t 1(month_t = m) \quad (2)$$

$$+ \eta_0 home_{zt} + \kappa devices_{zt} + e_{zt},$$

where $home_{zt}$ is the average median number of hours spent at home on each day.⁹ The term δ_{zymdh} represents controls for fixed effects by zone z , indicators year y , for month of year m , day of week times hour of day dh , and an intercept. The variable $devices_{zt}$ is a count of devices observed in each zone and controls for the purchase of new mobile devices, and e_{zt} is an error term. I allow the coefficients on the interaction between $home_{zt}$ and $post_t$ to vary by month. The estimates of interest are the sum of the coefficients on $home_{zt}$ for each month $\eta_0 + \eta_m$, which measure the marginal effect of stay-at-home behavior on electricity load throughout the pandemic. Similarly to before, I corroborate these results using a regression approach without the first-stage prediction in appendix section B and ultimately find similar results.

To further understand what is driving the load patterns, I turn to a panel analysis of utility-reported monthly electricity consumption by customer class (residential, industrial, and commercial customers).¹⁰ Changes in residential electricity consumption reflect new behaviors by individuals working and spending leisure time at home, while changes in industrial and commercial electricity consumption reflect changes in working hours and production intensity. Denoting Y_{imc} reported electricity consumption by utility i in month m by customer class c , I estimate the following fixed-effects regressions separately by customer class using

⁹That is, the average of the median hours for each census block in a zone, weighted by the number of devices observed.

¹⁰To simplify tables and figures, I do not present the analysis of consumption by transportation customer class, which remains a small (but growing) share of electricity consumption.

data from 2014 through 2020:

$$Y_{imc} = \mu_{imc} + \omega_c weather_{imc} + \sum_{m=3}^{12} \gamma_{mc} post_m 1(month_m = m) + \xi_{imc}. \quad (3)$$

μ_{imc} represents controls for utility fixed effects, month of year, and an intercept term. The variable $weather_{imc}$ is a vector of weather controls containing the monthly high, low, and average temperatures in the utility's service area. The treatment variable $post_m$ is an indicator variable equal to one for all months after March 2020, and similarly to the PJM-level analysis I allow the estimated coefficient to vary by month. ξ_{imc} is an error term. The estimates of γ_{mc} represent the heterogeneous effects of the COVID-19 pandemic by month for each customer class, controlling for differences in weather, monthly trends, and fixed differences across utilities. Thus, the variation used to estimate the effect of COVID-19 on electricity consumption patterns comes from the difference in electricity consumption before and after March 2020 within utility, within month of year, and controlling for weather.

Finally, I investigate the hypothesis that the relationship between outdoor temperature and electricity load has changed using the hourly zone-level consumption and weather data from PJM. To examine whether there is a changed relationship between temperature and electricity use, I return to the PJM hourly zone-level data. I estimate the following regression of logged hourly electricity consumption $\ln(MW_{zt})$ on a series of five-degree-Fahrenheit indicator variables $\sum_b 1(temp_{zt} \in b)$ for $b \in \{[10, 15), [15, 20), \dots, [90, 95), \geq 95\}$, allowing the relationship to vary for 2020 versus previous years:

$$\ln(MW_{zt}) = \delta_{zymdh} + \sum_b (\psi_b 1(temp_{zt} \in b) + \phi_b 1(temp_{zt} \in b) \times post_t) + \rho post_t + u_{zt}. \quad (4)$$

The term δ_{zymdh} includes controls for fixed effects by zone z , indicators by year y , indicators for month of year m , day of week times hour of day dh , and an intercept. The term u_{zt} is the residual. I allow the estimated coefficients on the temperature indicators to differ for 2020 and 2014-2019 and omit the temperature interval of 55-60 degrees Fahrenheit as the base

case, given it is the temperature bin with the least energy consumption.¹¹ Given this choice of omitted category, $\hat{\psi}_b$ estimates the difference in logged energy consumption for interval b relative to the 55-60 interval, and $\hat{\phi}_b$ estimates the difference in the relationship between temperature and electricity load in 2020 versus 2014-2019.

In the next section, I summarize the results from the difference-in-predicted differences model, the triple differences model using the cell phone location data, the consumer-class level estimates, and the analysis of the differential effect of temperature on electricity use in 2020.

4 Results

Table 1 displays the estimates of the difference-in-predicted-differences model. Standard errors account for multi-way clustering by zone and by day (Cameron et al., 2011).¹² These results confirm the graphical intuition from earlier that some electricity consumption may have been delayed to the summer months rather than displaced. Across all specifications, electricity consumption in March through May was below normal relative to the predicted baseline. In June, July, and August, electricity consumption was above the predicted baseline, though the difference is only statistically significant in August. In the fall, electricity consumption returned to pre-COVID levels before again increasing in December.

Figure 4 displays the estimates from the preferred specification that excludes March 1-15 (column 5). In March (beginning the 15th), April, and May, electricity consumption was 3.8%, 3.8%, and 2.7% lower relative to the predicted baseline levels in each month, respec-

¹¹Because outdoor temperature is all relative, there is no natural zero to choose as the base case. The choice of the interval 55-60 reflects that at this temperature, the least amount of electricity is used in PJM conditional on the aforementioned indicator variables. This likely reflects the temperature at which there is the least demand for electricity for space heating or cooling.

¹²These standard errors do not account for uncertainty from the first-stage matching procedure. It would be straightforward but computationally costly to bootstrap the entire estimation process; however, it is unclear whether this would provide an improvement in the standard errors. First, the bootstrap is known to fail with matching estimators (Imbens and Abadie, 2008). Second, in a similar application, Burlig et al. (2020) found that appropriately clustered standard errors are similar to fully bootstrapped standard errors. Finally, I am able to replicate the results using a regression approach with standard errors that account for all sources of uncertainty in that procedure (see appendix section B), lending confidence to the precision of the estimates here.

tively. In June and July, electricity consumption returned to baseline, with point estimates of 2.1% and 2.8% higher relative to the predicted baseline, although the confidence interval of these estimates contains zero. In August, electricity consumption spiked, coming in at 3.5% higher relative to the predicted baseline. In September, October, and November, electricity consumption returned to the baseline. In December, electricity consumption was 3.0% higher. After weighting by total load in each month, the difference-in-predicted-difference estimates suggest an overall 1.08% increase in electricity consumption after March 15th relative to the predicted baseline. In appendix section B, I corroborate the results using a regression-approach that does not involve a first-stage matching procedure. This regression approach returns estimates that follow a similar pattern and are larger in magnitude, but I prefer the difference-in-predicted differences approach given the improved ability of the matching algorithm to model electricity load (as shown in appendix section A). Finally, I replicate the difference-in-predicted differences procedure for 2019 in appendix section C and find that consumption in 2019 was lower than the expected baseline, which is consistent with gains in energy efficiency or increases in behind-the-meter solar generation. This suggests that the findings of summer increases in consumption may be a conservative lower bound if there would have been additional adoption of energy efficiency or behind-the-meter generation in 2020 had the pandemic not occurred.

Why was electricity consumption higher in July, August, and December? This pattern could reflect households resuming normal activities after the first wave of the pandemic ended or electricity consumption that was delayed rather than displaced by stay-at-home behavior. To analyze the role of stay-at-home behavior in driving these electricity consumption patterns, I incorporate stay-at-home metrics derived from SafeGraph cell-phone location data.

Table 2 displays the estimates from the triple-differences regression specified in equation 2 using the cell phone location data. Columns 1 and 2 present triple differences estimates and columns 3 and 4 allow the coefficients on the post period variable and interactions to vary by month. In columns 2 and 4 I control for the number of devices in each zone to account

Table 1: Difference-in-predicted-differences estimates

	(1)	(2)	(3)	(4)	(5)
Post × March	-0.010 (0.009)	-0.028* (0.007)	-0.038* (0.007)	-0.022* (0.009)	-0.031* (0.009)
Post × April	-0.035* (0.010)	-0.038* (0.009)	-0.038* (0.009)	-0.035* (0.010)	-0.035* (0.010)
Post × May	-0.024* (0.011)	-0.027* (0.011)	-0.027* (0.011)	-0.024* (0.011)	-0.024* (0.011)
Post × June	0.024 (0.014)	0.020 (0.014)	0.021 (0.014)	0.024 (0.014)	0.024 (0.014)
Post × July	0.032 (0.016)	0.028 (0.016)	0.028 (0.016)	0.032 (0.016)	0.032 (0.016)
Post × August	0.038* (0.013)	0.035* (0.013)	0.035* (0.013)	0.038* (0.013)	0.038* (0.013)
Post × September	0.020 (0.011)	0.016 (0.011)	0.016 (0.011)	0.020 (0.011)	0.020 (0.011)
Post × October	0.013 (0.010)	0.010 (0.009)	0.010 (0.009)	0.013 (0.010)	0.013 (0.010)
Post × November	-0.003 (0.010)	-0.007 (0.010)	-0.006 (0.009)	-0.003 (0.010)	-0.003 (0.010)
Post × December	0.034* (0.010)	0.030* (0.010)	0.030* (0.010)	0.034* (0.010)	0.034* (0.010)
Overall effect	1.10%	0.67%	0.71%	1.06%	1.08%
Treatment date	March 1	March 8	March 15	March 9	March 16
Excluded dates				March 1-8	March 1-15
Zone FE	Y	Y	Y	Y	Y
Year, Month, DOW X Hour	Y	Y	Y	Y	Y
Observations	368382	368382	368382	364350	360822

Dependent variable is the percent difference between hourly metered electricity load and predicted load. Estimates are the monthly difference in predicted differences before and after March 1st, 2020 and when multiplied by 100 are interpreted as the percentage difference in electricity consumption relative to the baseline prediction. Standard errors clustered to account for multi-way clustering by zone and by day. * p-value < 0.05.

