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

Dylan Brewer*

February 19, 2021

Abstract

Using hourly electricity consumption data from the PJM Interconnection in the United States in a difference-in-predicted-differences strategy, this article shows that while in the first months of the COVID-19 pandemic total electricity consumption declined by 3.7-5.4% relative to a predicted counterfactual, in July and August 2020 electricity consumption was 2.9-4.6% higher than the predicted counterfactual. 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-November 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 almost perfectly offset by later increases, implying that any expected "silver lining" of decreased emissions from electricity production may be smaller than previously thought.

Keywords: COVID-19, work-from-home, electricity, electricity markets

*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 Matt Oliver and from participants at the 2020 virtual Southern Economic 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, 2020a,b; Brewer, 2020; Carroll et al., 2020; Narajewski and Ziel, 2020; Percy and Mountain, 2020), noting that electricity demand could serve as a good proxy for overall economic activity (e.g., Cicala, 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. 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. After controlling for weather patterns and seasonality, these data show that during the initial surge of COVID-19 cases and stay-at-home policies from March through May 2020, total electricity use was 3.7-5.4% lower than predicted consumption each month; however, after these policies expired in June 2020, total electricity use was roughly equal to the predicted baseline and in July and August was 2.9 and 4.6% *higher* than predicted consumption.

There are several possible explanations for this increased electricity demand during the summer months. First, 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. Second, it 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. Third, 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.

To determine the source of the surplus demand, I analyze monthly nationwide reported electricity consumption data by sector at the utility level. These data show that in March 2020 at the beginning of the pandemic, 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 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 offset or reduced in the following months. Furthermore, it provides evidence that as individuals and businesses have adapted to the pandemic, the relationship between electricity consumption and underlying drivers such as temperature and time-of-day have changed.

This paper contributes to the emerging literature on the economic and environmental impacts of the COVID-19 pandemic and subsequent policy and behavioral responses.[1] 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 similar applications by Benatia (2020b); Graf et al. (2021).

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 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,

[1]Given that this literature is new and constantly changing, I do not attempt to summarize emerging work beyond those cited above.

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.[2]

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.[3]

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).[4] These data are available through November 2020 and 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 monthly temperatures and average monthly precipitation across the counties it serves.[5] Due to gaps in coverage of the nClimDiv weather

[2]Data 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. This interpolation does not change results but allows for the use of of Newey-West standard errors that are robust to heteroskedasticity and autocorrelation.

[3]0.1% 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.

[4]The 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.

[5]For 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.

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 weather patterns as measured using the Mahalanobis metric.[6] 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.[7] This approach has precedent: ISOs such as PJM use similar matching algorithms to forecast day-ahead load alongside other prediction algorithms (Anastasio, 2020).

After matching, the result is a panel of counterfactual hourly zone-level electricity consumption. I aggregate this up to the PJM level by day and plot 2020's consumption and

[6] 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.

[7] I resolve ties by averaging the load within all the matched hours when applicable.

