NBER WORKING PAPER SERIES

COVID19 AND THE MACROECONOMIC EFFECTS OF COSTLY DISASTERS

Sydney C. Ludvigson
Sai Ma
Serena Ng

Working Paper 26987
http://www.nber.org/papers/w26987

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2020

The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Board, the Federal Reserve System, or the National Bureau of Economic Research. Ng acknowledges support from the National Science Foundation under grant SES-1558623.

Covid19 and the Macroeconomic Effects of Costly Disasters
Sydney C. Ludvigson, Sai Ma, and Serena Ng
NBER Working Paper No. 26987
April 2020
JEL No. E17,I0,I3

ABSTRACT

The outbreak of covid19 has significantly disrupted the economy. This note attempts to quantify the macroeconomic impact of costly and deadly disasters in recent US history, and to translate these estimates into an analysis of the likely impact of covid19. A costly disaster series is constructed over the sample 1980:1-2019:12 and the dynamic impact of a costly disaster shock on economic activity and on uncertainty is studied using a VAR. Unlike past natural disasters, covid19 is a multi-month shock that is not local in nature, disrupts labor market activities rather than destroys capital, and harms the social and physical well being of individuals. Calibrating different shock profiles to reflect these features, we find that the effects of the event last from two months to over a year, depending on the sector of the economy. Even a conservative calibration of a 3-month, 60 standard deviation shock is forecast to lead to a cumulative loss in industrial production of 12.75% and in service sector employment of nearly 17% or 24 million jobs over a period of ten months, with increases in macro uncertainty that last five months.

Sydney C. Ludvigson
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10002
and NBER
sydney.ludvigson@nyu.edu

Serena Ng
Department of Economics
Columbia University
420 West 118th Street
New York, NY 10027
and NBER
Serena.Ng@columbia.edu

Sai Ma
Federal Reserve Board,
C Ave & 20th Street NW
Washington, DC 20551
sai.ma@frb.gov

1 Introduction

Short term fluctuations in a typical economic model are presumed to be driven by random shocks to preferences, factor inputs, productivity, or policies that directly impact the supply or demand of goods and services. While there is some scope for considering exogenous shocks due to natural disasters such as earthquakes and tsunamis, such "conventional" disaster shocks are typically assumed to be short-lived, with an initial impact that is local in nature. It is only when these shocks propagate across sectors, states, and countries that the aggregate effects are realized. These more typical disaster shocks are assumed to disrupt the economy, but not all affect the social and physical well being of individuals. This paradigm of modeling disaster-type exogenous economic fluctuations is, however, not suited for studying the impact of the coronavirus (COVID19). COVID19 is a multi-period shock that simultaneously disrupts supply, demand, and productivity, is almost perfectly synchronized within and across countries, and wherein health, social, and economic consequences are cataclysmic not just for the foreseeable few weeks after the crisis, but potentially for a long time period.

The ability to design policies to mitigate the economic impact of COVID19 requires reference estimates of the effects of the shock. This note provides some preliminary estimates of these effects. Our analysis has two ingredients. The first are a costly disaster (CD) and a deadly disaster (DD) time series, which we construct to gauge the magnitude of the COVID19 shock in relation to past disasters. The second is an analysis of the dynamic impact of a CD shock on different measures of economic activity and on a measure of macroeconomic uncertainty using a linear vector autoregression as the baseline. To shed light on more pessimistic scenarios, the baseline model is modified to allow for nonlinearities. We then design different shock profiles to represent COVID19 as a multi-period, large, and broad-based macroeconomic shock.

In summary, we find that the macroeconomic impact of COVID19 is larger than any catastrophic event that has occurred in the past four decades. Even under a fairly favorable scenario of a three month, 60 standard deviation shock, the estimates suggest that there will be a peak loss in industrial production of 5.82% and in service sector employment of 2.63% respectively, which translates into a cumulative ten-month loss in industrial production of 12.75% and an employment loss of nearly 17%, or 24 million jobs, before recovery starts. Results based on a non-linear model with a state-dependent intercept suggest steeper and longer losses for employment and scheduled flight departures but nonetheless indicate a possible rebound for after about a year. In all scenarios, there will be five or more months of heightened macroeconomic uncertainty.

2 Data and Methodology

Our analysis is based on monthly data on disasters affecting the U.S. over the last forty years taken from two sources. The first is NOAA, which identifies 258 costly natural events ranging from wildfires, hurricanes, flooding, to earthquakes, droughts, tornadoes, freezes, and winter storms spanning the period 1980:1-2019:12 for $T = 480$ data points, of which 198 months have non-zero cost values.[1] These data, which can be downloaded from `ncdc.noaa.gov/billions/events`, record both the financial cost of each disaster as well as the number of lives lost over the span of each disaster. As explained in Smith and Katz (2013), the total costs reported in NOAA are in billions of 2019 dollars and are based on insurance data from national programs such as on flood insurance, property claims, crop insurance, as well as from risk management agencies such as FEMA, USDA, and Army Corps. We take the CPI-adjusted financial cost series as provided by NOAA, and mark the event date using its start date. To obtain the monthly estimate, we sum the costs of all events that occurred in the same month.

The second source of data is the Insurance Information Institute (III), which reports the ten costliest catastrophes in the US reported in 2018 dollars. The data, available for download from `www.iii.org/table-archive/2142`, covers property losses only. Thus the cost for the same event reported in the III dataset is lower than that reported in the NOAA dataset. But in agreement with the NOAA data, the III dataset also identifies Hurricane Katrina as the most costly disaster in US history. The III dataset is of interest because it records 9/11 as the fourth most costly catastrophic event, arguably the most relevant historical event for the purpose of this analysis given the large loss of life involved. But as 9/11 is not a natural disaster, it is absent from the NOAA data. We therefore use the III data to incorporate the event into the NOAA data. To deal with the fact the two data sources define cost differently, we impute the cost of September 11 as follows. We first compute the ratio of cost (in 2018 dollars) of Katrina relative to 9/11 from the III data, which is 1.99. We then divide the cost of Katrina in NOAA data by this ratio to get the insurance-based estimate of 9/11 cost in the same units as those reported in NOAA.

An important limitation of the data needs to be made clear at the outset. With the exception of Hurricane Sandy, the natural disasters in our data have been concentrated in the southern states with FL, GA, or LA having experienced disasters most frequently. However, industrial production is concentrated in the New England area, the Great Lakes area, the mid-West, and the Mid-Atlantic States which have been much less impacted by natural disasters. The data

[1] The number of months with nonzero cost values is less than the number of events is because there were many events that occurred in the same month, and we sum them up.

