The Economic Costs of Forecasting Errors in the PJM Interconnection Due to the COVID-19 Quarantine

BY DYLAN BREWER

Dylan Brewer is with the School of Economics, Georgia Institute of Technology. e can be reached at brewer@gatech.edu

See footnotes at end of text.

Introduction

Doors are closed, lights are off, and normally busy rooms are silent in businesses around the world as individuals stay home and workplaces shut down to prevent the spread of COVID-19. Economic activity has declined significantly—and as a result, commercial and industrial demand for electricity have also declined. More people at home means increased residential demand for electricity, but preliminary data show that total electricity load has fallen nearly 10% in the U.S. and Europe, and the wholesale price of electricity in India is approaching an all-time low[1]. How has the decline in electricity demand affected electricity markets?

This unforeseen decline in electricity demand has disrupted electricity demand forecasts, increasing the costs of supplying electric power. On the grid, electricity supply (generation) and demand (load) must be balanced in real-time within a small band in order to avoid power disruptions. Because of this real-time balance, grid operators and utilities rely on forecasts of electricity demand in order to procure commitments in the day-ahead market from suppliers to generate the right amount of electricity at the right time. If the forecast predicts electricity demand will be low when it is high, the grid operator must procure "adjustment" generation at a higher cost to cover the extra demand. If the forecast predicts electricity demand will be high when it is low, the grid operator will have purchased too much electricity generation in the day-ahead market, increasing the price.

The declines in demand due to COVID-19 represent a massive shift in electricity load patterns, disrupting the accuracy of the day-ahead forecast. This article uses hourly data from the PJM Interconnection (PJM) to analyze the performance of the day-ahead load forecast during the COVID-19 period. First, I show that after matching on weather, day-of-week, and time-of-day, PJM electricity consumption during the COVID-19 period relative to the beginning of the year has declined 10.6%. Next, I show that PJM's day-ahead forecast error jumps to 3% on average during the COVID-19 period, on par with the worst performance of electricity forecasting models in the last five years. In addition, it appears that the forecast did not improve in the first week of April relative to March. Because of relatively mild temperatures in March and April 2020, the cost of forecast error is likely to be small; however, failure to adjust forecasting models could prove costly as the quarantine continues and temperatures warm.

The day ahead market and load forecasting in PJM

PJM is a U.S. regional transmission organization that operates a wholesale electricity market and transmission system serving parts of the Mid-Atlantic and Midwest regions. To purchase commitments of electricity generating capacity, PJM operates a wholesale day-ahead market in which generators bid to provide electricity and utilities bid to withdraw electricity in each hour of the following day. Each supplier bid stipulates the amount of generation (in MWh) offered during an hour and a minimum price required for the generator to operate. Each demand bid stipulates the amount of load (in MWh) and the maximum price required—in practice many bids do not specify a maximum price.

Bidding in the day-ahead market occurs the seven days prior to the day of generation. The day-ahead market closes at 10:30 AM ET the day prior to generation (PJM, 2017). Upon close, the market operator orders the bids from lowest price to highest price to create a supply or offer curve and accepts the bids required to meet demand. This ensures that the lowest-bid generating units are used to provide electricity.

Because it is impossible to know demand the next day, utilities purchasing electricity must forecast their load based upon their expectation of the next day's electricity demand. weather and economic activity. Typical load forecasts are based upon lagged load, forecasted weather (especially temperature), and time-of-day indicator variables meant to capture human and economic activity (see e.g., McCulloch and Ignatieva, 2020).

An extreme and novel event such as COVID-19 can break the link between past load and future load, causing forecasts to perform poorly. Fundamentally, forecasts predict future electricity load based upon the patterns observed in past load and thus rely on the assumption that future load patterns are comparable to what has been seen before. When an extreme event occurs, the relationship between the outcome variable and predictor variables can change. A statistician would say that the underlying data generating process has changed, degrading the performance of the forecast; others would liken the scenario to using a roadmap that is 10 years old or even flying in the dark.

Change in load during the COVID-19 period

How dramatically have electricity loads changed in PJM due to COVID-19? Between March 15th and April 4th, more than 10 percent of U.S. workers applied for unemployment benefits (Gould and Shierholz, 2020). Every state in the PJM area has a shelter-in-place order in effect (Key, 2020), and most schools and workplaces have switched to remote operation when applicable.

