Early Economic Impacts of COVID-19 in Europe: A View from the Grid

Steve Cicala*
University of Chicago
Click here for the latest version.

8 April, 2020

Abstract

This paper presents preliminary estimates of how electricity consumption has changed in the European Union since the spread of COVID-19, as a proxy for short-term changes in economic activity. I collect hourly data by country from European Network of Transmission System Operators for Electricity (ENTSO-E) from 2016-present, and match it with automated weather stations to adjust for heating and cooling demand. As of the week ending 4 April, 2020, power consumption is down roughly 10%, with large differences across countries reflecting the timing and stringency of lockdown policies.

*I am grateful to Francesco Decarolis, Tommaso Monacelli, and Justin Wolfers for helpful conversations. Iván Higuera-Mendieta provided excellent research assistance under unusual circumstances. Ari Anisfeld, Chinmay Lohani, and H.I. Park also pitched in on short notice. This research is funded by the Political Economics Initiative at the Becker Friedman Institute at the University of Chicago. All errors remain my own. e-mail: scicala@gmail.com

1 Introduction

Governments around the world have taken measures to reduce in-person interactions to reduce the speed of transmission of coronavirus disease 2019 (COVID-19). Mass quarantines and shutdowns are likely to have significant economic costs and require a substantial policy response to provide assistance to individuals who are unable to work while at home. Such a response naturally requires an understanding of the magnitude of the economic downturn.

While policies to reduce transmission have been instituted suddenly, standard economic indicators to measure their impact are collected over longer time horizons. This creates a costly window of uncertainty in which individuals are exposed to the full economic costs of shutdowns, but policy-makers are unable to measure these costs and respond accordingly.

Researchers have turned to a variety of high-frequency measures to fill this informational void by proxy. These include restaurant reservation databases, web searches, cellular phone locations, traffic, and many others. The challenge in interpreting such measures is in figuring out how they translate into the target statistics of interest–GDP and employment, for example. However the problem of low-quality economic statistics is not new—the salient issue has typically been one of data reliability rather than temporal delay.

In settings with low data quality, Henderson et al. (2012) show that night lights work well as a proxy for economic activity. It is intuitive that energy consumption follows economic activity closely because it is both widely-used throughout the economy and difficult to substitute away from (at least in the short-run). Night lights, electricity consumption, and other indicators of real economic activity have also been employed to detect manipulation of national accounts (Lyu et al. (2018); Chen et al. (2019)) following a similar logic. In concurrent work Cicala (2020) shows that grid-scale electricity consumption closely tracked economic activity during the 2008 financial crisis and subsequent recession in the United States: from peak to trough GDP fell 4.3%, while weather-adjusted electricity consumption fell about 5%.

In this paper I present early estimates of changes in electricity consumption during the COVID-19 pandemic in the European Union. These data are reported hourly, and I adjust each country's consumption for the variables that typically explain upwards of 90% of variation: hour of day, day of week, week of year, and temperature. In the work week ending 3 April, 2020, consumption was down roughly 10%, with significant differences across countries, reflecting the disparate impacts and policy responses to COVID-19.

It is important to distinguish the potential value of electricity consumption as a proxy for economic growth in the short-run versus the longer-run. Over time consumption responds to the price of consuming electricity services and technological innovations affect demand (among other drivers). Changes in productivity and the ability to substitute toward other inputs over time is likely to make electricity a less useful proxy over longer time horizons. Such dynamics are unlikely to be important drivers of consumption changes over the time horizon in which economic information is unavailable.

The paper is structured as follows: In the next section I describe the data, the third section describes the methods I use to adjust for confounding factors, the fourth section presents the results, and the fifth section concludes.

2 Data

Electricity consumption is observed at the country-hour from 2016-present (with Germany reporting four zones separately). These data are collected from the European Network of Transmission System Operators for Electricity (ENTSO-E) using the pyISO module from WattTime. Electricity consumption is reported as "system load." This is a measure of the amount of power drawn from the bulk transmission system, and includes the total of residential, commercial, and industrial consumption (not separated by customer class). Industrial plants that self-generate are only observed insofar as they draw additional power from the grid beyond their own production.

Temperature data is collected hourly from the U.S. National Weather Service's Automated Surface Observing Systems (ASOS), a network of automated weather stations that are typically located at airports. These stations are geolocated to the EU's Nomenclature of territorial units for statistics (NUTS)-level 3 codes, and the population of the stations' local territory is used to weight temperature readings up to the corresponding country or transmission zone. Data on major holidays over the sample period were collected from online resources.