Figure 1: Time Series of Disaster Series: 1980:1-2019:12

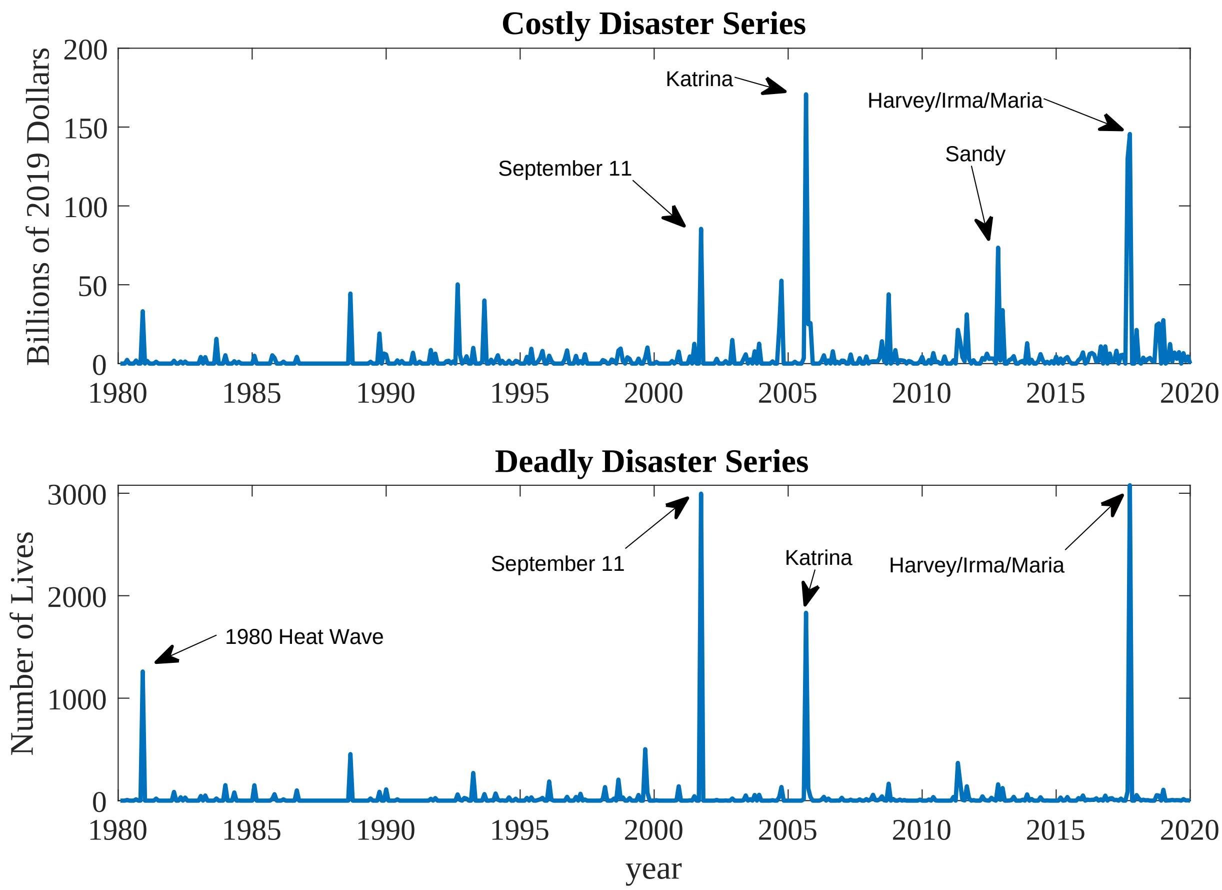

may not be able to establish a clear relation between industrial production and disasters.

The cost measures are based on monetary damages but do not include the value of lives lost, which is another measure of the severity of the disaster. Separately reported in NOAA is the number of deaths associated with each event. Since the number of deaths directly linked to 9/11 is known to be 2,996, we are able to construct a disaster series that tallies the number lives lost for all 259 events considered in the analysis.[2]

Figure 1 plots the resulting *costly disaster* (CD) series, in units of billions of 2019 dollars, and the *deadly disaster* (DD) series, in units of lives lost. Notably, there are four events in the CD series that stand out: Hurricanes Katrina in 2005, Harvey/Irma/Maria in 2017, Sandy in 2012, and 9/11 in 2001. As a point of reference, the value of CD at these four events are at least four standard deviations away from the mean of the series. In terms of the number of deaths, the sum of the DD series over the sample is 14,221, but three events, namely, Hurricane Harvey/Irma/Maria, 9/11, and Katrina, accounted for nearly two-thirds of the total deaths.

[2]Source: `https://www.cnn.com/2013/07/27/us/september-11-anniversary-fast-facts/`

Thus, both disaster series are heavy-tailed. We will return to this point below.

We will also make use of two additional pieces of information from these two data files. The first is the number of states being affected as reported in III. For example, Katrina directly impacted six states: AL, FL, GA, LA, MS, TN, while the direct impact of 9/11 was local to the city of New York and the D.C. region. The second is the duration of the event. As reported in NOAA, Katrina was a five-day event, Superstorm Sandy was a two-day event, while the 9/11 attack was a one-day event. From 1980 to 2019, the average duration of an event is 40 days and ranges from one day (e.g., 9/11 and 2005 Hurricane Wilma) to one year (e.g., the 2015 Western Drought). These statistics will be helpful in calibrating the size of the COVID19 shock subsequently.

To estimate the macroeconomic impact of a disaster shock, we begin as a baseline with a six-lag, $n = 3$ variable vector autoregression (VAR) in

$$\mathbf{X}_t = \begin{bmatrix} \mathrm{CD}_t \\ Y_t \\ U_t \end{bmatrix} = \begin{bmatrix} \text{Costly Disaster} \\ \log\text{ (Real Activity)} \\ \text{Uncertainty} \end{bmatrix},$$

where CD is our costly disaster series just constructed, U is a measure of uncertainty, and Y is one of four measures of real activity that will be discussed below. The reduced form VAR is

$$A(L)X_t = \eta_t.$$

The reduced form innovations η_t are related to mutually uncorrelated structural shocks e_t by

$$\eta_t = Be_t, \quad e_t \sim (0, \Sigma)$$

where Σ is a diagonal matrix with the variance of the shocks, and $\mathrm{diag}(B) = 1$. For identification, B is assumed to be lower triangular. In particular, the covariance matrix of VAR residuals is orthogonalized using a Cholesky decomposition with the variables ordered as above. The CD series is ordered first given that the disaster events are by nature exogenous. The resulting structural VAR (SVAR) has a structural moving average representation

$$X_t = \Psi_0 e_t + \Psi_1 e_{t-1} + \Psi_2 e_{t-2} + \dots, \tag{1}$$

with the impact effect of shock j on variable j measured by the j-th diagonal entry of Ψ_0, which is also the standard deviation of shock j. The dynamic effects of a one time change in e_t on X_{t+h} are summarized by the Ψ_h matrices which can be estimated directly from the VAR using Bayesian methods under flat priors, or by the method of local projections due to Jorda (2005). The goal of the exercise is to trace out the effect of COVID19 on itself, on economic activity Y over time, and on macroeconomic uncertainty U. This amounts to estimating the first columns of the 3 by 3 matrix Ψ_h at different horizons h.