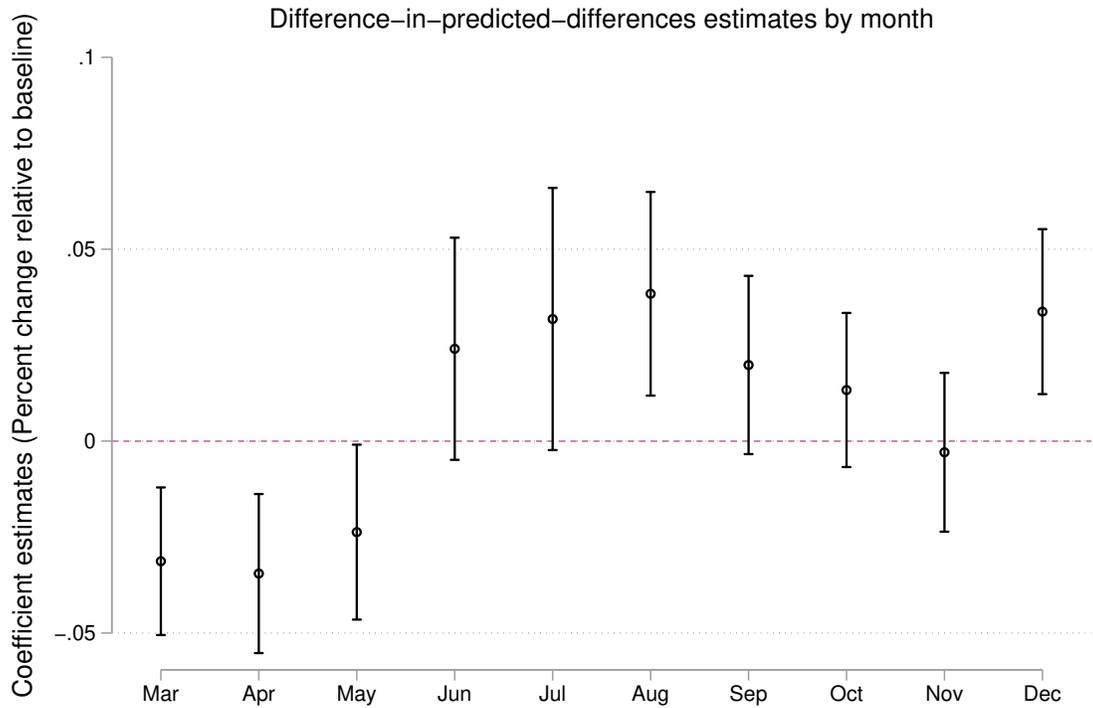


Figure 4: Estimates from column (5) of table 1 of the difference-in-predicted-differences model in equation 1. 95% confidence intervals account for multi-way clustering by zone and by day. Estimates are the monthly difference in predicted differences before and after March 1st, 2020 after omitting observations from March 1-15, and when multiplied by 100 are interpreted as the percentage difference in electricity consumption relative to the baseline prediction.

for the purchase of new devices as the year progressed. The overall effect of additional time spent at home on electricity consumption is the sum of the coefficients on home hours and home hours interacted with the post-pandemic variables. I plot the marginal effect of an additional hour spent at home in each month in figure 6 (the sum of the coefficients on home hours and home hours interacted with the monthly $Post_t$ variables). Overall, the effect of an additional hour at home on average is a reduction in electricity consumption of about 0.9%. This effect varies somewhat from month-to-month, but the point estimates are consistently negative. In appendix section B, I corroborate the results of the triple-differences analysis using a regression-approach that does not involve a first-stage matching procedure. This regression approach returns estimates of a similar pattern, but I prefer the difference-in-predicted differences approach given the improved ability of the matching algorithm to model electricity load (as shown in appendix section A).

In figure 5, I display separate estimates for urban zones corresponding to Chicago, Baltimore, Washington DC, Philadelphia, Cincinnati, and Pittsburgh. Each zone is below baseline consumption in the early months of the pandemic. Baltimore, Cincinnati, and Washington DC display increases in summer consumption over the baseline similar to the average trend (and other rural zones predominantly follow this pattern). Chicago and Pittsburgh do not exhibit the summer increase over baseline. The patterns in Chicago are particularly unique. Other than in June and December, Chicago's electricity consumption was below baseline. The estimated patterns for Chicago are similar to the estimates for New York City as shown by Benatia and Gringas (2021). Thus, while some cities saw persistent declines in electricity consumption, others saw recoveries and summer increases. Figure 15 in appendix E displays heat maps showing the percent change in electricity consumption by zone for each month of 2020.

Given the negative effect of time at home on electricity consumption and the relative decline of time spent at home beginning in June, July, and August, it appears that a reduction in stay-at-home behavior contributed to additional electricity consumption in the summer. Moreover, the effect of time at home changes from month to month, suggesting that house-

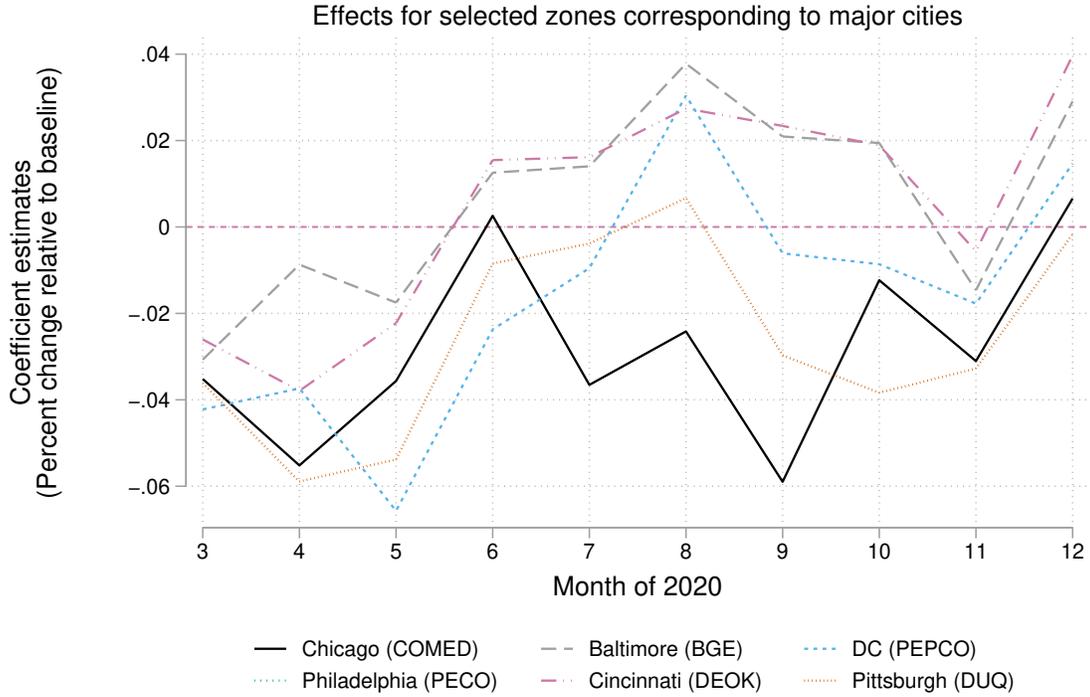


Figure 5: Heterogeneity for selected zones corresponding to major cities.

holds were spending time at home differently as the pandemic progressed. These results cannot explain all of the 2020 trends in electricity use; for instance, electricity consumption was above the expected value for December, but time spent at home was comparable to December 2019. Thus, I turn to an analysis by sector.

Table 3 displays the estimates from equation 3 using the nationwide panel of reported monthly consumption by class, and figure 7 displays these estimates graphically. These results confirm the overall consumption trends estimated with the metered PJM data. The total effect on consumption in these sectors is the sum of the estimates in each month; thus, in the early months of the pandemic, total electricity consumption was down relative to normal with slight increases in residential consumption. In April through June, increased residential electricity demand was cancelled out by large reductions in commercial and industrial electricity use. July and August featured much higher than typical residential electricity consumption, and while commercial and industrial electricity consumption was closest to its baseline prediction since the start of the pandemic, it was still slightly lower than in typical

Table 2: Triple-differences estimates

	(1)	(2)	(3)	(4)
Post	0.062 (0.030)	0.061 (0.030)		
Home hours	-0.005* (0.002)	-0.005* (0.002)	-0.005* (0.002)	-0.005* (0.002)
Post × Home hours	-0.004 (0.003)	-0.004 (0.003)		
Device count (100k devices)		-0.002 (0.003)		-0.001 (0.003)
Post × March × Home hours			-0.004 (0.002)	-0.004 (0.002)
Post × April × Home hours			0.009* (0.003)	0.009* (0.003)
Post × May × Home hours			-0.004 (0.003)	-0.003 (0.003)
Post × June × Home hours			-0.007 (0.007)	-0.007 (0.007)
Post × July × Home hours			-0.012 (0.012)	-0.012 (0.012)
Post × August × Home hours			0.004 (0.006)	0.003 (0.006)
Post × September × Home hours			-0.009 (0.006)	-0.009 (0.006)
Post × October × Home hours			0.003 (0.003)	0.003 (0.003)
Post × November × Home hours			0.001 (0.005)	0.001 (0.005)
Post × December × Home hours			-0.009 (0.004)	-0.009 (0.004)
Post X Month			Y	Y
Zone FE	Y	Y	Y	Y
Year, Month, DOW X Hour	Y	Y	Y	Y
Observations	350841	350841	350841	350841

Dependent variable is the percent difference between hourly metered electricity load and predicted load. Home hours is the average median time spent at home by device users in census blocks for each PJM zone and Device count controls for the number of devices in each zone (units in hundred thousands). Standard errors clustered to account for multi-way clustering by zone and by day. * p-value < 0.05.