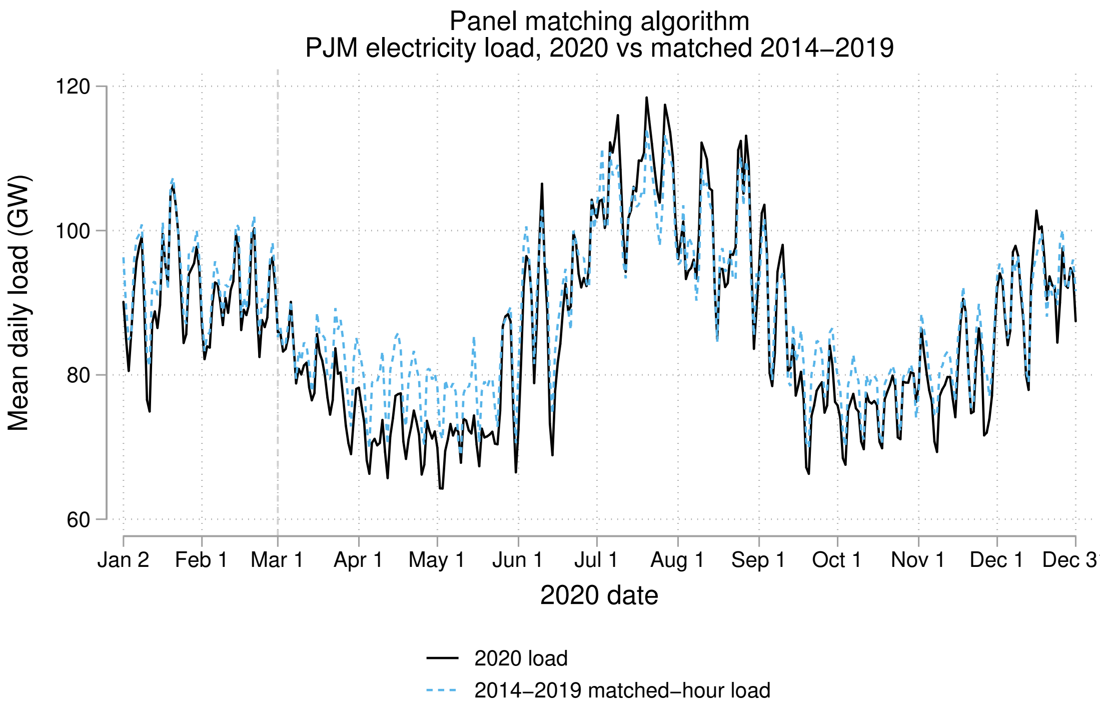

Figure 1: 2020 Electricity load matched to the most-similar weather hour from the previous five years, matching within zone exactly on the month-of-year, hour-of-day, and day-of-week. Figure displays the series in units of mean daily load for ease of interpretation.

counterfactual consumption in figure 1. 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 predicting aggregate load for all of PJM. I find that using 2014-2018 to predict hourly load in 2019 at the zone level provides a very close fit.

Looking at the matched load and realized load in figure 1, 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 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 2020 using a least-squares regression:

$$\frac{Y_{zt} - \hat{Y}_{zt}}{\hat{Y}_{zt}} = \alpha + \beta_{month} post_t + \varepsilon_{zt}, \tag{1}$$

where $post_t$ is an indicator variable equal to one for all hours after the beginning of the

COVID-19 period, and β_{month} varies by month. Estimates of β_{month} 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.[8] My preferred specification is that which excludes March 1 - 15 from the sample and thus avoids measurement error in the treatment variable.

This difference-in-predicted-difference estimator is similar to the estimator implemented in Burlig et al. (2020) and the synthetic difference-in-differences estimator described in a working paper by Arkhangelsky et al. (2019). The estimated coefficients reveal whether earlier reductions in electricity consumption were offset by later increases when individuals began returning to work.

Next, to gain an understanding of 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).[9] 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

[8]Note, it is also possible to control for zone fixed effects and temporal effects such as hour of day, but this algebraically has no effect on the coefficients (or standard errors) estimated in equation 1 because there is no variation in the treatment variables by zone or hour of day. An indicator for day-of-week does have a small, but not meaningful effect on the estimates due to a slight imbalance in days of the week.

[9]To 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.

class using data from 2014 through 2020:

$$Y_{imc} = \mu_{imc} + \omega_c weather_{imc} + \gamma_{mc} post_m + \xi_{imc}. \tag{2}$$

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 $b \in \{[10, 15), [15, 20), ..., [90, 95), \geq 95\}$, allowing the relationship to vary for 2020 versus previous years:

$$\ln(MW_{zt}) = \delta_{zymdh} + \sum_b \big(\psi_b 1(temp_{zt} \in b) + \phi_b 1(temp_{zt} \in b) \times post_t\big) + \rho\; post_t + u_{zt}. \tag{3}$$

The term δ_{zymdh} includes controls for fixed effects by zone, indicators by year, indicators for month of year, day of week times hour of day, 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.[10] 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 consumer-class level estimates, and the analysis of the differential effect of temperature on electricity use in 2020.