We will consider four measures of real activity Y: industrial production (IP), initial claims for unemployment insurance (IC), number of employees in the service industry (ESI), and scheduled plane departures (SFD). The first three variables are taken from FRED, and the last from the Bureau of Transportation Statistics. IP is a common benchmark for economic activity, while unemployment claims are perhaps the most timely measure of the impact on the labor market. In the data, initial claims one month after Katrina (i.e., September 2005) increased by 13.3% compared to its level the previous year. The variable ESI is considered because non-essential activities such as going to restaurants, entertainment, repairs, and maintenance can be put on hold in the event of a disaster, and these are all jobs in the service sector. Disasters tend to disrupt travel due to road and airport closures. Data constraints limit attention to air traffic disruptions, as measured by the number of scheduled flight departures, SFD.

3 Results

For each measure of Y, we estimate a VAR and compute the response coefficient Ψ_h scaled so that it corresponds to a one standard deviation increase in the innovation to CD. Since the dynamic responses of CD and U to a CD shock are insensitive to the choice of Y and U, we only report these two impulse response functions using the VAR with IP as Y and the JLN measure of uncertainty constructed in Jurado, Ludvigson, and Ng (2015) as U. In all results to follow, the blue line depicts the median response and the dotted lines refer to 68 percent confidence bands.

The top left panel of Figure 2 shows that the impact of CD shock on itself dies out after two months, suggesting that the CD is a short-memory process that does not have the autoregressive structure typically found in SVARs for analyzing supply and demand shocks. The top right panel of Figure 2 shows that JLN uncertainty rises following a positive CD shock, and the heightened uncertainty persists for three months. The bottom panel replaces the JLN measure of macro uncertainty by the measure of LMN financial uncertainty developed in Ludvigson, Ma, and Ng (2019). A CD shock raises financial uncertainty for one month but quickly becomes statistically insignificant. The bottom right panel uses the measure of policy uncertainty (EPU) in Baker, Bloom, and Davis (2016). A costly disaster shock increases policy uncertainty for about three months, similar to the duration of the impact on JLN uncertainty. In both cases, uncertainty is highest one month after the shock. These results suggest that short-lived disasters have statistically significant adverse effects on uncertainty that persist even after the shock subsides.

Figure 2: Dynamic Response of CD and U to a σ Shock

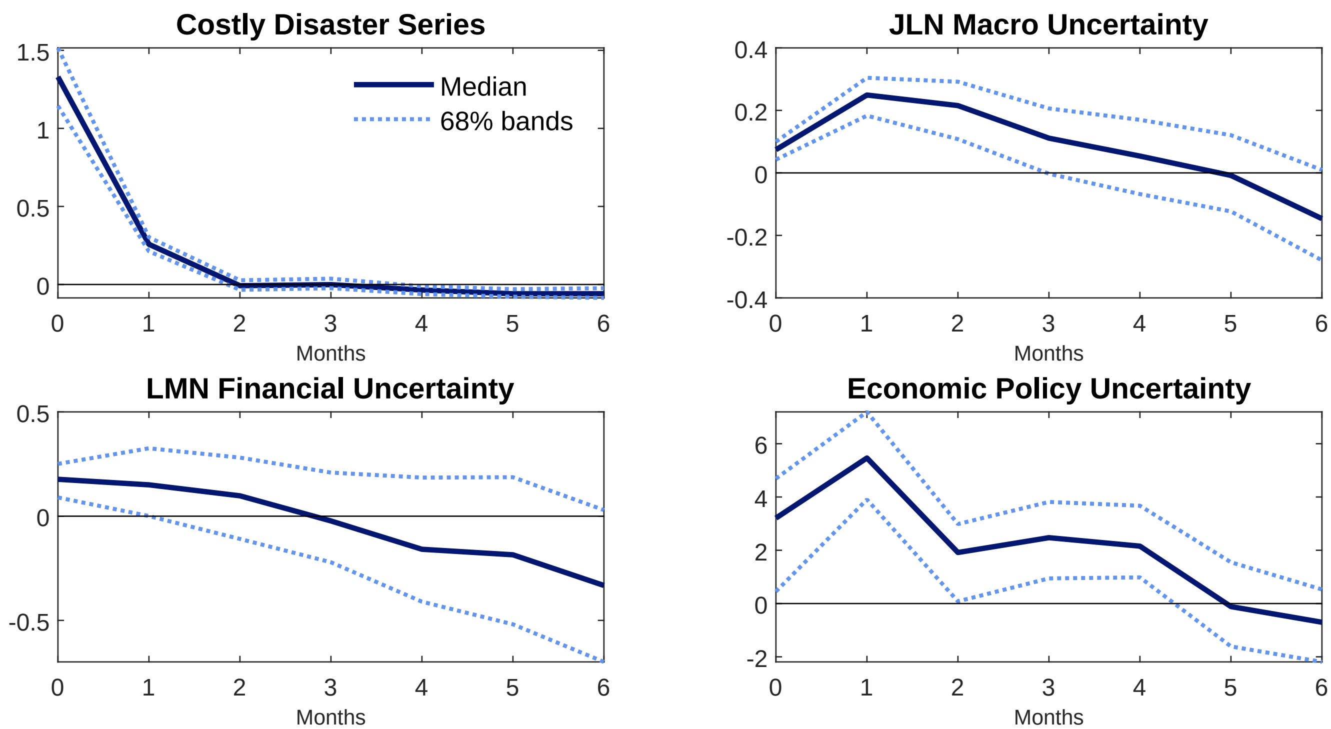

Note: The figure plots the dynamic responses to costly disaster shock. The posterior distributions of all VAR parameters are estimated using Bayesian estimation with flat priors and the 68% confidence bands are reported in dotted lines. The sample spans 1980 Jan. to 2019 Dec.

Next, we consider the effect of a one standard deviation CD shock on four measures of Y. The left top panel of Figure 3 shows that monthly IP immediately drops by 0.05% on impact but becomes statistically insignificant after two months, which as seen from Figure 2, is also the duration needed for the CD series to return to zero. There is, however, some evidence of a strong rebound in the economy but the effect is not statistically well determined. The small effect of CD on IP may seem surprising, but one possible explanation is that in our data, the natural disasters have not had much direct impact on industrial production. The top right panel shows that a CD shock triggers a statistically significant rise in unemployment claims IC for about three months with a statistically significant decline in claims (ie. a rebound in employment) thereafter.

The bottom left panel of Figure 3 shows that a CD shock leads to an immediate and statistically significant drop in the number of employed workers in the service industry, ESI. However, unlike results using IP and IC as Y, the baseline ESI response is more persistent, with the effect bottoming out at about 4 months. It is worth noting that ESI is a national measure of service employment and may mask the higher impact in some regions. The bottom right panel shows that a CD shock forces an immediate and persistent decline in the number of scheduled flights, SFD. Of all the measures of real activity, the impact effect of a CD shock on SFD is not only the largest, but also more sustained. Though recovery follows right after the shock, the process is slow, taking up to six months for the effect to become statistically insignificant.