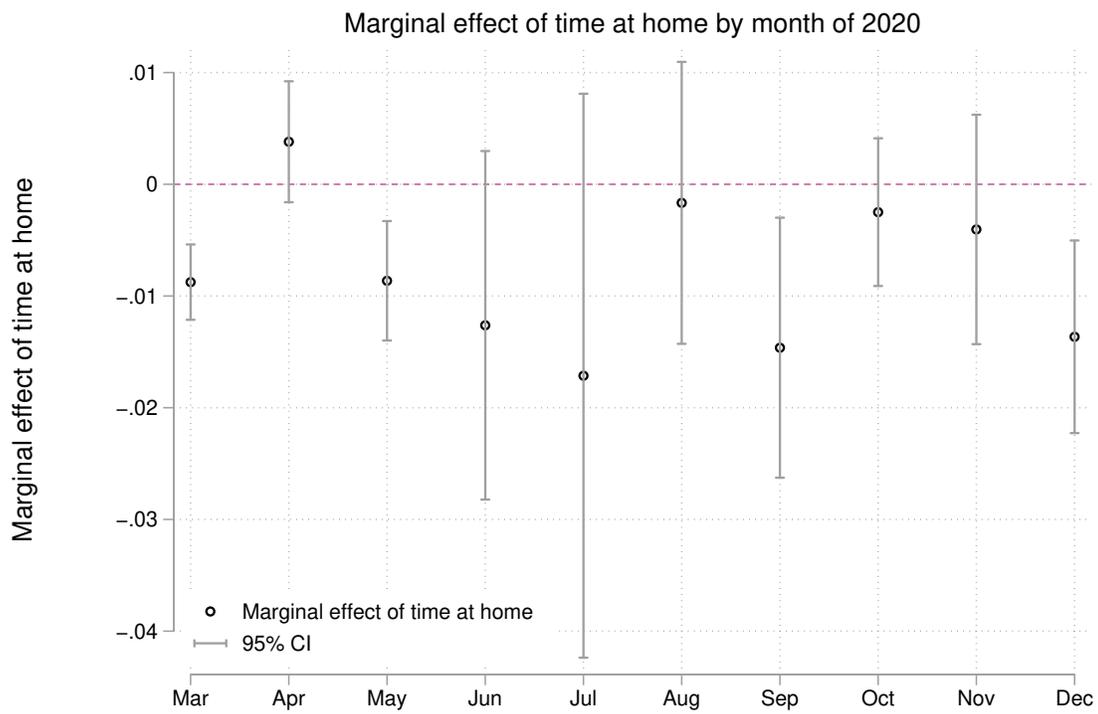


Figure 6: Marginal effect of an additional hour of time spent at home on the gap between predicted and actual electricity consumption in PJM.

Table 3: Monthly utility panel estimates

	(1)	(2)	(3)
	Residential	Commercial	Industrial
March	-9.21*	-8.65*	-0.30
	(2.75)	(3.03)	(1.49)
April	5.21*	-22.2*	-13.5*
	(2.38)	(3.37)	(2.51)
May	7.02*	-27.2*	-14.9*
	(2.54)	(4.07)	(2.55)
June	14.6*	-16.4*	-12.1*
	(3.57)	(3.42)	(2.23)
July	32.0*	-4.42*	-5.95*
	(2.88)	(2.10)	(1.57)
August	24.7*	-10.4*	-6.43*
	(4.08)	(4.33)	(1.96)
September	-0.0018	-15.5*	-8.04*
	(2.15)	(3.47)	(2.10)
October	3.80*	-11.2*	-5.15*
	(1.43)	(2.90)	(1.71)
November	14.6*	-4.74	-2.09
	(5.72)	(3.21)	(1.54)
December	21.3*	-4.67*	-4.44*
	(3.07)	(2.22)	(1.96)
State-utility FE	Yes	Yes	Yes
Month-of-year indicators	Yes	Yes	Yes
Weather controls	Yes	Yes	Yes
Observations	35354	35351	35347

Each observation is monthly electricity consumption by customer class at the state-utility level, with units in gigawatt hours. Weather controls include monthly high, low, and average temperatures as well as average monthly precipitation across each utility's counties. 95% confidence intervals constructed from standard errors accounting for multi-way clustering by state-utility and by month. Observations vary due to some utilities not reporting commercial or industrial consumption. * p-value < 0.05.

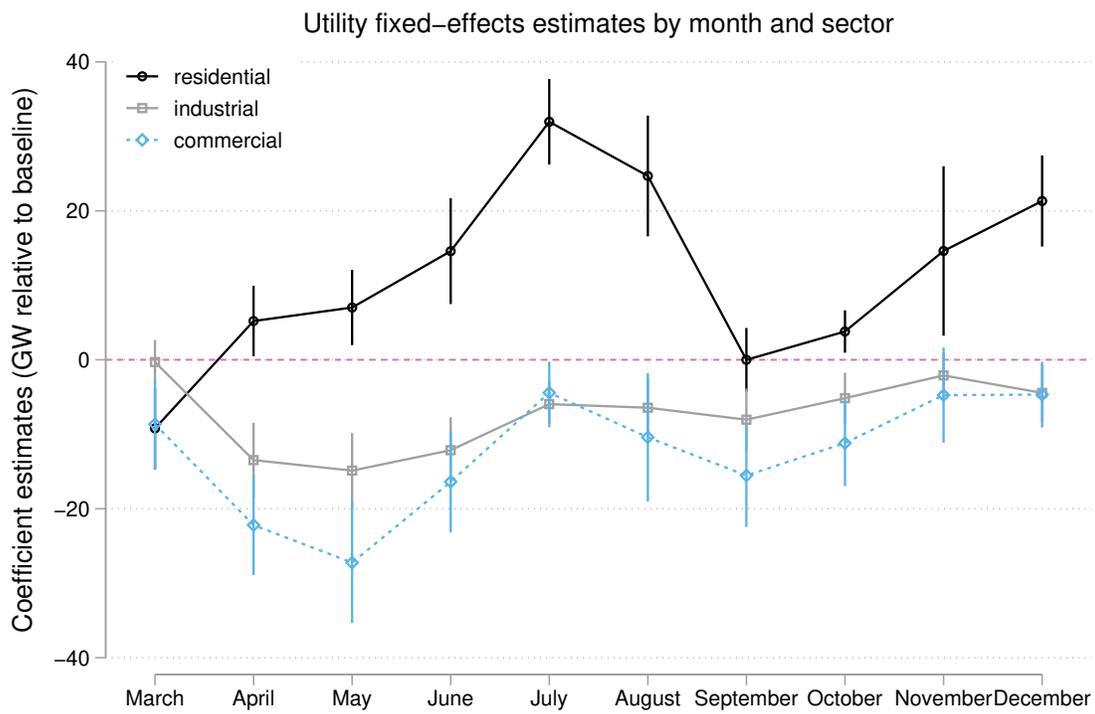


Figure 7: Estimates by customer class at the state-utility level from equation 3. 95% confidence intervals account for multi-way clustering by utility-state and by month.

years. I tested the equivalence of the coefficients across sectors using two-way Z-tests for equivalence of coefficients across models (Clogg et al., 1995; Paternoster et al., 1998). The residential coefficients are statistically different than the commercial and industrial coefficients in each month with the exception of March’s residential and industrial coefficients, which I fail to reject any difference between. The commercial and industrial coefficients are only statistically different in March through May. In June through December, I fail to reject a null hypothesis that industrial and commercial consumption changed differentially.

These patterns are consistent with a slight recovery in commercial and industrial activities beginning in July and August, but do not provide evidence that businesses were making up for missed production from earlier in the pandemic. It is still possible that this “making-up” behavior occurred in some sectors (such as manufacturing and mining), while the service sector continued to consume less electricity relative to normal, and while other sectors consumed the same or more electricity (such as the information technology sector). Unfortunately, with aggregated data it is not possible to test this follow-up hypothesis.

There is more support for the hypothesis that summer cooling of residential homes with individuals working from home during the hottest hours is less energy efficient. Despite controlling for outdoor temperature in the analysis, it is possible that the relationship between temperature and residential electricity consumption changed in 2020 relative to previous years and that this unobserved change can account for some of the increased summer consumption. I investigate this possibility by regressing hourly zone-level electricity consumption from June-August on indicators for five-degree temperature intervals, controlling for fixed effects by zone, year, month of year, and day of week times hour of day as described in equation 4.

Figure 8 plots the coefficient estimates on temperature interval interacted with the indicator variable for March-December 2020 from regression 4. These estimates represent the difference in the relationship between temperature and electricity load during the COVID-19 period compared to during 2014-2019. The confidence intervals on most of the estimates contain zero, except for the 80-85°F interval, but an F-test of the joint hypothesis that the

coefficients are zero returns a p value < 0.0001 . The estimates in the 95+ interval and the 5-10 and 10-15 intervals are very imprecise due to observing very few days at these extremes.

To gain a better understanding of how the relationship between temperature and load differs during the COVID-19 period, I plot the marginal effect on log load of being in a temperature interval relative to the omitted interval of 55-60°F in figure 9 (thus, the 55-60°F estimate is zero and the other estimates are relative to it). Higher temperatures result in more electricity consumption while low temperatures result in slightly less electricity consumption in 2020 than in previous years, though these results do not hold in the 95+ interval and the 5-10 and 10-15 intervals. Again, a test of joint significance rejects that the marginal effects are the same, but at the individual level, only the effect of being in the 80-85°F interval are statistically different. Thus, the relationship between temperature and electricity consumption changed with the pandemic; however, the difference is largely manifested at high temperatures. While this difference may account for some of the increased electricity use in July and August, it cannot explain the increase of similar magnitude in December.