4 Results

Table 4 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).[11] 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 June was below normal relative to the predicted baseline. In July and August, electricity consumption was above the predicted baseline. In the fall, electricity consumption returned to pre-COVID levels before again increasing in December.

Figure 2 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 3.7%, 5.4%, and 4.0% lower relative to the predicted baseline levels in each month, respectively. In June, electricity consumption had begun to recover and was only 0.2% lower relative to the predicted baseline. In July and August, electricity consumption spiked, coming

[10]Because 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.

[11]These 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.

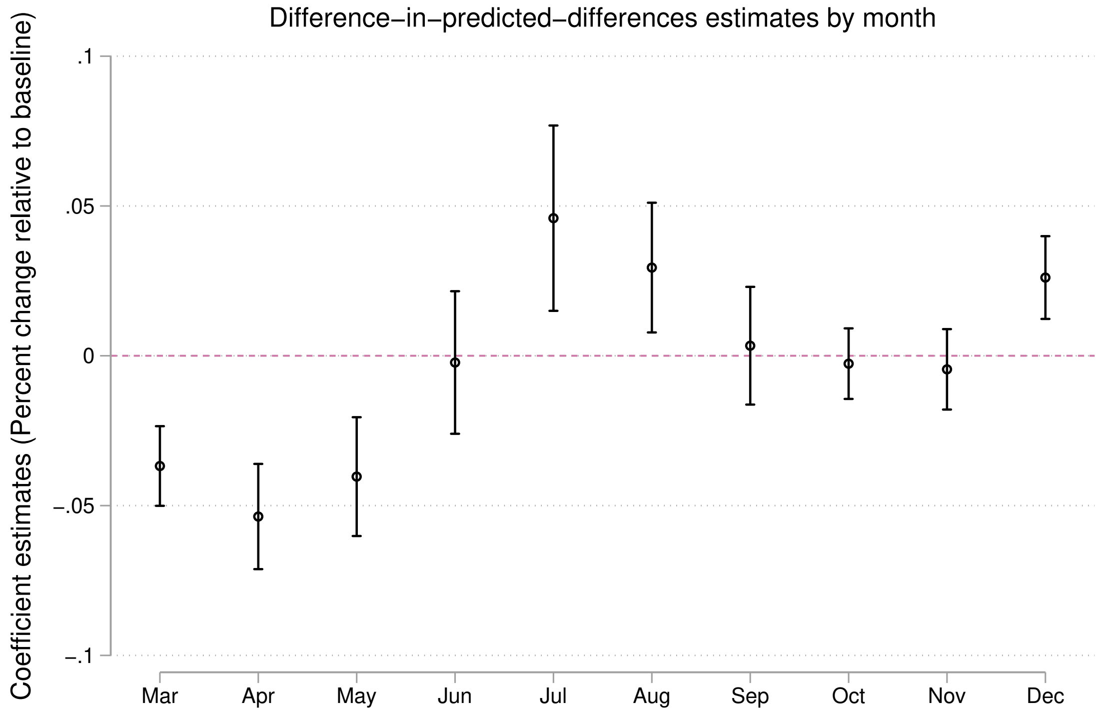

Figure 2: Estimates from column (5) of table 4 of the difference-in-predicted-differences model in equation 1. 95% confidence intervals account for multi-way clustering by zone and by day. Estimates are the monthly difference in predicted differences before and after March 1st, 2020 after omitting observations from March 1-15, and when multiplied by 100 are interpreted as the percentage difference in electricity consumption relative to the baseline prediction.

in at 4.6% and 2.9% higher relative to the predicted baseline. In September, October, and November, electricity consumption was within 0.5% of normal at 0.34% higher, 0.26% lower, and 0.45% higher respectively, relative to the baseline. In December, electricity consumption was 2.6% higher. After weighting by total load in each month, the difference-in-predicted-difference estimates suggest an overall 0.09% increase in electricity consumption after March 15th relative to the predicted baseline.