Many more households are at home during the day—either unemployed or working from home—and many industrial and commercial sites are unoccupied. One would thus expect overall load to decline due to decreased economic activity and for the load profile to change as schedules become more flexible.

To test these hypotheses, I examine hourly-metered load data from PJM and hourly weather data (temperature, dew point, precipitation, and snow depth) from NOAA's Integrated Surface Database from January 1, 2014 to April 7, 2020. I map weather stations to PJM's electricity transmission zones using weights published by PJM (PJM, 2018). It is important to control for weather to be sure that any pattern in electricity

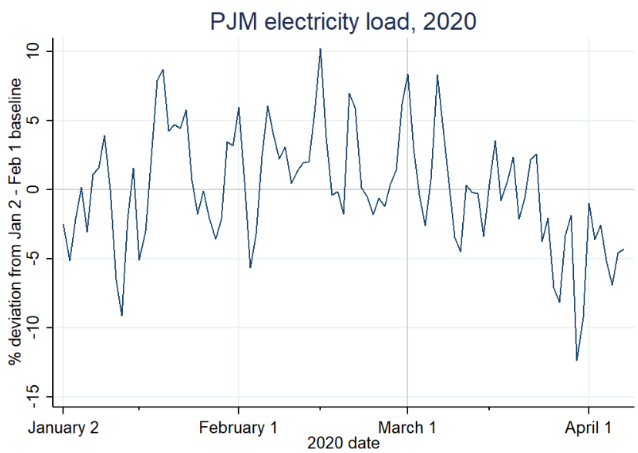

Figure 1: PJM hourly electricity load, controlling for weather, month-of-year, day-of-week, and hour-of-day variation. Presented at the average daily level and normalized to deviations from zero for the baseline period January 2nd - February 29th.

consumption observed in March and April 2020 is not being caused by changing weather.

Figure 1 plots total load in PJM from January 2nd to April 7th, 2020, controlling for the weather in each of PJM's zones, month-of-year, day-of-week, and hour-of-day variation in

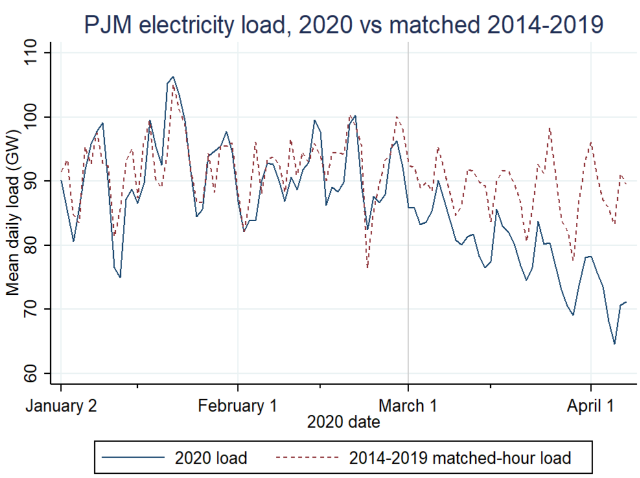

Figure 2: Raw electricity load matched to the most-similar weather hour from the previous five years, matching exactly on the hour-of-day and day-of-week.

load patterns.[2] A slight trend break can be seen around March 1st, but there is still a lot of residual noise in the series and it is difficult to make concrete statements about the size of the change in electricity consumption.

Another way to see the change in demand patterns is to compare March and April 2020 electricity consumption to similar-weather hours from prior years. Figure 2 plots the raw load data from 2020 versus similar-weather hours from 2014-2019. While electricity demand was fairly normal for the first two months of the year, a clear and growing decrease in load can be seen beginning around March 1st. To construct this figure, I match each hour of electricity demand in 2020 to the most-similar weather hour from the previous five years, matching exactly on the hour-of-day and day-of-week[3]. For example, I compare the hour on Wednesday April 1st from 12-1pm to all Wednesday 12-1pm hours in the first four months of 2014 to 2019 and match it to the hour with the most

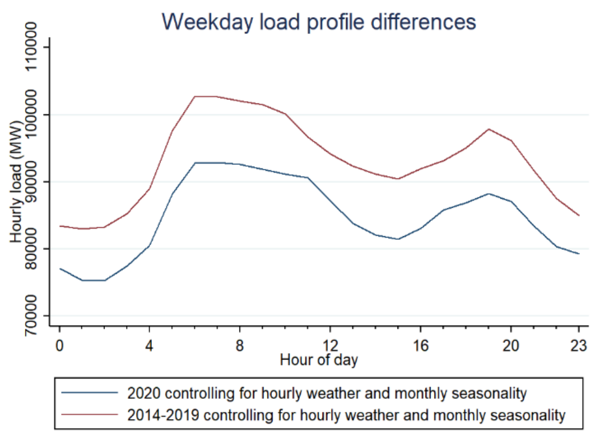

Figure 3: Average March and April weekday load by hour-of-day, controlling for weather and monthly seasonality.