5 Implications and conclusions

Using hourly electricity consumption data in a difference-in-predicted-differences strategy, this article shows that while electricity consumption declined by 3.8, 3.8, and 2.7% in the first three months of the COVID-19 pandemic, electricity use was 3.5% higher in August 2020. Electricity consumption in September through November was roughly normal compared to the predicted baseline, while consumption in December was 3.0% higher than the predicted baseline. Cell-phone location data on stay-at-home behavior show that more time spent at home decreases electricity consumption and that time spent at home substantially decreased after May. Nationwide monthly data on electricity consumption by load class reveals that commercial and industrial consumption was below its expected baseline from March-November 2020, while residential consumption was above its expected baseline,

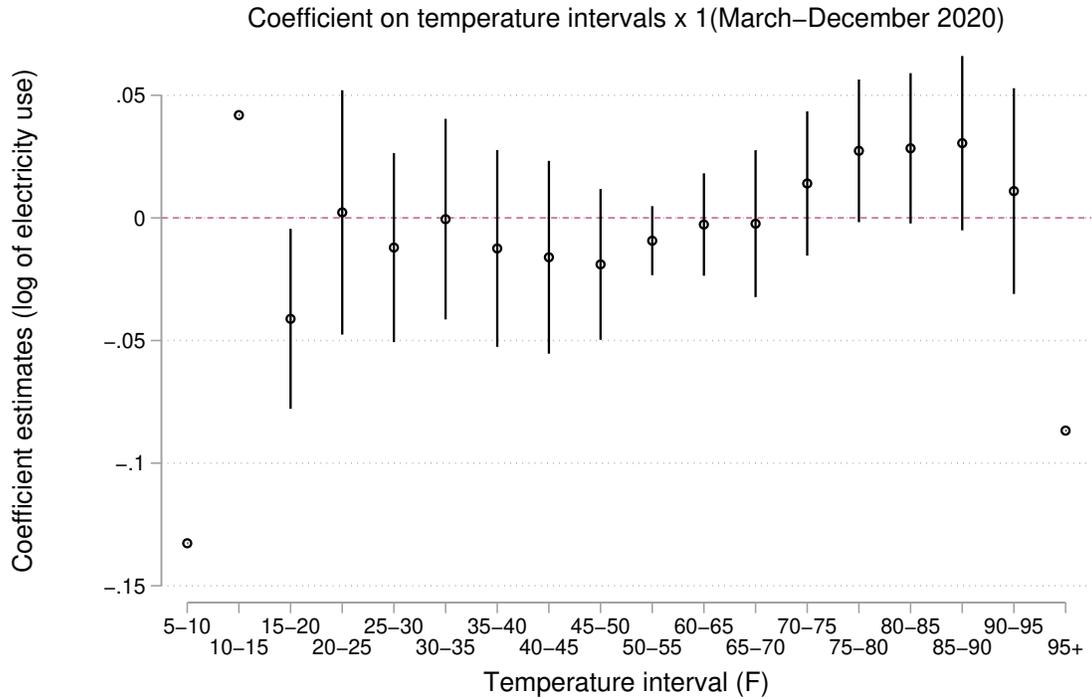


Figure 8: Coefficient estimates on temperature interval interacted with an indicator variable for March-December 2020 from regression 4. Temperature interval 55-60°F is the omitted category. Estimates represent the difference in the relationship between temperature and electricity load in 2020 versus 2014-2019 after controlling for fixed effects by zone and seasonality via indicators for year, month, and day of week times hour of day. 95% confidence intervals account for multi-way clustering by zone and by day. Confidence intervals suppressed on 5-10, 10-15, and 95+ intervals for scale. An F-test of the joint hypothesis that the coefficients are zero can be rejected with a $p\text{-value} < 0.000$.

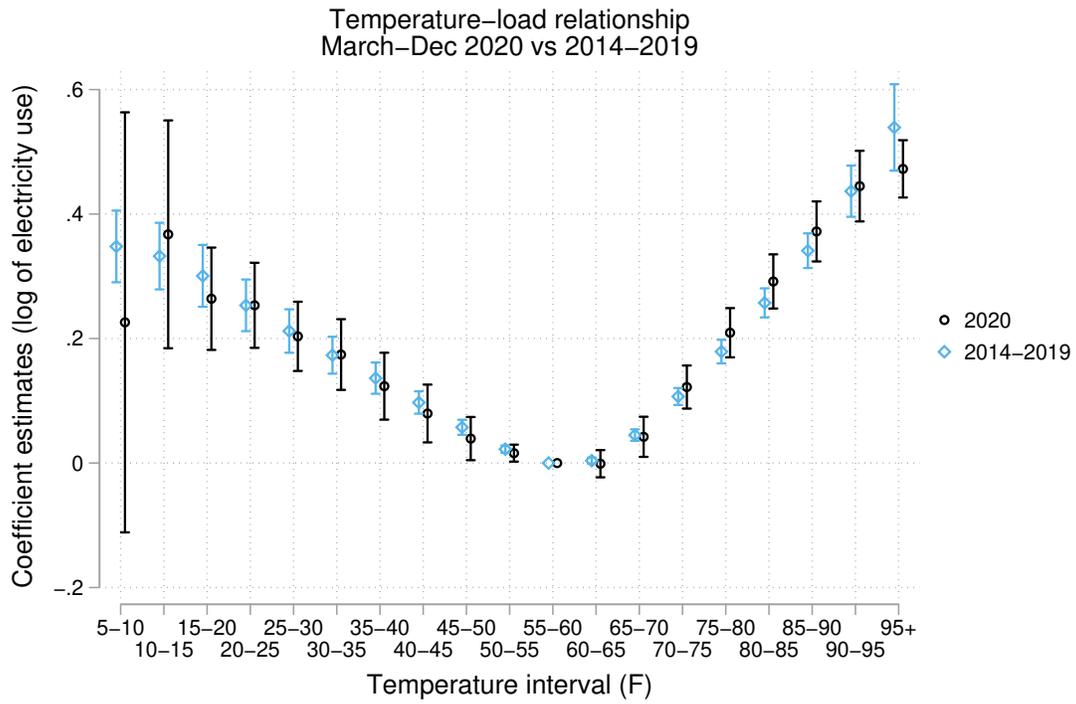


Figure 9: Marginal effect on log load of being in a temperature interval relative to the excluded 55-60°F interval. 95% confidence intervals account for multi-way clustering by zone and by day.

peaking in July. As a whole, the early reductions in electricity consumption were offset by the summer increases, with an overall effect of an increase in electricity consumption of 1.08%. The zone serving electricity to Chicago experienced consistent declines in electricity consumption, while zones covering rural areas saw the largest summer rebounds.

These empirical findings suggest that initial reductions in electricity consumption were offset by increased consumption by residential users. This increased consumption is at least partially due to cooling responses to high summer temperatures. When more people are at home during the day, home HVAC systems are used more, drawing additional electricity. This increased residential consumption more than made up for reductions in industrial and commercial consumption, resulting in the pandemic counterintuitively increasing electricity consumption in the United States during the hottest months of the year. This explanation cannot explain increased demand in December, as cold weather did not have as large of an effect on electricity consumption as in previous years.

Increased electricity consumption has important implications for the growing literature examining “silver linings” of the pandemic. To the extent that electricity generation contributes to local and global air and water pollution, the gains will be smaller than expected due to increased demand for cooling in the summer months. Future work in this area should focus on air and water quality improvements from reduced commuter traffic and should acknowledge that the pandemic did not uniformly decrease electricity consumption.

Early in the pandemic, some scholars noted that electricity consumption changes were a better real-time measure of macroeconomic conditions than traditional metrics such as quarterly earnings reports (e.g. Bui and Wolfers, 2020). This was likely true in the short run when the relationship between electricity consumption and its underlying causes remained the same. At that time, a change in electricity consumption only reflected the reduction in the level of economic activity after controlling for other underlying determinants of electricity consumption. In the long run, behavior adapted to new conditions and the relationship between electricity consumption and its underlying causes changed. For example, a change in electricity consumption then reflects differences in the level of economic activity and

the new relationship between temperature and electricity consumption. To use electricity consumption as a valid metric of macroeconomic health, the proposed metric should account for changes in the relationship between electricity consumption and its primitives in a full decomposition approach.

These results have several key policy implications for future pandemics or for a work-from-home environment where households spend substantially more time at home. First, policies targeting residential electricity consumption will target a larger proportion of load in these scenarios. More residential load means residential energy efficiency has additional value relative to previously, which may increase the payoff of residential energy efficiency programs. Similarly, residential demand response programs will also be more important for reducing peak demand. At the same time, commercial electricity consumption may be low-cost to conserve and more elastic in industries amenable to work-from-home. Electricity prices already incentivize reductions in commercial demand, but time-varying electricity prices may be able to induce greater reductions in commercial electricity demand. Research using meter-level consumption data should be conducted to investigate the extent to which “the lights were left on” in empty commercial buildings during the pandemic to determine the extent to which policy can reduce this consumption further.

References

- Abrams, D. S. (2021). COVID and crime: An early empirical look. *Journal of Public Economics* 194, 104344.
- Anastasio, E. (2020). Short-term load forecasting. Technical report, PJM. Presentation, January 9, 2020, <https://www.pjm.com/-/media/committees-groups/committees/oc/2020/20200109/20200109-item-11-load-forecasting-education.ashx> accessed 13-01-2021.
- Auray, S. and A. Eyquem (2020). The macroeconomic effects of lockdown policies. *Journal of Public Economics* 190, 104260.