Why was electricity consumption higher in July, August, and December? This pattern could reflect electricity consumption that was delayed rather than displaced by stay-at-home behavior, but it is also possible that this pattern is due to other behavioral changes related to spending more time at home. For example, it could be that the warm summer

Table 1: Difference-in-predicted-differences estimates

	(1)	(2)	(3)	(4)	(5)
Post x March	-0.015*	-0.027*	-0.038*	-0.027*	-0.037*
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Post x April	-0.054*	-0.056*	-0.055*	-0.054*	-0.054*
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Post x May	-0.040*	-0.042*	-0.042*	-0.040*	-0.040*
	(0.010)	(0.009)	(0.009)	(0.010)	(0.010)
Post x June	-0.002	-0.004	-0.004	-0.002	-0.002
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
Post x July	0.046*	0.044*	0.044*	0.046*	0.046*
	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Post x August	0.029*	0.028*	0.028*	0.029*	0.029*
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Post x September	0.003	0.001	0.002	0.003	0.003
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Post x October	-0.003	-0.005	-0.004	-0.003	-0.003
	(0.006)	(0.006)	(0.005)	(0.006)	(0.006)
Post x November	-0.005	-0.006	-0.006	-0.005	-0.005
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Post x December	0.026*	0.024*	0.024*	0.026*	0.026*
	(0.007)	(0.007)	(0.007)	(0.007)	(0.007)
Overall effect	0.13%	-0.10%	-0.08%	0.08%	0.09%
Treatment date	March 1	March 8	March 15	March 9	March 16
Excluded dates				March 1-8	March 1-15
Observations	184464	184464	184464	180432	176904

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.

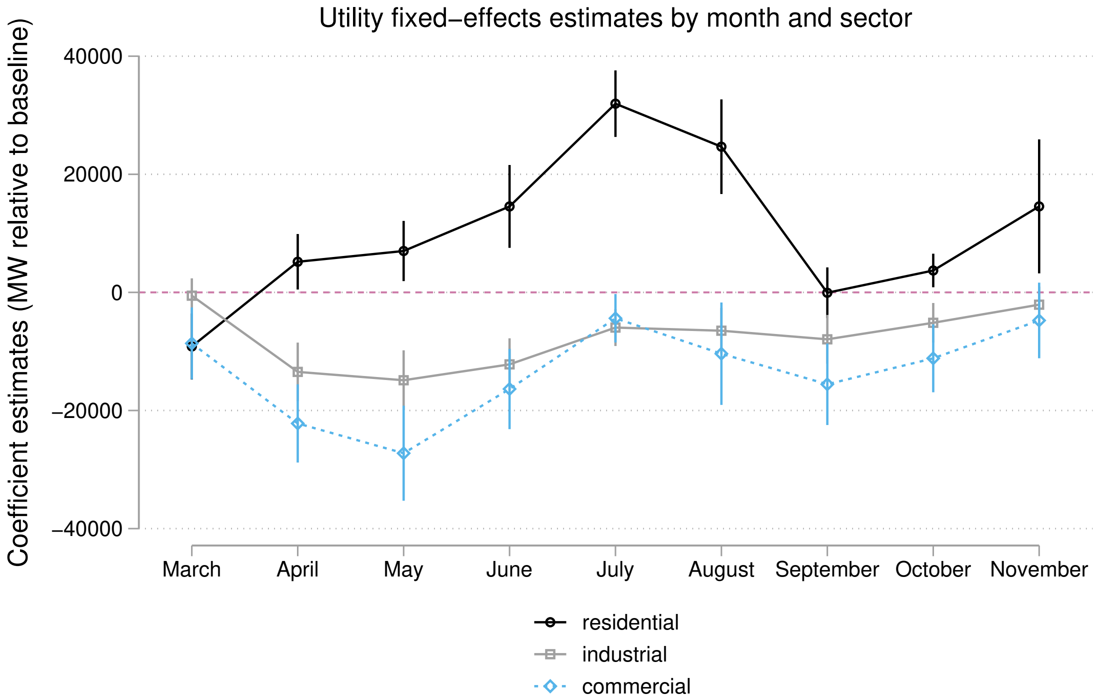

Figure 3: Estimates by customer class at the state-utility level from equation 2. 95% confidence intervals account for multi-way clustering by utility-state and by month.