Figure 1 shows the extent to which electricity consumption is both highly variable and predictable. Panel (a) plots total system load by day in 2019. In contrast to the United States, where air conditioning is widely adopted, the EU has a "winter peaking" grid, in which electric power is used to heat many homes during the colder months. The more high-frequency cyclical pattern follows days of the week: consumption falls nearly 20% on the weekends. Holidays also stand out, reflecting time away from work at the start and end of the year, Christian holidays in late April and May, as well as summer vacations in August. Panel (b) shows how electricity consumption follows the course of the day (in Universal Coordinated Time). There is a baseload level of consumption when most people are asleep, and load subsequently rises following the work day (with a lull in the afternoon).

Figure 1: E.U. Electricity Consumption in 2019

(A) Total Daily Load

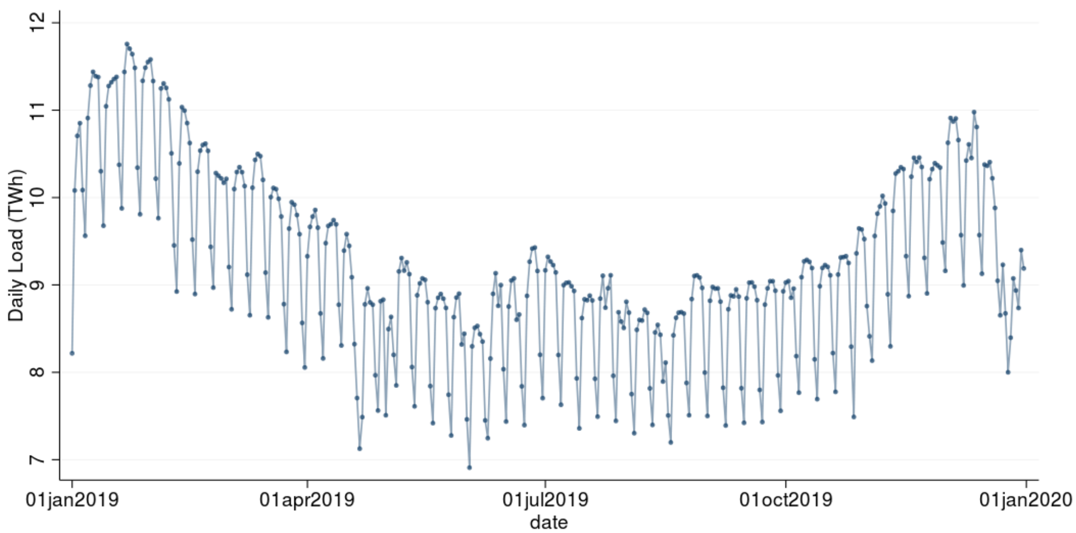

(B) Mean Load by Hour of Day

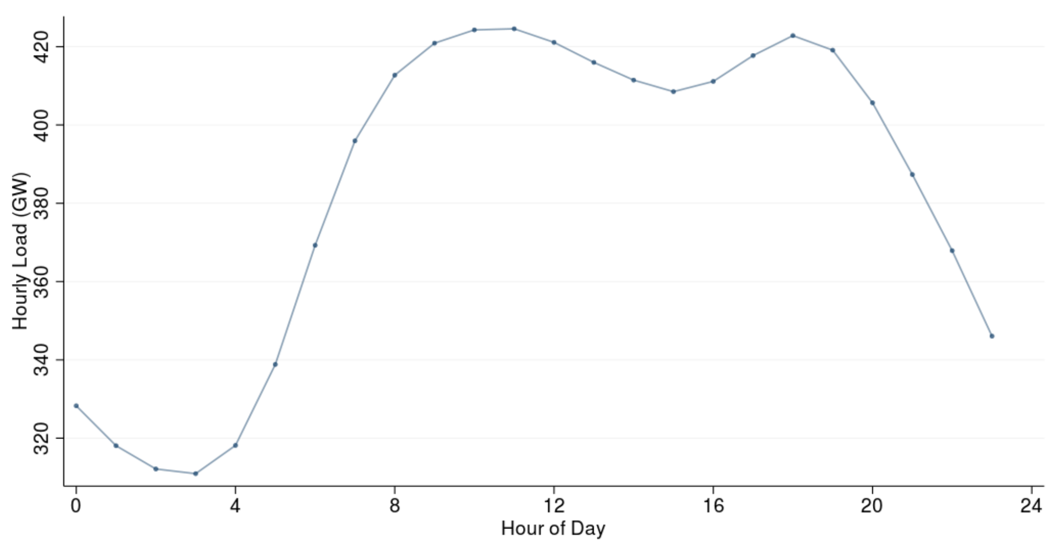

Source: ENTSO-E.