- Baron, E. J., E. G. Goldstein, and C. T. Wallace (2020). Suffering in silence: How COVID-19 school closures inhibit the reporting of child maltreatment. *Journal of Public Economics* 190, 104258.
- Benatia, D. (2020). Electricity markets under lockdown: Insights from New York. *IAEE Energy Forum Special Covid-19 Edition 2020*, 55–58.
- Benatia, D. (2022). Ring the alarm! electricity markets, renewables, and the pandemic. *Energy Economics* 106.
- Benatia, D. and S. Gringas (2021). Reaching new lows? The pandemic’s consequences for electricity markets. USAEE Working Paper No. 20-454, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3623966, accessed June 2022.
- Bover, O., N. Fabra, S. García-Uribe, A. Lacuesta, and R. Ramos (2022). Firms and households during the pandemic: what do we learn from their electricity consumption? *The Energy Journal* (forthcoming).
- Brewer, D. (2020). The economic costs of forecasting errors in the PJM interconnection due to the COVID-19 quarantine. *IAEE Energy Forum Special Covid-19 Edition 2020*, 116–119.
- Brodeur, A., N. Cook, and T. Wright (2021). On the effects of COVID-19 safer-at-home policies on social distancing, car crashes and pollution. *Journal of Environmental Economics and Management* 106, 102427.
- Bui, Q. and J. Wolfers (2020). Another way to see the recession: Power usage is way down. *The New York Times: The Upshot*, April 8, 2020. <https://www.nytimes.com/interactive/2020/04/08/upshot/electricity-usage-predict-coronavirus-recession.html>, accessed July 2022.
- Burlig, F., C. Knittel, D. Rapson, M. Reguant, and C. Wolfram (2020). Machine learning

- from schools about energy efficiency. *Journal of the Association of Environmental and Resource Economists* 7(6), 1181–1217.
- Cameron, A. C., J. B. Gelbach, and D. L. Miller (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics* 29(2), 238–249.
- Carroll, J., K. Conway, A. Shannon, and E. Denny (2020). The impact of COVID-19 restrictions on electricity demand: Evidence from two jurisdictions. *IAEE Energy Forum Special Covid-19 Edition 2020*, 51–54.
- Cicala, S. (2022). Powering work from home. *Journal of Urban Economics: Insights (forthcoming)*.
- Cicala, S., S. Holland, E. Mansur, N. Muller, and A. Yates (2020). Expected health effects of reduced air pollution from COVID-19 social distancing. NBER Working Paper Series 27135, National Bureau of Economic Research.
- Clogg, C. C., E. Petkova, and A. Haritou (1995). Statistical methods for comparing regression coefficients between models. *American Journal of Sociology* 100(5), 1261–1293.
- Couch, K. A., R. W. Fairlie, and H. Xu (2020). Early evidence of the impacts of COVID-19 on minority unemployment. *Journal of Public Economics* 192, 104287.
- Dang, H.-A. H. and T.-A. Trinh (2021). Does the COVID-19 lockdown improve global air quality? new cross-national evidence on its unintended consequences. *Journal of Environmental Economics and Management* 105, 102401.
- Fabra, N., A. Lacuesta, and M. Souza (2022). The implicit cost of carbon abatement during the COVID-19 pandemic. *European Economic Review (forthcoming)*.
- Forsythe, E., L. B. Kahn, F. Lange, and D. Wiczer (2020). Labor demand in the time of COVID-19: Evidence from vacancy postings and UI claims. *Journal of Public Economics* 189, 104238.

- Gillingham, K., C. Knittel, J. Li, M. Ovaere, and M. Reguant (2020). The short-run and long-run effects of Covid-19 on energy and the environment. *Joule* 4(7), 1337 – 1341.
- Graf, C., F. Quaglia, and F. A. Wolak (2021). (Machine) learning from the COVID-19 lockdown about electricity market performance with a large share of renewables. *Journal of Environmental Economics and Management* 105, 102398.
- Hensvik, L., T. Le Barbanchon, and R. Rathelot (2021). Job search during the COVID-19 crisis. *Journal of Public Economics* 194, 104349.
- Imbens, G. and A. Abadie (2008). On the failure of the bootstrap for matching estimators. *Econometrica* 76(6), 1537–1557.
- Irwin, N. B. and M. R. Livy (2021). Measuring environmental (dis)amenity value during a pandemic: Early evidence from maryland. *Journal of Environmental Economics and Management* 106, 102417.
- Isphording, I. E. and N. Pestel (2021). Pandemic meets pollution: Poor air quality increases deaths by COVID-19. *Journal of Environmental Economics and Management* 108, 102448.
- Jacob, D. (2020). Cross-fitting and averaging for machine learning estimation of heterogeneous treatment effects. arXiv pre-print available at <https://arxiv.org/abs/2007.02852>, accessed October 2022.
- Kong, E. and D. Prinz (2020). Disentangling policy effects using proxy data: Which shut-down policies affected unemployment during the COVID-19 pandemic? *Journal of Public Economics* 189, 104257.
- Leslie, E. and R. Wilson (2020). Sheltering in place and domestic violence: Evidence from calls for service during COVID-19. *Journal of Public Economics* 189, 104241.
- Narajewski, M. and F. Ziel (2020). Changes in electricity demand pattern in europe due to COVID-19 shutdowns. *IAEE Energy Forum Special Covid-19 Edition 2020*, 44–47.

National Centers for Environmental Information (2014-2020a). Integrated surface database. Data version current as of February, 11, 2021. <https://www.ncdc.noaa.gov/isd>.

National Centers for Environmental Information (2014-2020b). NOAA's Climate Divisional Database (nCLIMDIV). <https://www.ncei.noaa.gov/metadata/geoportal/rest/metadata/item/gov.noaa.ncdc:C00005/html>.

Paternoster, R., R. Brame, P. Mazerolle, and A. Piquero (1998). Using the correct statistical test for the equality of regression coefficients. *Criminology* 36(4), 859–866.

Percy, S. and B. Mountain (2020). Covid-19 and social distancing: Does it show up in the demand for electricity? *IAEE Energy Forum Special Covid-19 Edition 2020*, 38–39.

Persico, C. L. and K. R. Johnson (2021). The effects of increased pollution on COVID-19 cases and deaths. *Journal of Environmental Economics and Management* 107, 102431.

PJM (2014-2020). Hourly load: Metered. https://dataminer2.pjm.com/feed/hrl_load_metered/definition.

PJM (2018). PJM manual 19: Load forecasting and analysis. Technical Report Revision 33, PJM.

PJM (2019). 2019 pjw annual report. <https://www.pjm.com/-/media/about-pjm/newsroom/annual-reports/2019-annual-report.ashx?la=en>.

United States Energy Information Administration (2014-2020a). Annual Electric Power Industry Report, Form EIA-861 detailed data files. <https://www.eia.gov/electricity/data/eia861/>.

United States Energy Information Administration (2014-2020b). Form EIA-861M (formerly EIA-826) detailed data. <https://www.eia.gov/electricity/data/eia861m/>.

Appendices

A Goodness of fit

To assess whether a regression approach, matching at the zone level, or matching at the aggregate PJM level is a better approach, I predict 2019 electricity load using load from 2014 through 2018 using each method. I consider three approaches: using linear regression to forecast load for each zone, matching at the aggregate PJM level, and matching at the zone level. This is a type of cross-validation exercise to select which approach performs better and to provide assurance that the chosen approach provides accurate predictions (and thus a valid counterfactual).

The regression approach estimates the following regression separately for each zone using least-squares using the data from 2014-2018:

$$MW_{zt} = \delta_{zymdh} + \beta weather_{zt} + \varepsilon_{zt}, \quad (5)$$

where the dependent variable is the natural log of the megawatts of electricity used in zone z in hour t . The term δ_{zymdh} represents indicators for month of year, day of week times hour of day, and an intercept. The variable $weather_{zt}$ includes controls for temperature and precipitation, and ε_{zt} is the residual. I then take the fitted values from the regression and forecast load for 2019.

The first matching algorithm performs the matching at the aggregate PJM level, using weather in all zones as variables to match on (within month-of-year, hour-of-day, and day-of-week). For each hour of 2020, the algorithm finds the hour within the same month-of-year, hour-of-day, and day-of-week with the most similar temperature and precipitation as measured using the Mahalanobis metric. The Mahalanobis metric is the generalized Euclidean distance between two vectors X and Y : $\sqrt{(X - Y)V^{-1}(X - Y)'} where V^{-1} is the variance matrix. It scales each dimension by the variance and takes the covariance structure$

of the data into account to control for correlation between two dimensions. Denote X_t the vector of temperature and precipitation for hour t in 2020 and X_ℓ the vector of temperature and precipitation for hour $\ell \in L_t$, where L_t is the set of hours in the same month-of-year, hour-of-day, and day-of-week in 2014-2018. The matching hour $m(t)$ solves

$$m(t) = \arg \min_{\ell \in L_t} \sqrt{(X_t - X_\ell)V^{-1}(X_t - X_\ell)'} \quad (6)$$

The result is a single time-series of hourly counterfactual PJM electricity consumption $\hat{Y}_t = MW_{m(t)}$ where $MW_{m(t)}$ is the electricity consumption matched to hour t by applying equation 6. I refer to this algorithm as the time-series matching algorithm.

The second matching algorithm performs the same matching at the zone level, using same-zone weather to match on (within month-of-year, hour-of-day, and day-of-week). The result is a panel of zone-level hourly electricity consumption, $\hat{Y}_{zt} = MW_{zm(t)}$ where $MW_{zm(t)}$ is the zone-level electricity consumption matched to hour t by applying equation 6 zone-by-zone. I refer to this algorithm as the panel matching algorithm.