Figure 3: Dynamic Response of Real Activities to a σ Shock

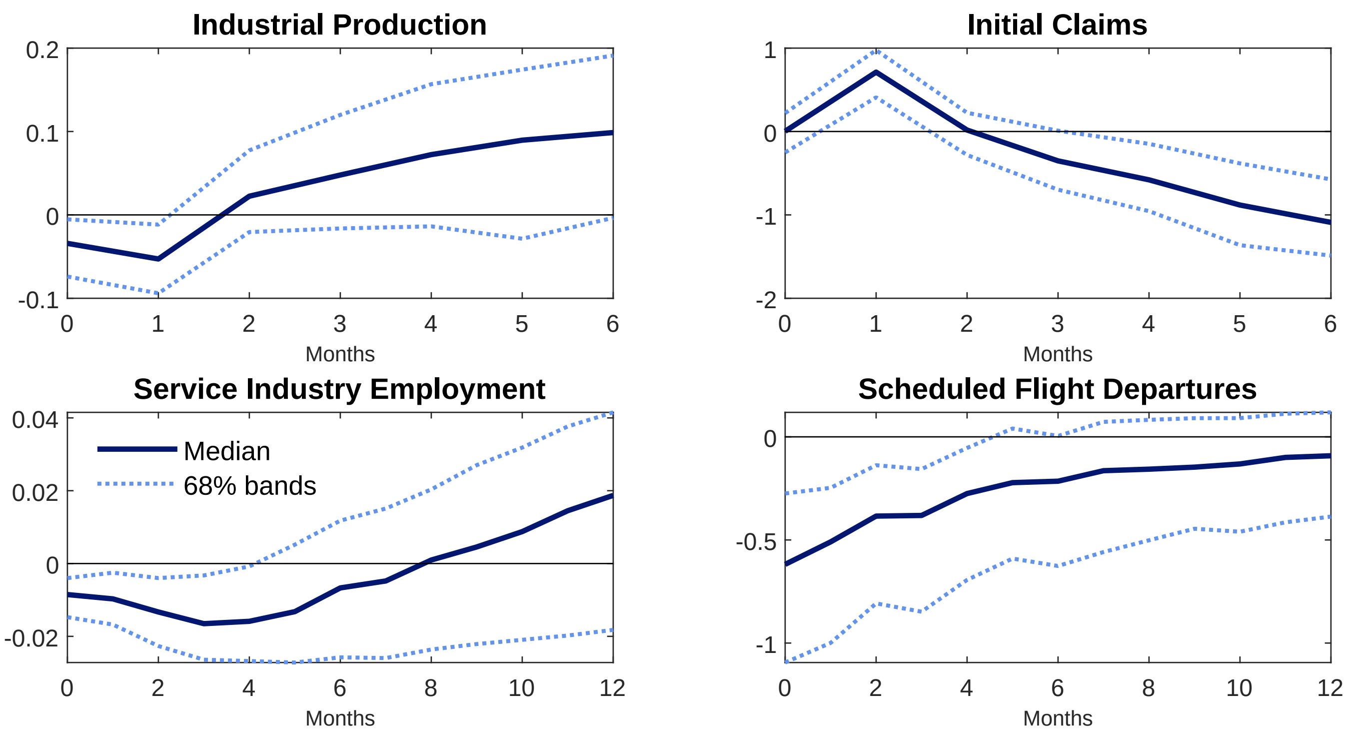

Note: The figure plots the dynamic responses to costly disaster shock. The posterior distributions of all VAR parameters are estimated using Bayesian estimation with flat priors and the 68% confidence bands are reported in dotted lines. The sample spans 1980 Jan. to 2019 Dec.

Taken together, this baseline estimation using pre-COVID19 data suggests that a one-period, one-standard-deviation increase in CD will have statistically significant adverse effects on real economic activity. Though there are variations in how long the impact will last, for all four real activity measures considered, the effects of the one period shock will die out within a year.

However, COVID19 differs from historical disasters in several dimensions. For one thing, the initial impact of the historical disasters had been local in terms of both the geographical area and population affected. In fact, never in the 30 years of data was there a disaster that involved more than one of the five largest states in the country simultaneously. For another, the historical disasters were short-lived, and with the exception of a drought that lasted over a year, they have, on average, only been one month long. Even with 9/11, the North American airspace was closed for a few days while Amtrak stopped service for two days, but activity resumed by September 14, albeit gradually.

The same cannot be said of COVID19. First, COVID19 is a global pandemic and the effects traverse across states and countries. As of April 1, 2020, 93% of the world population lives in countries with restricted travel.[3] Second, the most disastrous events in history in terms of loss of lives were Katrina and 9/11, but the number of deaths due to COVID19 already exceeded 5,000 at the time of this writing, higher than the deaths due to Katrina and 9/11 combined. Third, one month into the pandemic, the crisis had yet to reach its peak. Moreover, there is a good deal of uncertainty as to whether normalcy will return by the summer of 2020. Fourth, social distancing was not imposed in past disasters, and Gascon (2020) documents that the consequence of social distancing may be particularly harsh for those employed in the service sector. Fifth, past disasters created destruction in physical capital, while COVID19 creates no such damage. Instead, the labor force is constrained from working efficiently, and resources are diverted to unanticipated uses. Finally, as mentioned above, industrial production was not severely impacted by past natural disasters. These considerations suggest that the dynamic effects of CD need to be computed for shock profiles that reflect COVID19.

4 Effects of Large and Prolonged Shocks

This section considers dynamic responses to multi-period large shocks. To do so, we take as a benchmark that Katrina was an 11 standard deviation (σ) shock while 9/11 was a 5.5σ shock. So how big is COVID19? By the end of March 2020, 10 million Americans had made initial unemployment insurance claims, which is a 900% increase compared to February 2020, comparable in magnitude to that during the Great Depression. Furthermore, COVID19 is now projected to kill 100,000 to 240,000 Americans, more fatalities than the Vietnam War (90,229) and Korea War (54,246) combined.[4] We thus consider shocks that are larger than 11 standard deviations.

Next, we need to allow for the possibility that COVID19 is not an impulse of one-month. But what is a plausible duration? Ideally, this is the life of the virus which is not only unobservable, but potentially endogenous. Now a COVID19 shock can be thought of as an economic shock that

[3] Source: https://www.pewresearch.org/fact-tank/2020/04/01/more-than-nine-in-ten-people-worldwide-live-in-countries-with-travel-restrictions-amid-covid-19/

[4] `https://www.cnbc.com/2020/04/01/coronavirus-could-kill-more-americans-than-some-wars.html`

constrains consumers and producers from conducting economic activities. We use the expected duration of the 'stay-at-home' policy as the government's expected duration of the shock. At this moment, there is little doubt that this number is at least two months.