temperatures caused individuals to run their home air conditioning units more and that this accounts for the summer bump rather than increased production. Similarly, more electricity may have been used to run home furnaces in December. To investigate this possibility, I compare the amount of electricity consumed by residential versus industrial, commercial, and transportation end users in 2020 versus the previous five years.

Table 2 displays the estimates from equation 2 using the nationwide panel of reported monthly consumption by class, and figure 3 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

Table 2: Monthly utility panel estimates

	(1) Residential	(2) Commercial	(3) Industrial
March	-9153.5*	-8631.6*	-554.0
	(2815.8)	(3060.4)	(1474.1)
April	5186.4*	-22189.5*	-13457.9*
	(2362.4)	(3330.1)	(2491.4)
May	7008.6*	-27236.1*	-14874.0*
	(2562.0)	(4045.2)	(2547.1)
June	14556.3*	-16357.3*	-12188.3*
	(3524.3)	(3416.2)	(2217.3)
July	31956.6*	-4382.0*	-5954.8*
	(2830.9)	(2063.0)	(1572.4)
August	24669.1*	-10381.3*	-6488.9*
	(4028.5)	(4358.1)	(1939.5)
September	-59.46	-15559.4*	-7944.7*
	(2156.4)	(3468.9)	(2079.9)
October	3708.3*	-11163.7*	-5144.8*
	(1428.2)	(2890.8)	(1683.9)
November	14565.8*	-4749.9	-2078.4
	(5703.6)	(3217.4)	(1535.1)
State-utility FE	Yes	Yes	Yes
Month-of-year indicators	Yes	Yes	Yes
Weather controls	Yes	Yes	Yes
Observations	34935	34932	34929

Each observation is monthly electricity consumption by customer class at the state-utility level. 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.

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 years.

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 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 3.

Figure 4 plots the coefficient estimates on temperature interval interacted with the indicator variable for March-December 2020 from regression 3. 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. 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 5 (thus, the 55-60°F estimate is zero and the other estimates are relative