The cross-validation procedure evaluates the goodness-of-fit of the three prediction approaches by predicting electricity consumption in 2019 using training data from 2014-2018. This simulates the prediction task of predicting a year ahead into 2020—the goal is to use the prediction method that performs best at predicting a year ahead.¹³ Figure 10 displays daily metered 2019 load and the forecasted load using the regression approach. The root-mean-squared prediction error for the regression approach is 81.4, corresponding to an absolute percent error of 2.9%. Figure 11 displays daily metered 2019 load and the matched load using the nearest-neighbors algorithm at the PJM level. The root-mean-squared prediction

¹³In a canonical prediction approach, one separates the training sample randomly into a training and hold-out evaluation sample (potentially several times in a k -fold cross validation procedure). The canonical procedure assumes that the training and prediction data are randomly assigned. In our context, the prediction set contains data the following year only, which differs from the canonical case. Evaluating the goodness of fit on a random split of data would be an artificially easier prediction task, because same-year electricity load would then be used for prediction. For example, predicting August 16, 2019 load using 2014-2018 load data is much more difficult than predicting August 16, 2019 load with a training sample that includes August 15, 2019 load. This task is artificially easier than predicting the next year’s load and would lead to an overstatement of the performance of the matching algorithm and potentially lead to the choice of an inferior prediction approach.

error for the time-series matching algorithm is 170.7, corresponding to an absolute percent error of 6.3%. Figure 12 displays daily metered 2019 load and the matched load using the nearest-neighbors algorithm at the zone level. The root-mean-squared prediction error for the panel matching algorithm is 71.3, corresponding to an absolute percent error of 2.5%.¹⁴ Comparing these figures and the corresponding root-mean-squared prediction errors, it is clear that matching at the zone level provides a much better prediction. For this reason, I prefer the results from the zone-level matching algorithm and use it throughout the paper. Finally, mean daily load in 2019 was 2,155 GWh (std dev = 287), so the prediction errors are quite small relative to the typical variation in load.

¹⁴As an additional check suggested by a referee, I examined the cross-validation performance of the algorithms with an additional layer of cross-fitting in which I split the training sample into five groups, train each prediction algorithm on each group, and take the median prediction (Jacob, 2020). This results in a RMSE of 294.9 (absolute percent deviation 12.9%) for the regression approach, 175.0 (absolute percent deviation 6.3%) for the time-series matching algorithm, and 100.6 (absolute percent deviation 3.4%) for the panel matching algorithm. Thus, the preferred approach by this metric remains the panel matching algorithm.

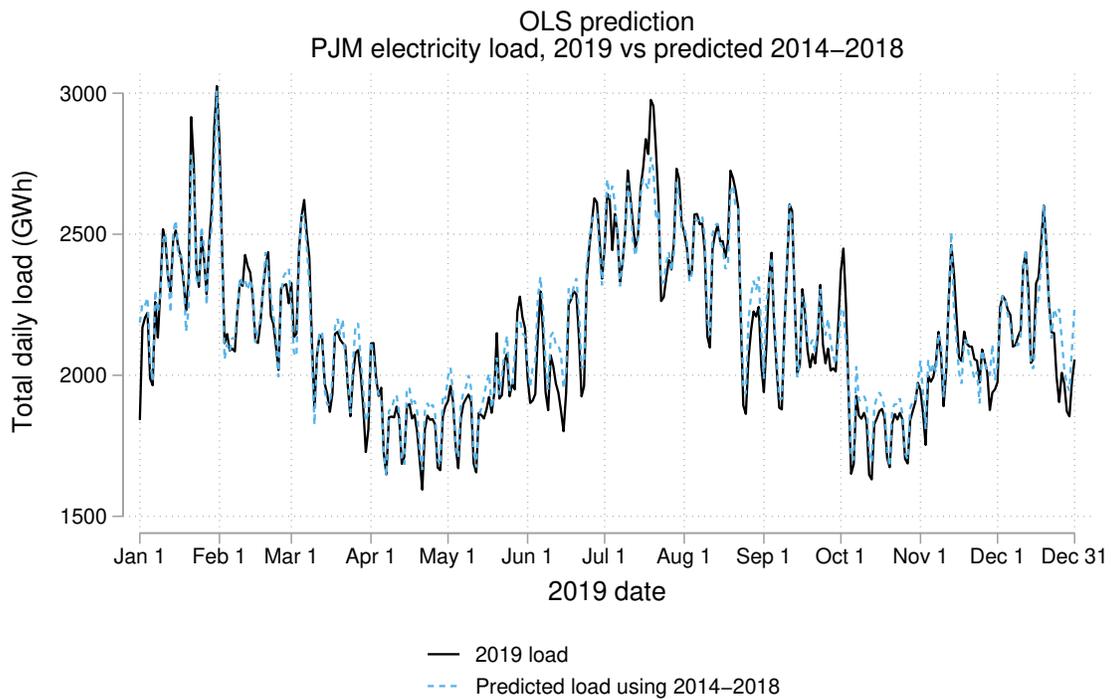


Figure 10: 2019 electricity load predicted using a linear regression of 2014-2018 load on a third-degree polynomial of temperature, precipitation, and indicators for month of year, and hour of day times day of week. Figure displays the series in units of total daily load for ease of interpretation.

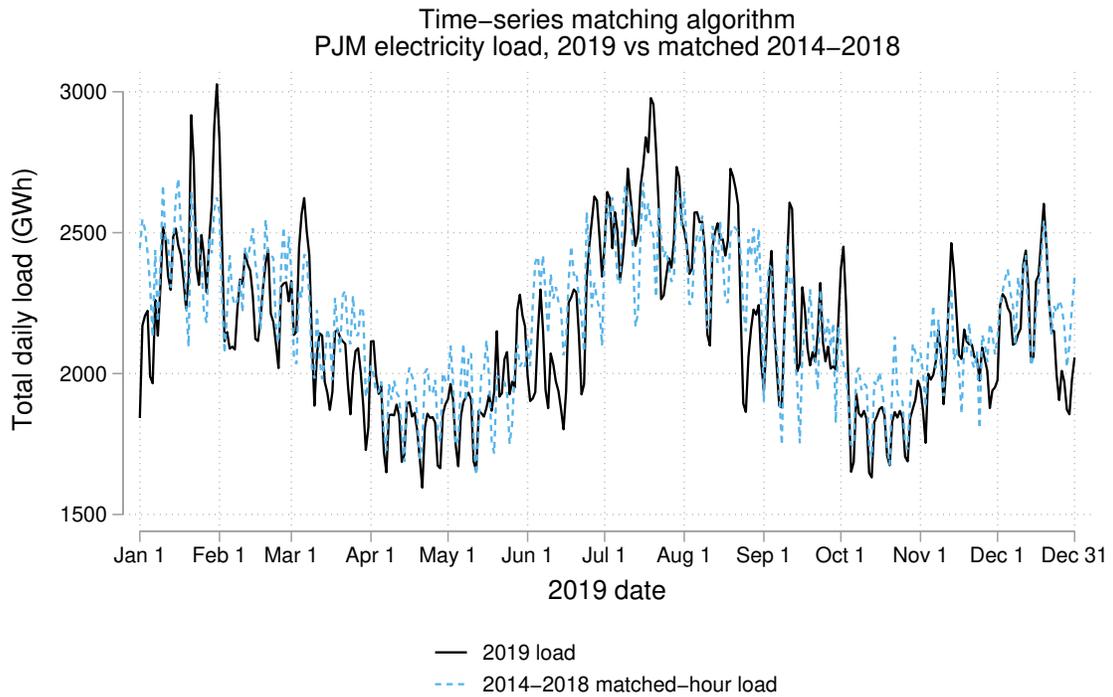


Figure 11: 2019 electricity load matched to the most-similar weather hour from 2014-2018, matching exactly on the month-of-year, hour-of-day, and day-of-week at the aggregate PJM level. Figure displays the series in units of total daily load for ease of interpretation.

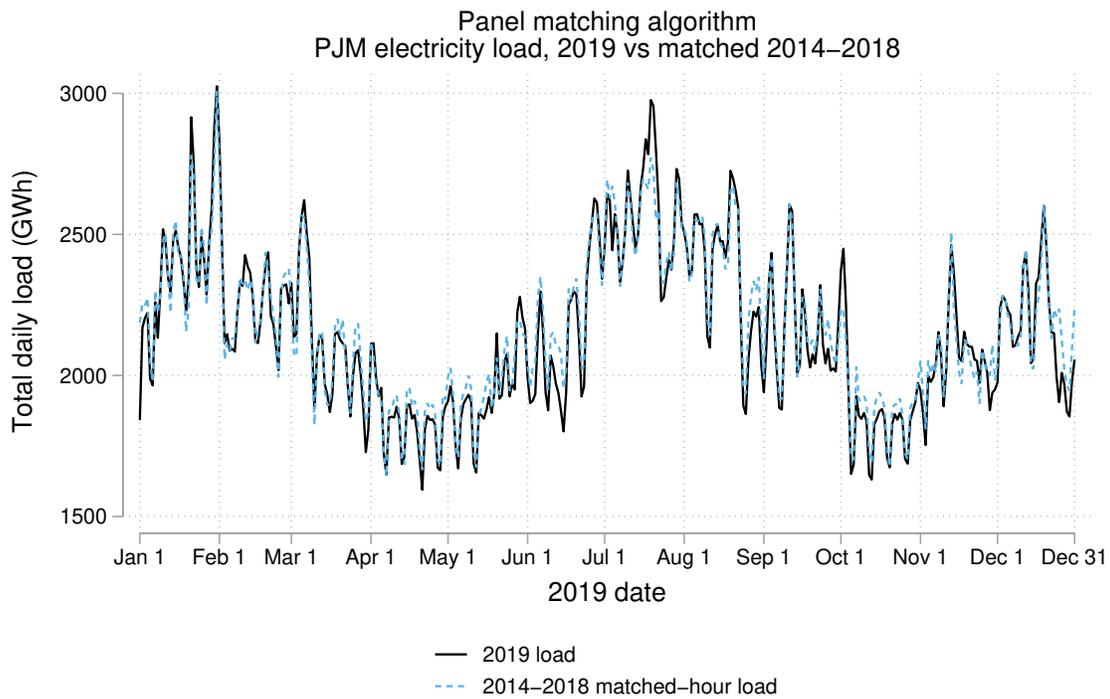


Figure 12: 2019 electricity load matched to the most-similar weather hour from 2014-2018, matching within zone exactly on the month-of-year, hour-of-day, and day-of-week. Figure displays the series in units of total daily load for ease of interpretation.

B Alternative regression analysis

In this section, I provide a regression-based analysis that is an alternative to the difference-in-predicted differences estimation and triple-differences estimation presented in the main text. I present this analysis to corroborate and verify the earlier matching results as a sort of robustness check. Given the matching procedure’s significant advantage in modeling counterfactual electricity consumption demonstrated in appendix section A, I prefer the results in the main text over the results below.