3 Methods

I conduct my analysis at the country/zone (i)-by-hour (t) level. I run regressions of the form

$$Log(Load_{it}) = \tau_{id} + \Omega_i + \Psi_i + \eta_i + \Gamma_i + \sigma_i heating_{it} + \kappa_i cooling_{it} + u_{it}$$

The covariates are a set of dummies for day of week (Ω_i), hour of day (Ψ_i), holidays (η_i), week of year (Γ_i), as well as season/hour-specific coefficients for heating and cooling degrees.[1] The target of interest is a set of dummy variables, τ_{id}, that indicate specific dates of interest–typically daily dummies for the final year (or months) of the sample. Having adjusted for temperature fluctuations, these dummies measure by how much load on date d differed from prior years within the same week of the year, day of week, hour of day and temperature. Note that this set of dummies cannot span the entire sample period without making the week of year dummies collinear. I therefore employ a rolling window of analysis: I estimate daily means τ_{id} for year y by pooling data for years $y-3$ through y, making the daily fixed effects an interaction between day of year and an indicator for the final year of the sample.[2] When presenting weekly results instead of daily, the τ_{iw} are indicators for week of sample rather than date, and carry a similar interpretation–the mean reduction in that load relative to the same week in prior years, all else equal. To ease interpretation I normalize coefficients relative to a baseline period before salient events. Without any normalization these coefficients measure the change relative to prior years–the normalization saves from needing to do an additional subtraction before calculating the statistic of interest. One can renormalize the results to taste.

These regressions are run separately by country/zone, so the i subscripts indicate that coefficients are specific to each area. The error term u_{it} is likely to be serially correlated, so standard errors are clustered at the country/zone-month. In these preliminary results I have estimated each country separately, then calculate the EU-wide average using mean load in 2019 as weights.

4 Results

Figure 2 presents the main results by calendar date for the European Union. The coefficients are normalized so that the mean of the estimates between 1 February and 21 February, 2020 is zero. The first regional quarantines in Northern Italy began with the closure of that window. It shows a sharp break in the second week of March as large-scale shutdowns and quarantines took effect, with most recent days indicating a roughly 10% decline in electricity consumption relative to baseline.

The data reveal an important pattern of declines over the course of the week: Because electricity consumption (and economic activity) are already low on weekends, there is less room to fall. This creates an illusion of 'smaller drops' on the weekend. This is effectively because power consumption is falling quite a lot during the work week, but there are fewer shops and factories to close on weekends beyond those that normally close on weekends already. The declines are therefore smaller on weekends.

Table 1 present the results by country, averaging over the entire week ending 4 April, 2020.[3] The decline in electricity consumption is widespread, but also heterogeneous. There are enormous

[1] A heating-degree in hour t is defined as the number of degrees the ambient temperature is below 18^oC: $\max\{18-temperature_{it}, 0\}$. It is defined analogously for cooling degrees when the ambient temperature exceeds 18^oC.

[2] I have also conducted the analysis using data from years $y-3$ to $y-1$, and calculating out of sample residuals for year y (of which I then average by date). The results are broadly similar: the distinction between them is whether data from year y contributes to the estimation of the heating and cooling coefficients, σ_i and κ_i. The main advantage of using dummies instead of predicting out of sample is the standard inference on the daily means.

[3] A few countries that report to ENSO-E are missing from the table due to data quality issues.

declines in the countries known to have been particularly hard-hit, with Italy down nearly 25% and Spain down 15%. A broad swath of the EU has fallen closer to the overall average, including France, Belgium, the U.K., and parts of Germany. Nordic countries have not registered declines, with Norway, Sweden, and Denmark all operating close to historical levels.

Figure 3 presents the daily picture for Italy, which has been especially hard-hit by the virus, and has implemented extraordinary measures to reduce the speed of transmission. As with the EU-wide figure, there is less room to fall on weekends–but weekday reductions have been down by over a quarter since the third week of March. The differences in shutdown timing is visible in the daily plots of individual countries, as shown for the UK in Figure 4. The UK announced lockdowns on 23 March, 2020. Their consumption profile was reasonably flat prior to that date, and has fallen precipitously since.