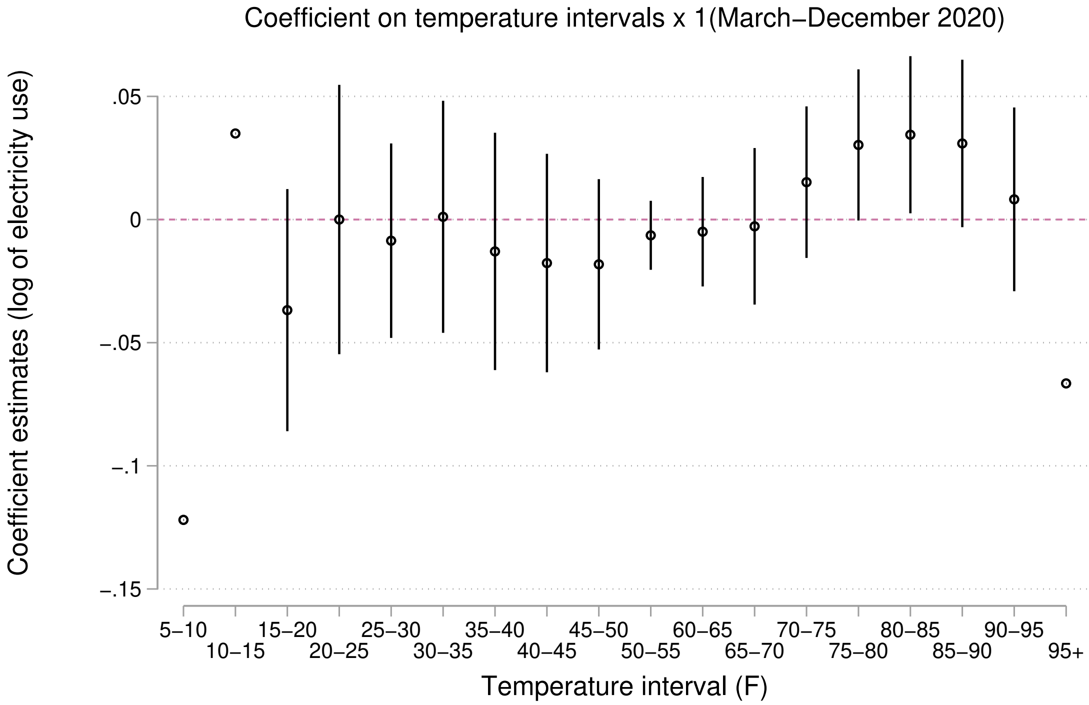

Figure 4: Coefficient estimates on temperature interval interacted with an indicator variable for March-December 2020 from regression 3. Temperature interval 55-60°F is the omitted category. Estimates represent the difference in the relationship between temperature and electricity load in 2020 versus 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-value<0.000.

to it). Higher temperatures are more costly while low temperatures are somewhat less costly in 2020 than in previous years. 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.

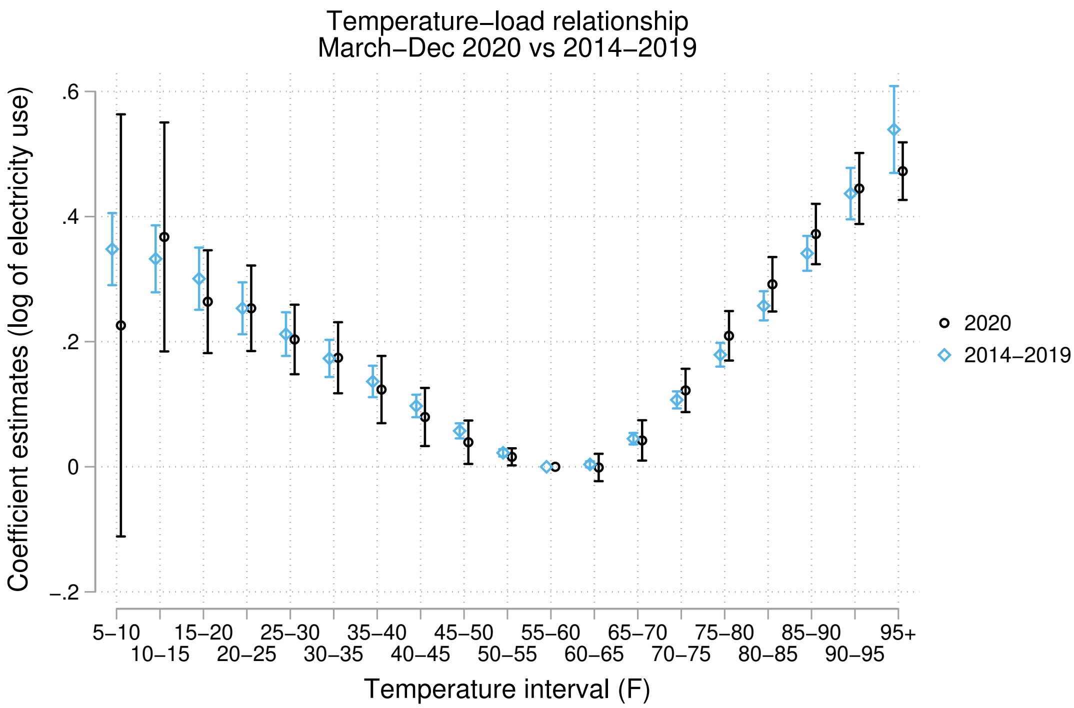

Figure 5: Marginal effect on log load of being in a temperature interval relative to the excluded 55-60°F interval. 95% confidence intervals account for multi-way clustering by zone and by day.