First, I corroborate the difference-in-predicted differences estimates using a regression that models electricity use before and after the pandemic. Using the full hourly consumption data from PJM, I estimate the following model using ordinary least squares:

$$\ln(MW_{zt}) = \delta_{zymdh} + \sum_{m=3}^{12} \beta_m post_t 1(month_t = m) + \gamma weather_{zt} + \varepsilon_{zt}, \quad (7)$$

where the dependent variable is the natural log of the megawatts of electricity used in zone z in hour t . The term δ_{zymdh} represents controls for fixed effects by zone z , indicators for year y , month of year m , day of week times hour of day dh , and an intercept. The variable $post_t$ is equal to one for all hours after the beginning of the COVID-19 period, and its coefficient β_m varies by month. $weather_{zt}$ includes controls for precipitation and a third-degree polynomial of temperature, and ε_{zt} is the residual. Given the log-linear specification, estimates of β_m are approximately percent changes in electricity use from month-to-month during the pandemic holding the control variables fixed. As in the main analysis, I allow the treatment date to vary by specification and include specifications that exclude observations in the first week and two weeks of the pandemic.

Table 4 displays the results of the regression equation 7. As in the main analysis, the results indicate that electricity demand was down in the early months of the pandemic before increasing during the summer months and again in December. The estimates are similar in sign and magnitude, though the regression results tend to be larger than those in the main results. While I favor the difference-in-predicted differences estimates given the matching

algorithm’s improved mean-squared error, this alternative analysis does have the advantage that the standard errors account for all sources of sampling uncertainty.

Next, I corroborate the triple-differences estimates of the effect of stay-at-home behavior using a regression-based difference-in-differences estimate. I estimate the following regression on the sample from 2019-2020 (the dates when social distancing data are available) using least squares:

$$\begin{aligned} \ln(MW_{zt}) = & \delta_{zymdh} + \sum_{m=3}^{12} \beta_m post_t 1(month_t = m) + \eta_m home_{zt} \cdot post_t 1(month_t = m) \quad (8) \\ & + \eta_0 home_{zt} + \kappa devices_{zt} + \omega weather_{zt} + e_{zt}, \end{aligned}$$

where the dependent variable is the natural log of the megawatts of electricity used in zone z in hour t . The term δ_{zymdh} represents controls for fixed effects by zone, indicators for month of year and day of week times hour of day, and an intercept. The variable $post_t$ is equal to one for all hours after the beginning of the COVID-19 period, and $home_{zt}$ is the average median number of hours spent at home on each day. The variable $devices_{zt}$ is a count of devices observed in each zone and controls for the purchase of new mobile devices, $weather_{zt}$ includes controls for precipitation and a cubic polynomial in temperature, and e_{zt} is the residual. I allow the coefficients on the interaction between $home_{zt}$ and $post_t$ to vary by month in the preferred specification. The estimates of interest are the sum of the coefficients on $home_{zt}$ for each month $\eta_0 + \eta_m$, which measure the marginal effect of stay-at-home behavior on electricity load throughout the pandemic.

Table 5 displays the estimates from regression equation 8. As in the main analysis, the effect of time spent at home is negative on average with negative or statistically-zero month-specific estimates. The estimates differ somewhat from the main results displayed in table 2, but confirm the general finding that time spent at home reduces electricity consumption. Figure 13 displays the coefficient estimates from column (4) graphically to compare to figure 4 in the main text.

Table 4: Regression estimates: overall effect

	(1)	(2)	(3)	(4)	(5)
Post × March	-0.025*	-0.042*	-0.050*	-0.041*	-0.052*
	(0.009)	(0.008)	(0.008)	(0.008)	(0.009)
Post × April	-0.043*	-0.045*	-0.044*	-0.043*	-0.043*
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Post × May	-0.057*	-0.060*	-0.058*	-0.057*	-0.057*
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Post × June	-0.001	-0.003	-0.002	-0.000	-0.000
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Post × July	0.053*	0.050*	0.052*	0.053*	0.053*
	(0.012)	(0.012)	(0.012)	(0.012)	(0.012)
Post × August	0.035*	0.032*	0.034*	0.035*	0.035*
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
Post × September	-0.012	-0.014	-0.013	-0.011	-0.011
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Post × October	-0.005	-0.008	-0.006	-0.005	-0.005
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Post × November	-0.007	-0.010	-0.008	-0.007	-0.007
	(0.009)	(0.008)	(0.008)	(0.009)	(0.009)
Post × December	0.029*	0.027*	0.028*	0.029*	0.029*
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Overall effect	-0.30%	-0.71%	-0.53%	-0.38%	-0.36%
Treatment date	March 1	March 8	March 15	March 9	March 16
Excluded dates				March 1-8	March 1-15
Weather	Y	Y	Y	Y	Y
Zone FE	Y	Y	Y	Y	Y
Year, Month, DOW X Hour	Y	Y	Y	Y	Y
Observations	1288539	1288539	1288539	1284507	1280979

The outcome variable is logged hourly electricity consumption at the zone level in PJM from 2014-2020. Weather variables include a cubic polynomial of hourly temperature and a control for precipitation. 95% confidence intervals constructed from standard errors accounting for multi-way clustering by zone and by day. * p-value < 0.05.

Table 5: Regression estimates: Stay-at-home data

	(1)	(2)	(3)	(4)
Post	0.021 (0.032)	0.021 (0.032)		
Home hours	-0.009* (0.002)	-0.009* (0.002)	-0.008* (0.003)	-0.008* (0.003)
Post \times Home hours	-0.001 (0.003)	-0.001 (0.003)		
Device count (100k devices)		0.000 (0.006)		0.001 (0.006)
Post \times March \times Home hours			-0.004 (0.003)	-0.004 (0.003)
Post \times April \times Home hours			0.006 (0.004)	0.006 (0.004)
Post \times May \times Home hours			-0.008 (0.005)	-0.008 (0.005)
Post \times June \times Home hours			0.014 (0.010)	0.014 (0.010)
Post \times July \times Home hours			0.032* (0.015)	0.032* (0.015)
Post \times August \times Home hours			0.020 (0.012)	0.020 (0.012)
Post \times September \times Home hours			-0.005 (0.010)	-0.005 (0.011)
Post \times October \times Home hours			0.004 (0.006)	0.004 (0.006)
Post \times November \times Home hours			-0.001 (0.006)	-0.001 (0.006)
Post \times December \times Home hours			-0.016* (0.006)	-0.016* (0.006)
Post X Month			Y	Y
Weather	Y	Y	Y	Y
Zone FE	Y	Y	Y	Y
Year, Month, DOW X Hour	Y	Y	Y	Y
Observations	350841	350841	350841	350841

The outcome variable is logged hourly electricity consumption at the zone level in PJM from 2019-2020. Weather variables include a cubic polynomial of hourly temperature and a control for precipitation. 95% confidence intervals constructed from standard errors accounting for multi-way clustering by zone and by day. * p-value < 0.05.

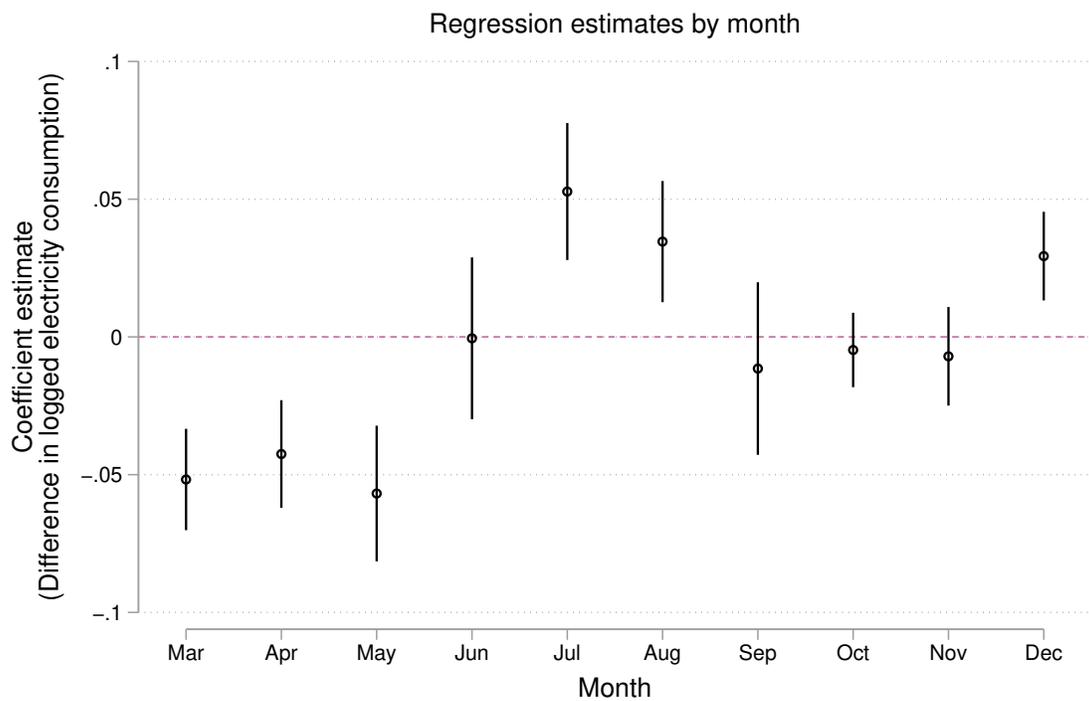


Figure 13: Estimates from column (5) of table 4 of the regression model specified in equation 7. 95% confidence intervals account for multi-way clustering by zone and by day. Estimates are the monthly difference in logged electricity consumption before and after March 1st, 2020 after omitting observations from March 1-15, and are approximately the percentage difference in electricity consumption relative to baseline consumption in past years.