Let $\mathbb{X}^t$ collects all information in X at time t and at all lags. From the moving-average representation of the SVAR given in (1), we see that if there are two consecutive shocks of one standard deviation, the dynamic response of X_{t+h} is

$$\mathbb{E}\Big[X_{t+h}\big|e_{1t} = \sigma, e_{1t-1} = \sigma; \mathbb{X}^t\Big] - \mathbb{E}\Big[X_{t+h}\big|e_{1t} = 0, e_{1t-1} = 0; \mathbb{X}^t\Big] = \Psi_h + \Psi_{h+1}.$$

If the shock in t is of size $.5\sigma$, and the one at $t+1$ is of size 2σ, the desired response matrix is $.5\Psi_h + 2\Psi_{h-1}$. Scaling and summing the Ψ_h coefficients allows us to evaluate all the dynamic responses to each of the shocks at a magnitude deemed appropriate. The idea is akin to the one used in Box and Tiao (1975) to study the effect of interventions on a response variable in the presence of different dependent noise structure, or the innovational outlier model studied in Fox (1972). We are only interested in the effect of a disaster shock and so only need to estimate the first column of Ψ_h for $h = 1, ...H$.

Figure 4: Dynamic Response of CD and U to a one σ Shock

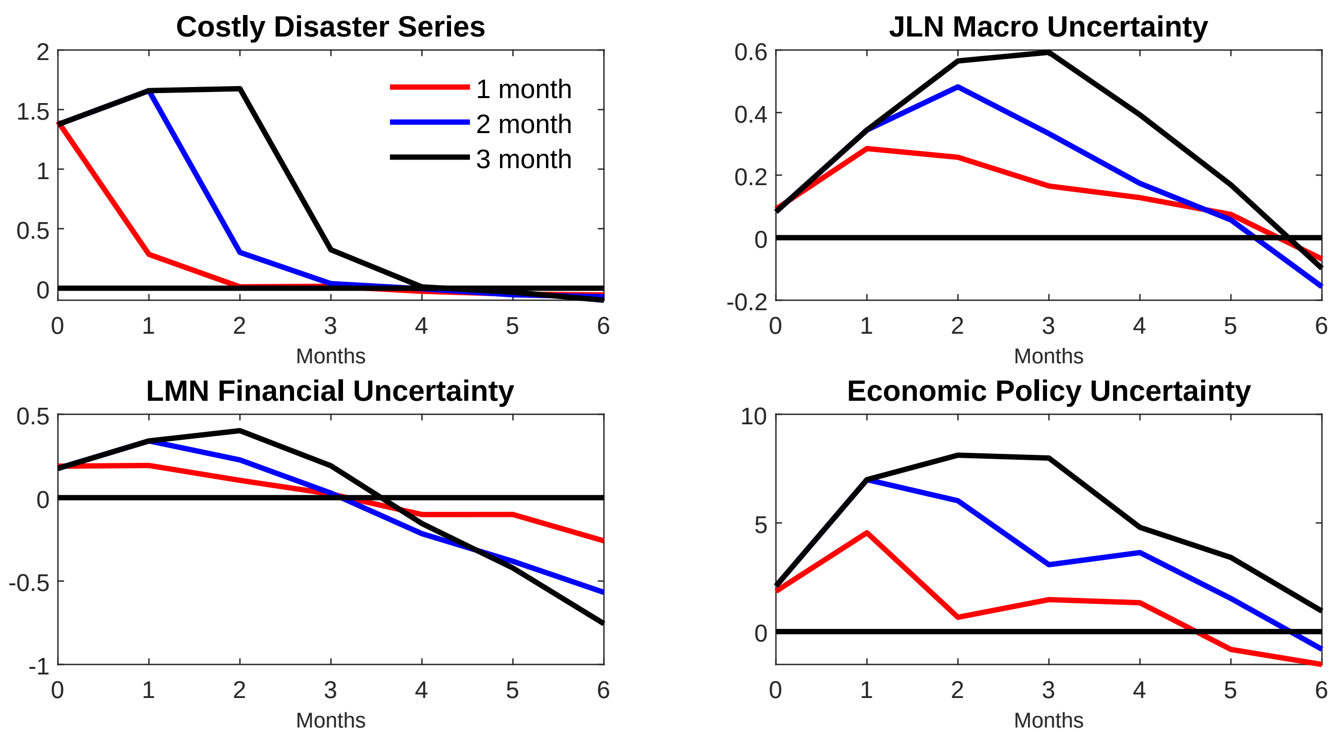

Note: The figure plots the dynamic responses to multi-period costly disaster shocks. The sample spans 1980 Jan. to 2019 Dec.

To separate the effect of a multiple period shock from that of a multi-standard-deviation shock, Figure 4 reports the response of CD and U, similar to Figure 2, except that there are now consecutive one-standard deviation shocks.[5] The red line is the same as the one period shock reported in Figure 2 and serves as a benchmark. Evidently, the CD series now requires three months to die out after a two-period shock, and four months after a three-period shock. The effects on all measures of uncertainty become larger and more persistent. Taking the JLN measure as an example, U peaks after three months instead of one, and is four times larger.