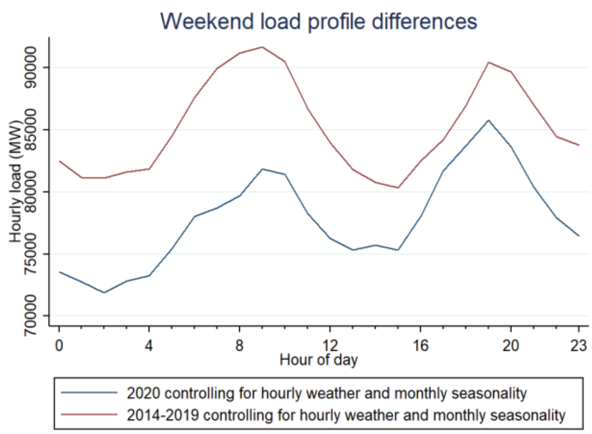

Figure 4: Average March and April weekend load by hour-of-day, controlling for weather and monthly seasonality.

similar weather. Overall, the difference in average PJM electricity consumption during the COVID-19 period from March 1st to April 7th relative to the matched baseline is 10.6% lower than this difference between January 2nd and February 29th.[4]

Finally, one would expect intra-day load patterns to change given the spike in unemployment and work-from-home responses. Figures 3 and 4 plot the average electricity use for each hour of March and early April 2020 and each hour of March and April on weekdays and weekends for the prior five years, controlling for hourly weather variation and monthly seasonality. Surprisingly, while there is a mean-shift in hourly load, it is difficult to detect any change in patterns other than a slight elongation and flattening of the mid-morning peak.

Forecasting error

While each bidder's forecasting method is private and their day-ahead bids are kept confidential for six months, PJM constructs and releases an independent public forecast that is likely similar to forecasts used by bidders (or may be an input to private forecasts). The forecast uses a mix of prediction algorithms (including three different neural nets, a matching algorithm, and time-series regression) created both in-house and by a third-party vendor (Anastasio, 2017). The forecasting team often uses a weighted average of each algorithm's prediction but has found that different algorithms work better for extreme-weather days and holidays. PJM's day-ahead forecast updates every half hour on the quarter; thus, the final two forecasts before the day-ahead market closes update at 9:45 AM ET and 10:15 ET. PJM's database contains historical 9:45 AM forecasts. I examine the performance of PJM's 9:45 AM forecast in March 2020 versus its historical performance to understand how well prediction models are performing given the COVID-19 shock to electricity load.

I consider two measures of forecasting error. The first measure is the mean percent prediction error, which is simply the absolute hourly forecasting error as a percent of load, averaged by month. This measure captures the relative performance of PJM's prediction algorithm in the COVID-19 period relative to its past performance. The top panel of Figure 5 displays monthly average percent prediction errors from PJM's forecast from January 2014 through April 2020. March

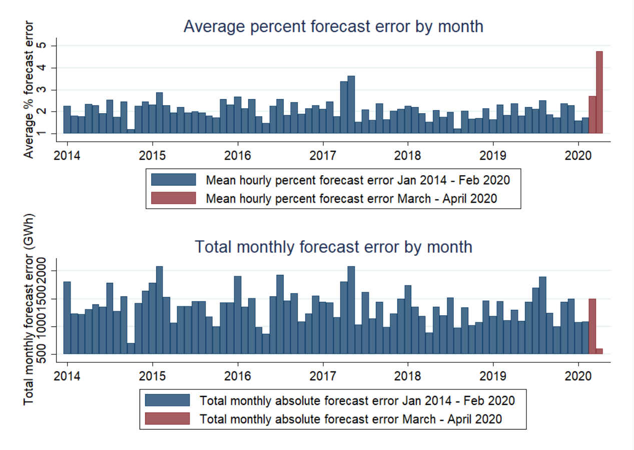

Figure 5: Mean hourly percent forecast error by month (top) and total absolute forecast error by month (bottom) from January 1, 2014 though April 7, 2020.