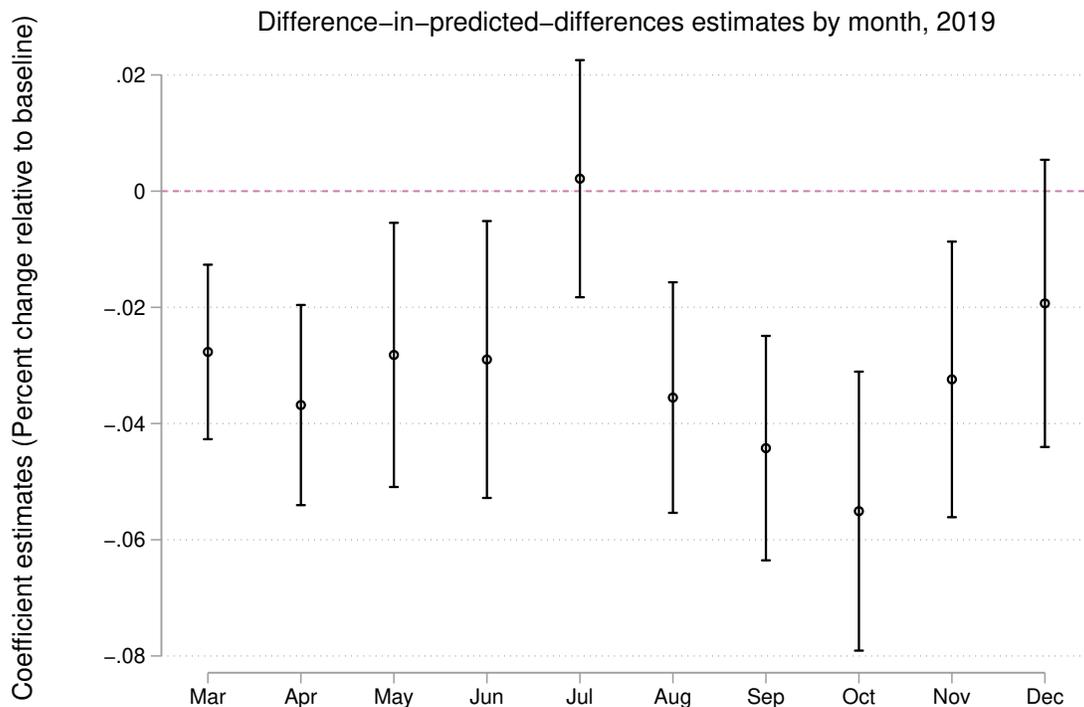


Figure 14: Estimates of the change in electricity consumption in 2019 relative to previous years. Estimates are the monthly difference in logged electricity consumption before and after March 1st, 2020 after omitting observations from March 1-15, and are approximately the percentage difference in electricity consumption relative to baseline consumption in past years.

C Change in 2019 consumption

In this section, I replicate the difference-in-predicted differences analysis for 2019, estimating the change in electricity relative to baseline. Figure 14 shows that electricity consumption was lower than the baseline. This may reflect increases in energy efficiency or behind-the-meter solar generation. If, absent the pandemic, 2020 would have exhibited a decline in energy consumption due to energy efficiency or renewable generation that was instead not adopted, our estimates of the summer increases may be a lower bound of the potential effect of the pandemic.

D Holiday robustness check

In this section, I control for US federal holidays in the second-stage regression to test the robustness of the results to controlling for holidays. One may suspect that the matching algorithm’s prediction errors are larger during holidays in which electricity consumption is anomalous. Differences in these prediction errors may potentially result in spurious estimates. To investigate whether this is an issue, I estimate the following regression:

$$\frac{Y_{zt} - \hat{Y}_{zt}}{\hat{Y}_{zt}} = \delta_{zymdh} + \sum_{m=3}^{12} \beta_m post_t 1(month_t = m) + \gamma holiday_t + \varepsilon_{zt}, \quad (9)$$

where $holiday_t$ is a vector of indicator variables for each US federal holiday.¹⁵ I use the specification from column (5) of table 1 from the main text that omits March 1-15. Table 6 displays the results of this regression, which do not differ substantially from the results in the main text.

¹⁵Each indicator variable is equal to one during the holiday and day before and after. I include all federal holidays celebrated in 2019-2020: New Year’s Day, Martin Luther King Jr. Day, Memorial Day, Independence Day, Labor Day, Columbus Day, Veterans Day, Thanksgiving, and Christmas.

Table 6: Robustness check: Holidays

	(1)
Post × March	-0.031* (0.009)
Post × April	-0.035* (0.010)
Post × May	-0.024* (0.011)
Post × June	0.024 (0.014)
Post × July	0.032 (0.016)
Post × August	0.039* (0.012)
Post × September	0.020 (0.011)
Post × October	0.013 (0.009)
Post × November	-0.004 (0.009)
Post × December	0.034* (0.008)
Overall effect	1.07%
Treatment date	March 16
Excluded dates	March 1-15
Zone FE	Y
Year, Month, DOW X Hour	Y
Holiday indicators	Y
Observations	360822

Dependent variable is the percent difference between hourly metered electricity load and predicted load. Estimates are the monthly difference in predicted differences before and after March 15th, 2020 and when multiplied by 100 are interpreted as the percentage difference in electricity consumption relative to the baseline prediction. Standard errors clustered to account for multi-way clustering by zone and by day. * p-value < 0.05.

E Heat maps by zone

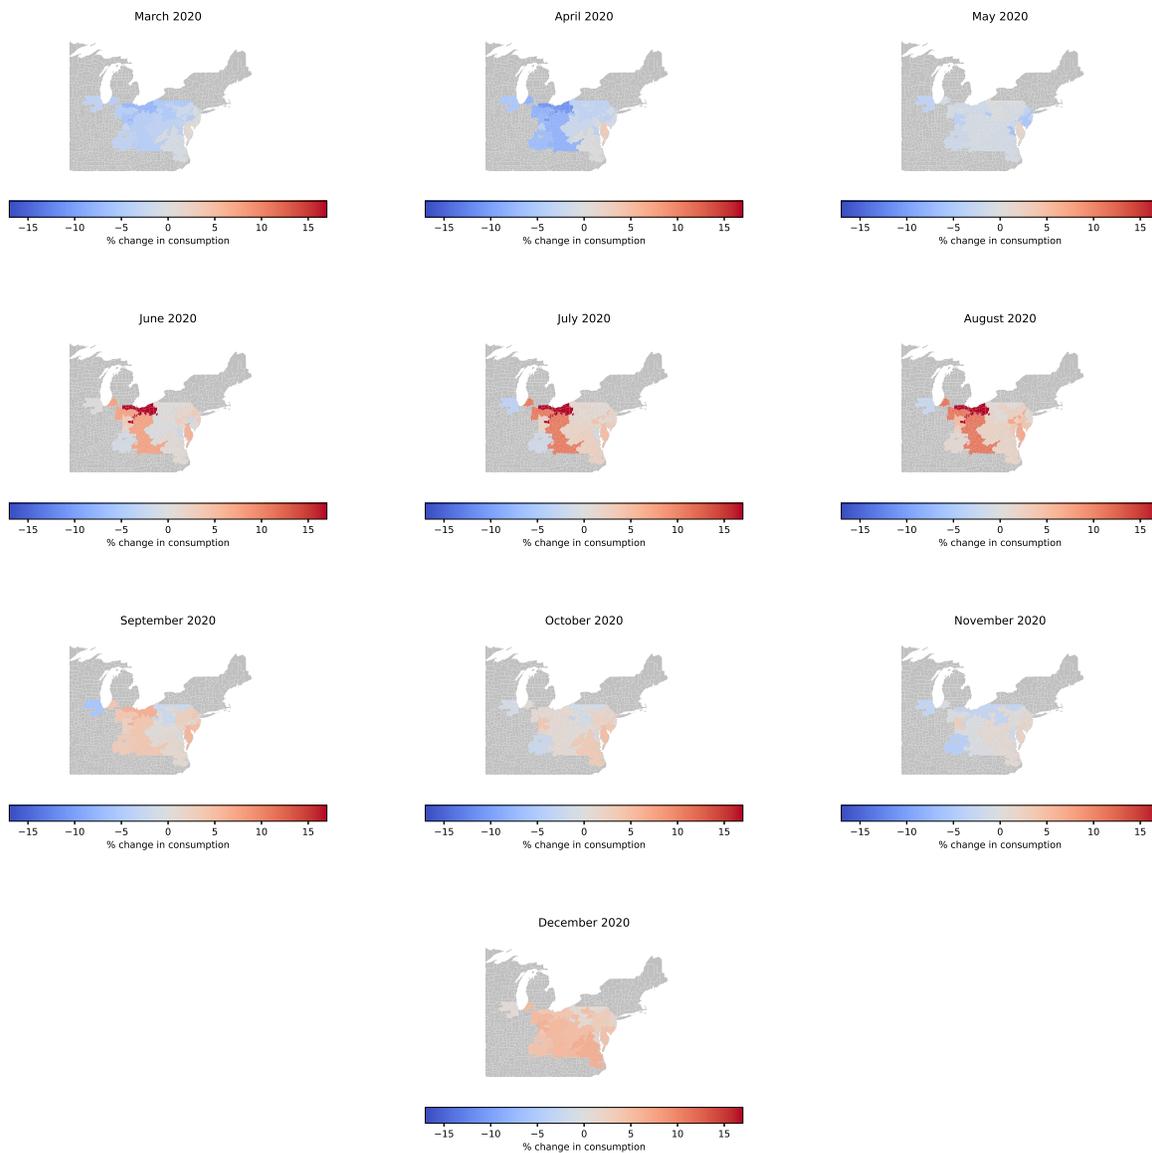


Figure 15: Heat maps of changes in electricity consumption by zone, estimated by applying ordinary least squares to each zone separately using equation 1.