Figure 5: Dynamic Response of Real Activities to a one σ Shock

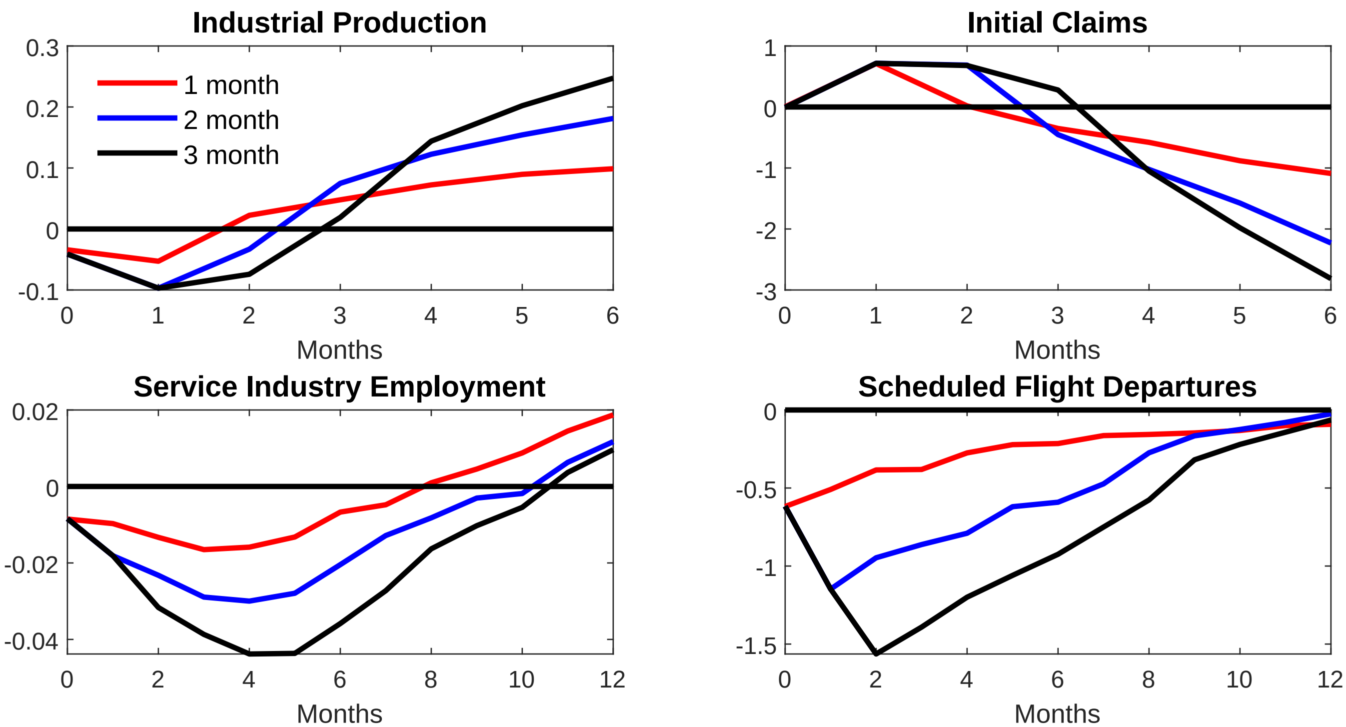

Note: The figure plots the dynamic responses to multi-period costly disaster shocks. The sample spans 1980 Jan. to 2019 Dec.

Figure 5 reports the dynamic responses of the four measures of Y to the multi-period shock of one standard deviation each period. The red lines are identical to the ones plotted in Figure 3 for a single period shock. For IP, the adverse effects are prolonged but are not significantly magnified. For IC, the maximum increase is the same in the multi-period shock as it is for a single period, and always occurs one month after the shock. However, multi-period shocks slow the time to recovery from two months to four. For ESI and SPD, there is a clear amplification effect due to consecutive shocks. At the worst of times, employment loss in the service sector is tripled that due to a one-shot shock, and the series is not back to control for well over three

[5] To avoid clutter, the confidence bands are not plotted. Note that for all except that ESI and SFD, the responses are well determined up to about three periods plus the duration of the shock. For ESI and SFD, they are well determined up to about six periods plus the duration of the shock.

quarters. Similarly, instead of an immediate recovery, multi-period shocks reduce scheduled flight departures by two more months before a slow recovery begins.

To understand the effects of large, multi-period shocks, note first that a large shock shifts up the dynamic responses relative to a one-standard-deviation shock, while a multi-period shock shifts the dynamic responses to the right. Since the effects are multi-dimensional, there are many ways to summarize them. We report in Table 1 the maximum response in a 12-month period, where the location of the maximum can be inferred from Figure 5. Table 1 also reports the cumulative loss over the months with negative responses.[6] These maximum and cumulative losses are reported for 1, 10, 30, 60, 100 standard deviation shocks, and for shock durations of 1,2, and 3 months.

Our baseline profile of COVID19 is based on the fact that Hurricane Katrina, which was a 11σ shock, resulted in 1.8 million jobless claims,[7] which by way of comparison, is merely *one-six* of the unemployment claims recorded in March 2020. Assuming that non-essential travel restrictions will be in place for at least another month leads to a characterization of COVID19 as a 3-month 60σ shock. This estimate is conservative considering that Katrina directly impacted only four southern states, while COVID19 is affecting all states and is particularly lethal along the two coasts where the population is concentrated. Furthermore, Katrina lasted five days, not five weeks.

With this in mind, the bottom panel of Table 1 shows that a shock of duration three months and magnitude 50σ will lead to a maximum drop in industrial production of 5.8% occurring after one month, a 2.6% maximum loss in service sector employment (over 3 million jobs) occurring after four months, and a 94% reduction in scheduled flights after two months. The reduction in ESI is not trivial because over 75% of workers (or over 140 million) are employed in the service sector. The implied cumulative reduction of 16.77%, or loss of nearly 24 million service sector jobs before the onset of recovery is staggering.

It is of interest to ask how the responses change if shocks are spread over more periods. Figure 6 plots the dynamic responses of a $(60, 60, 60, 0, 0)\sigma$ shock profile in dark blue, which is the same in shape as Figure 5, but the magnitude is multiplied by 60 since we consider 60 instead of one σ shock. Plotted next in dotted blue is a five-month $(60, 40, 30, 30, 20)\sigma$ shock profile, We see that changing from a three to a five-month shock holding the size of the initial shock fixed does not change the dynamic responses in any significant way. Next, we consider a $(15, 30, 90, 30, 15)\sigma$ profile plotted in a thin blue line. The largest shock is now delayed to

[6] The cumulative responses could be overestimated because the response can be statistically zero at lags much earlier than the point estimate of the response crosses the zero line.

[7] Media coverage includes `https://money.cnn.com/2005/09/15/news/economy/initial\_claims/`

Table 1: Maximum Negative Response to Disaster Shock: Linear Model

	1-month Shock			
Shock Size	Industrial Production	Initial Claims	Service Employment	Flights
1σ	-0.06%	0.72%	-0.02%	-0.62%
10σ	-0.56%	7.23%	-0.15%	-6.17%
30σ	-1.68%	21.69%	-0.46%	-18.50%
60σ	-3.35%	43.37%	-0.92%	-37.00%
100σ	-5.59%	72.29%	-1.53%	-61.66%
Cumu. Loss (10σ)	-0.97%	7.16%	-0.90%	-34.02%
Cumu. Loss (60σ)	-5.82%	42.97%	-5.39%	-204.09%
Cumu. Loss (100σ)	-9.70%	71.62%	-8.98%	-340.20%
	2-month Shock			
Shock Size	Industrial Production	Initial Claims	Service Employment	Flights
1σ	-0.10%	0.72%	-0.03%	-1.15%
10σ	-0.97%	7.23%	-0.30%	-11.48%
30σ	-1.94%	14.46%	-0.60%	-22.96%
60σ	-5.82%	42.97%	-1.80%	-68.87%
100σ	-9.70%	72.29%	-3.00%	-114.79%
Cumu. Loss (10σ)	-1.71%	13.96%	-1.83%	-67.16%
Cumu. Loss (60σ)	-10.28%	83.76%	-10.98%	-402.98%
Cumu. Loss (100σ)	-17.13%	139.61%	-18.29%	-671.62%
	3-month Shock			
Shock Size	Industrial Production	Initial Claims	Service Employment	Flights
1σ	-0.10%	0.72%	-0.04%	-1.56%
10σ	-0.97%	7.23%	-0.44%	-15.64%
30σ	-1.94%	14.46%	-0.88%	-31.27%
60σ	-5.82%	42.97%	-2.63%	-93.81%
100σ	-9.70%	72.29%	-4.38%	-156.36%
Cumu. Loss (10σ)	-2.12%	16.68%	-2.79%	-99.76%
Cumu. Loss (60σ)	-12.75%	100.07%	-16.77%	-598.61%
Cumu. Loss (100σ)	-21.25%	166.78%	-27.94%	-997.60%

Note: This table shows maximum negative dynamic response of real activity from VAR $\mathbf{X}_t = (CD_t, Y_t, U_{Mt})'$ for one-month, two-month, and three-month shocks. The size of the shock is indicated in the first column. The "cumu. loss" is the sum of all negative (positive for IC) responses within 12 months. The sample spans 1980 Jan. to 2019. Dec.