Figure 2: Changes in EU Electricity Consumption: 1 February - 6 April, 2020

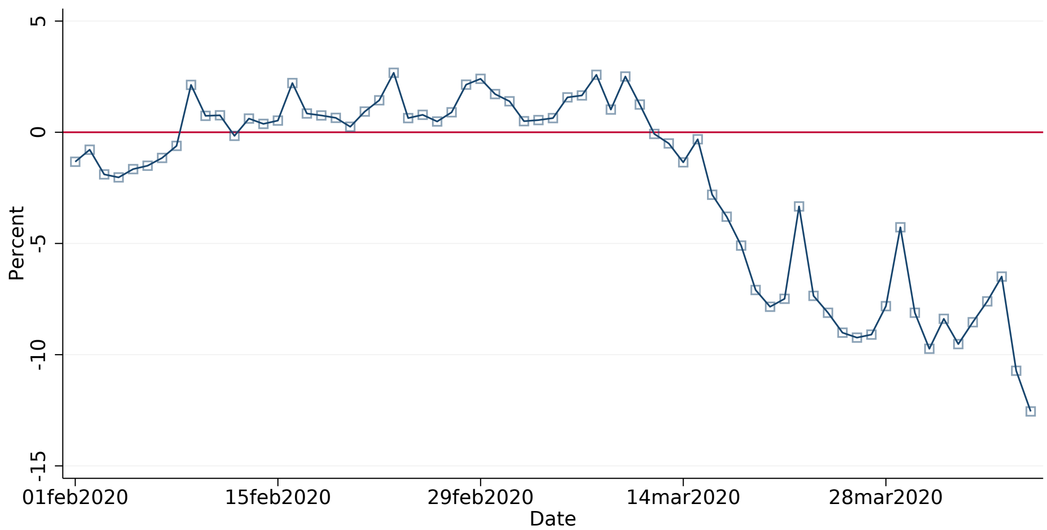

Note: Coefficients are normalized to be mean zero between 1 February and 21 February, 2020. The EU total is calculated by weighting country/zone-specific estimates in proportion to mean load in 2019.

Figure 3: Changes in Italian Electricity Consumption: 1 February - 6 April, 2020

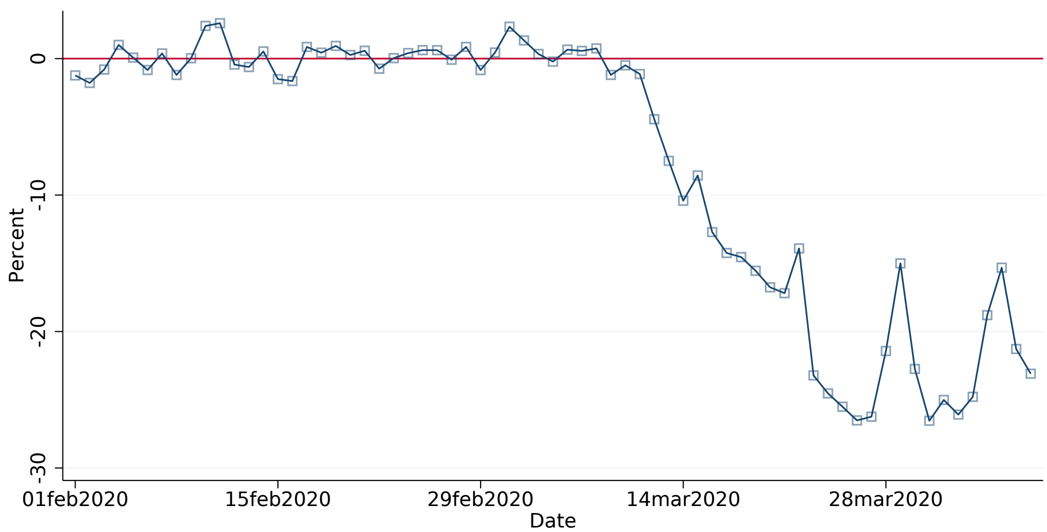

Note: Coefficients are normalized to be mean zero between 1 February and 21 February, 2020. The EU total is calculated by weighting country/zone-specific estimates in proportion to mean load in 2019.

Table 1: Percent Change in Adjusted Electricity Consumption: Week Ending 4 April 2020

Country	Percent Change
Austria	-6.6
	(1.2)
Belgium	-10.2
	(.88)
Bulgaria	2.9
	(1.7)
Czech Republic	-7.9
	(1.7)
Germany (50Hz)	-1.3
	(3.4)
Germany (Amprion)	-8.0
	(1)
Germany (TenneT GER)	-0.3
	(2.8)
Germany (TransnetBW)	-10.1
	(1.9)
Denmark	-0.6
	(1.5)
Estonia	-3.2
	(1.3)
Spain	-14.9
	(1.6)
France	-8.0
	(1.3)
Croatia	-5.5
	(2.8)
Hungary	-9.3
	(2)
Italy	-23.1
	(2)
Latvia	-2.1
	(.82)
Norway	2.4
	(1.3)
Poland	-7.4
	(2)
Portugal	-7.7
	(1.7)
Romania	-9.5
	(1.7)
Sweden	3.2
	(.66)
Slovakia	-9.6
	(1.6)
United Kingdom	-7.0
	(2.4)

Note: Percent changes are calculated relative to the two-week baseline period from 2 February, 2020 to 15 February, 2020. Standard errors in parentheses are clustered at the country/zone-month. Estimates are from separate regressions.