2020 was the fourth-worst forecasting month in six years while the beginning of April 2020 is on track to be the worst forecasting month in six years by mean percent prediction error. While not displayed in this article, controlling for weather increases the March and April relative forecasting error size.

The second measure of forecasting error is the sum of absolute prediction error, which is simply the total absolute hourly forecasting error. This measure roughly corresponds to the relative economic costs of the forecasting error in PJM during the COVID-19 period because it reflects the size of the missed predictions. The bottom panel of Figure 5 displays total monthly prediction errors from PJM's forecast from January 2014 through April 2020. Despite the large percentage error, the total absolute forecast error is relatively small because March is a relatively low-demand month for electricity and April has only begun. Even though the forecast is performing poorly, the cost of the error is relatively small. Of the 76 months since January 2014, March 2020 had the 23rd-highest total absolute forecasting error.

Conclusions

Self-quarantine efforts in response to COVID-19 have induced a 10 percent decline in PJM electricity load from March 1st through April 7th relative to baseline levels when accounting for weather, hourly, and monthly variation. This sudden change in the underlying data-generating process for electricity load has reduced the predictive power of PJM's day-ahead forecast. The forecasting model performed poorly relative to previous performance—March 2020 was the forecast's fourth-worst month in terms of percent absolute forecast error. Despite this, it is likely that over-procurement of generation in the day-ahead market was small due to the relatively low electricity load in March 2020. Failure to adjust forecasting models (or for the learning-based algorithms to self-adjust) may become more costly as summer loads increase.

As of the first week, April's forecasting error is on track to be the worst in recent history, thus it is not clear whether the model is improving. An adjustment to a complex forecasting model is not simple. It requires human time to experiment with changes to underlying parameters of the model and significant computation time. In addition, electricity load patterns continue to change as conditions change, so it is possible that new forecasts would soon be obsolete. Forecasts with increased weight on recent dates (during the COVID-19 period) may perform better, but this may not prove true as the seasons change from spring to summer and conditions continue to change

Footnotes

[1] US: https://www.nytimes.com/interactive/2020/04/08/upshot/electricity-usage-predict-coronavirus-recession.html. Europe: https://www.wsj.com/articles/plunge-in-italys-electricity-use-hints-at-coronavirus-risks-facing-u-s-11584532801. India: https://economictimes.indiatimes.com/industry/energy/power/power-demand-drops-as-offices-stay-plugged-out/articleshow/74819188.cms?from=mdr.

[2] This method is similar to that taken by Steve Cicala in his analysis

described by Bui and Wolfers (2020).

[3] Matching up to the ten most-comparable hours does not substantially change the results.

[4] This estimate is the difference-in-predicted-differences estimate, similar to the estimators considered in Burlig et. al (2020).

References

Anastasio, E., (2017), Load Forecasting at PJM. Presentation, January 10, 2017. https://www.pjm.com/~/media/committees-groups/committees/oc/20170110/20170110-item-06-load-forecasting-at-pjm.ashx.

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.

Burlig, F. et al, (2020), "Machine Learning from Schools about Energy Efficiency." R&R at the Journal of the Association of Environmental and Resource Economists.

Key, P., (2020), RTOs, ISOs Seeing Load Changes, Adjusting Models Due to Pandemic. EnergyCentral. https://energycentral.com/c/em/rtos-seeing-load-changes-adjusting-models-due-pandemic?utm_medium=eNL&utm_campaign=em_net&utm_content=0&utm_source=2020_04_08.

Gould, E. and H. Shierholz, (2020), 16.8 million apply for unemployment benefits in three weeks. Economic Policy Institute. https://www.epi.org/press/16-8-million-apply-for-unemployment-benefits-in-three-weeks-the-labor-market-has-been-upended/.

McCulloch, J. and K. Ignatieva, (2020), Intra-day Electricity Demand and Temperature. The Energy Journal 41 (3): forthcoming. http://www.iaee.org/en/publications/ejarticle.aspx?id=3506.

PJM, (2017), Day-Ahead Energy Market. Presentation, June 12, 2017. https://www.pjm.com/-/media/training/nerc-certifications/markets-exam-materials/generation-itp/day-ahead-energy-market.ashx?la=en.

PJM, (2018), PJM Manual 19: Load Forecasting and Analysis, revision 33. https://www.pjm.com/~/media/documents/manuals/archive/m19/m19v33-load-forecasting-and-analysis-10-25-2018.ashx.