Figure 6: Dynamic Response to Six Different Profiles

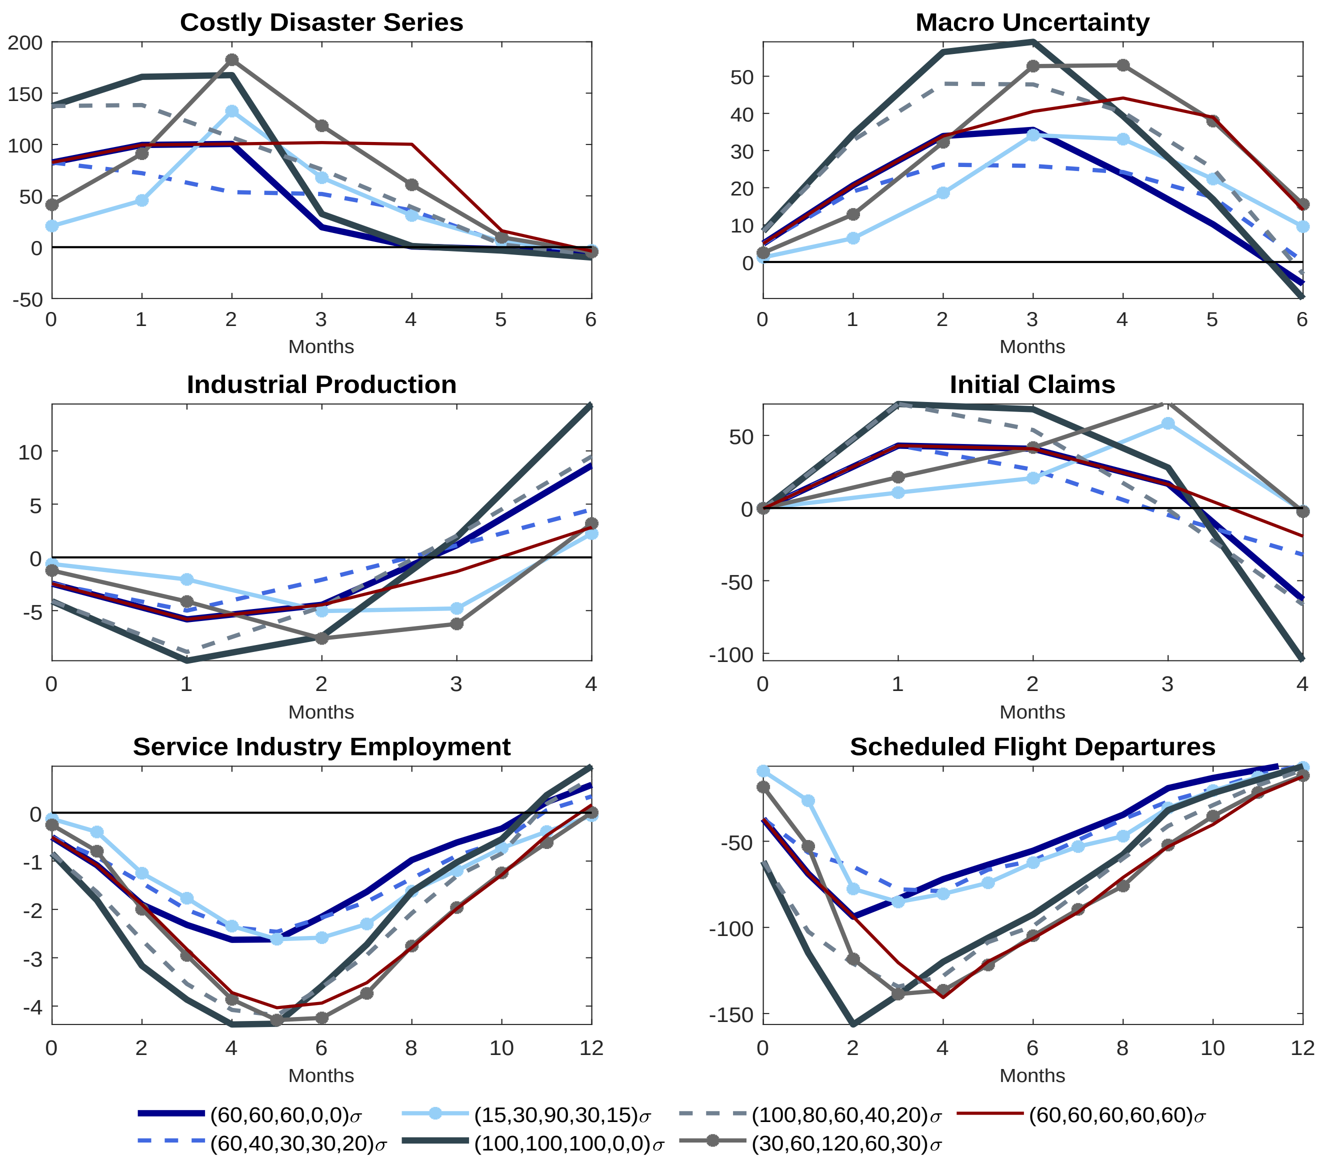

month three, but for all measures of Y, the shape of the dynamic responses are similar. The three cases considered so far have a cumulative magnitude of 180σ. The next case considered is a $(100, 100, 100, 0, 0)\sigma$ shock which has a cumulative magnitude of 300σ over three months. As seen from the thick black line, uncertainty is higher than the same shock with a smaller magnitude, and the response for all measures of Y are steeper than the previous three profiles. Plotted next in dotted gray is the same smooth decay as the dotted blue line except that the shocks are larger each period. Next, the profile that generates the thin gray line is the same as delayed peak that generates the thin blue line, except that the shocks are twice as big each period. For all measures of Y, the blue lines (with total shock magnitude of 300) are always larger losses than the gray lines (with total shock magnitude of 180). Finally, we also consider a five period shock of 60σ each period for a total of 300σ. It is not has powerful as the three period 100σ in gray, but it lasts longer. As seen from the line plotted in brown, the maximum effect in terms on ESI and SFD are recorded slightly after the $(100, 100, 100, 0, 0)\sigma$ profile in dark gray, but the effects are not quite as deep.