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 3.7, 5.4, and 4.0% in the first three months of the COVID-19 pandemic, electricity use was 2.9 and 4.6% higher in July and August 2020. Electricity consumption in September through November was roughly normal compared to the predicted baseline, while consumption in December was 2.6% higher than the predicted baseline. 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, peaking in July. As a whole, the early reductions in electricity consumption were almost perfectly offset by the increases in July, August, and December, with estimates of the overall effect centered on zero across specifications, ranging from -0.10% to +0.13%.

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. The spike in December can only be attributed to unobserved differences in electricity consumption.

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.

References

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.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2019). Synthetic difference in differences. NBER Working Paper Series 25532, National Bureau of Economic Research.

Benatia, D. (2020a). Electricity markets under lockdown: Insights from new york. *IAEE Energy Forum Special Covid-19 Edition 2020*, 55–58.

Benatia, D. (2020b). Reaching new lows? the pandemic's consequences for electricity markets. USAEE Working Paper 20-454, United States Association for Energy Economists.

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.

Bui, Q. and J. Wolfers (2020, 4). Another way to see the recession: Power usage is way down. *The New York Times The Upshot*.

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. (2020). Powering work from home. NBER Working Paper Series 27937, National Bureau of Economic Research.

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.

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.

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.

Imbens, G. and A. Abadie (2008). On the failure of the bootstrap for matching estimators. *Econometrica 76*(6), 1537–1557.

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.

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.

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.

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

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). The result is a single time-series of hourly counterfactual PJM electricity consumption. I refer to this algorithm as the time-series matching algorithm.

The second matching algorithm performs the 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. I refer to this algorithm as the panel matching algorithm.

Figure 6 displays metered 2019 load and the matched load using the nearest-neighbors algorithm at the PJM level. Figure 1 displays metered 2019 load and the matched load using the nearest-neighbors algorithm at the zone level. Comparing these figures, 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.