Figure 4: Changes in U.K. Electricity Consumption: 1 February - 6 April, 2020

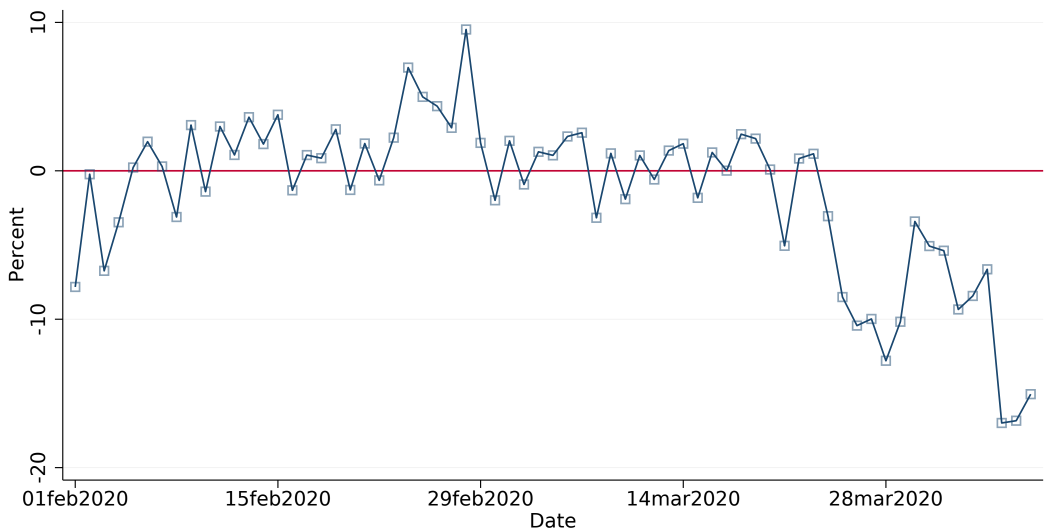

Note: Coefficients are normalized to be mean zero between 1 February and 21 February, 2020. The EU total is calculated by weighting country/zone-specific estimates in proportion to mean load in 2019.

5 Conclusion

Electricity consumption in Europe is down roughly 10% since various shutdown policies have been implemented to slow the transmission of COVID-19. The drop in power consumption broadly reflects the timing and stringency of lockdowns.

How should one interpret these results in the context of the wider economy? Preliminary work on the historical record indicates an approximately 1-for-1 short-term relationship with standard economic indicators (Cicala (2020)), which would suggest a dire situation if it were to persist for an extended period of time. On the other hand, the nature of disease transmission-slowing shutdowns is not quite the same as other types of economic shocks. Dingel and Neiman (2020) have recently calculated that roughly one third of U.S. jobs can prospectively be done from home. If work from home requires the same energy profile as in a centralized office, then changes in system-level consumption would continue to accurately reflect changes in economic activity. If the same amount of work can be accomplished at home with less energy consumption, then the fall in employment or GDP may not be as large. Similarly, differences between home and work in whether they use gas or electricity for heating and cooling would cause the indices to diverge. These considerations may suggest modest adjustments, but the overall picture is one of a sudden and large decline in economic activity in Europe.

References

Chen, Wei, Xilu Chen, Chang tai Hsieh, and Zheng (Michael) Song, "A Forensic Examination of China's National Accounts," *Brookings Papers on Economic Activity*, 2019, (Spring).

Cicala, Steve, "Electricity Consumption as a Real Time Indicator of Economic Activity," 2020.

Dingel, Jonathan I. and Brent Neiman, "How Many Jobs Can be Done at Home?," *Covid Economics: Vetted and Real-Time Papers*, 2020, *1* (1), 16–24.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil, "Measuring Economic Growth from Outer Space," *American Economic Review*, 2012, *102* (2), 994–1028.

Lyu, Changjiang, Kemin Wang, Frank Zhang, and Xin Zhang, "GDP management to meet or beat growth targets," *Journal of Accounting and Economics*, 2018, *66* (1), 318–338.