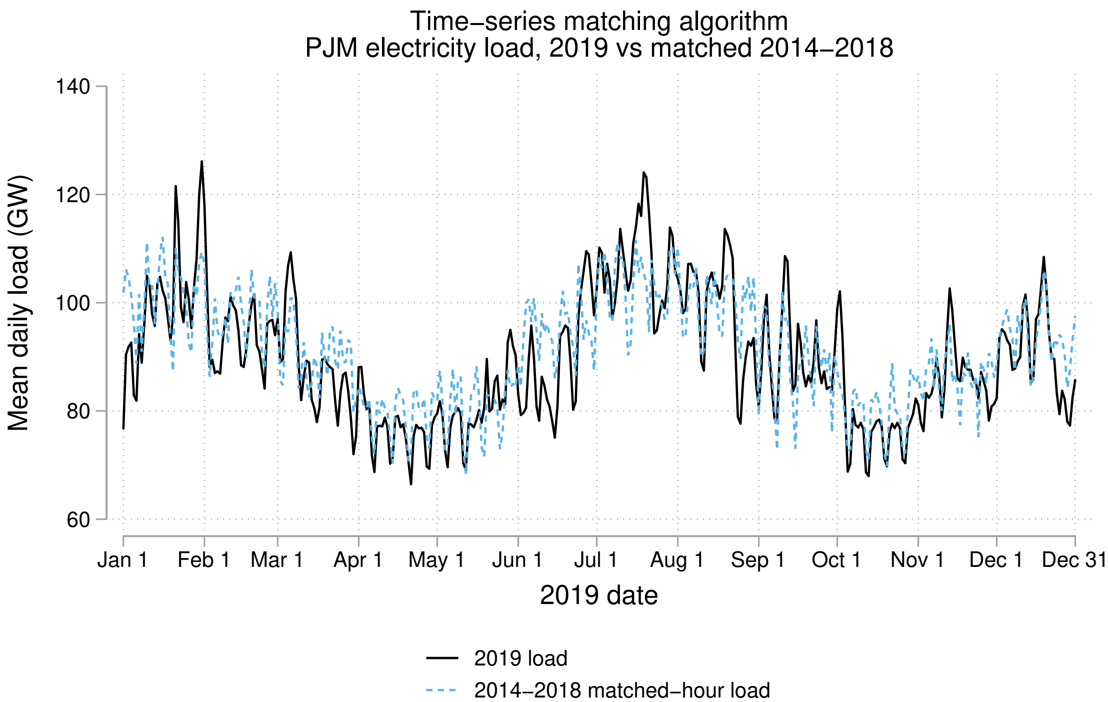

Figure 6: 2019 Electricity load matched to the most-similar weather hour from 2014-2018, matching exactly on the month-of-year, hour-of-day, and day-of-week at the aggregate PJM level. Figure displays the series in units of mean daily load for ease of interpretation.

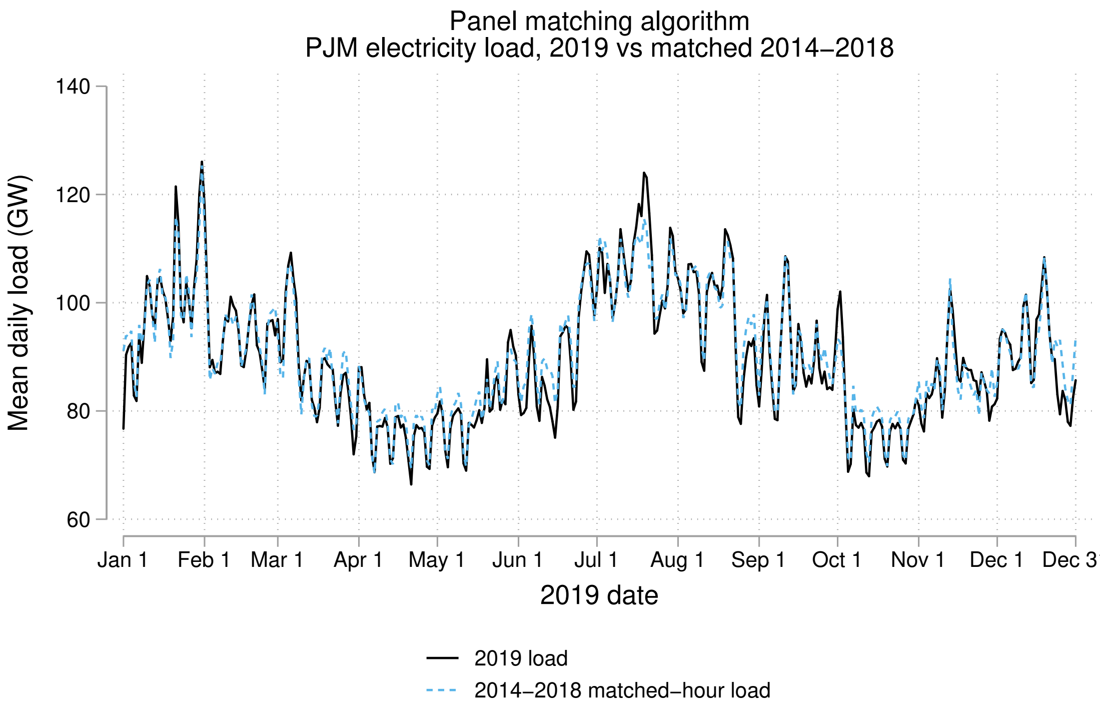

Figure 7: 2019 Electricity load matched to the most-similar weather hour from 2014-2018, matching within zone exactly on the month-of-year, hour-of-day, and day-of-week. Figure displays the series in units of mean daily load for ease of interpretation.