The picture that emerges from Figure 6 is that the total magnitude of the shock is a more important determinant of cumulative loses than the duration of the shock or the magnitude in any one period. In all profiles considered, the effects on U, IP, and IC peak and subside quickly, like the response of CD itself. The losses for ESI and SFD are more persistent, and though discouraging large, they do return to control after about one year.

5 Nonlinearities

While there were 259 disasters in our data, most of these were small. A linear model may underestimate the effect of large shocks. We therefore consider a model that allows the coefficients to be different for large disasters. Let S_t be an observable variable. We estimate a series of single equation regressions, one for each h, to obtain the dynamic response at lag $h \geq 1$:[8]

$$Y_{t+h} = \alpha_0 + \beta^h(L)'\mathbf{X}_{t-1}(L) + S_{t-1}\left(\delta_0^h + \delta_1^{h'} X_{t-1}\right) + e_{t+h},$$

where $S_t = \frac{\exp(-\gamma \text{DD}_t)}{1+\exp(-\gamma \text{DD}_t)}$ is a logistic function in the number of deaths normalized to be mean zero and variance one. After some experimentation, we set the vector of coefficients δ_1 to zero. In other words, the model has a state-dependent intercept with constant slope coefficients.

Figure 7 plots the dynamic responses to a one-period, one standard deviation shock constructed from the non-linear model. For IP and IC, the responses of the non-linear model (in

[8]This procedure has been called the "local projection" method by Jorda (2005).

red) are similar to the linear model (in blue). Both responses peak almost immediately after the shock. For SFD, the negative responses are larger and more persistent. The ESI losses are larger than those the linear model, but even in the non-linear model, the effects are statistically insignificant after one year.

Figure 7: Dynamic Response of Real Activities to a σ Shock: Non-linear Model

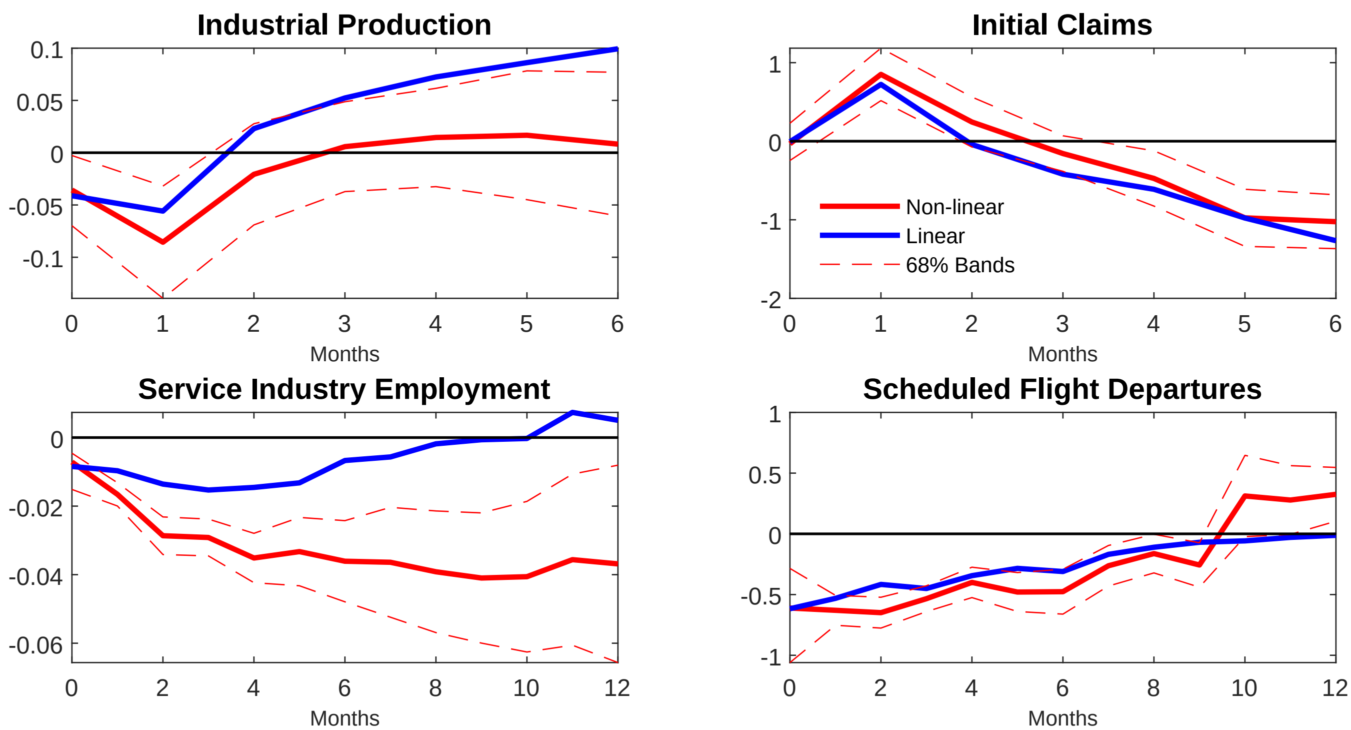

Note: The figure plots the dynamic responses to costly disaster shock from the non-linear model. The posterior distributions of all parameters are estimated using Bayesian estimation with flat priors and the 68% confidence bands are reported in dotted lines. The sample spans 1980 Jan. to 2019 Dec.

Table 2 summarizes the maximum and cumulative responses based on the non-linear model. Compared to estimates from linear model reported in Table 1, the maximum impact of the disaster shock is larger for all measures of activity, and particularly so when the shock extends more than one period. The baseline profile of a (3-month 60σ) shock now leads to a maximum one-month reduction in IP of 8.5%, a 113% reduction of scheduled flights, and service employment loss of in month eight of 7.7% which is roughly 10 million jobs. As of April 03, 2020, 8.5 million more people are on unemployment benefits than there were two weeks ago.[9] The cumulative losses are much larger than the linear scenario, with a 3-month 60σ shock generating a 22% drop in IP and a 63% drop in service sector employment.

[9] Source: `https://www.nytimes.com/2020/04/03/upshot/coronavirus-jobless-rate-great-depression.html`

Table 2: Maximum Negative Response to Disaster Shock: Non-linear Model

Shock Size	Industrial Production	Initial Claims	Service Employment	Flights
	1-month Shock			
Shock Size	Industrial Production	Initial Claims	Service Employment	Flights
1σ	-0.09%	0.85%	-0.04%	-0.65%
10σ	-0.86%	8.51%	-0.41%	-6.49%
30σ	-2.57%	25.52%	-1.23%	-19.47%
60σ	-5.14%	51.03%	-2.46%	-38.95%
100σ	-8.56%	85.05%	-4.10%	-64.91%
Cumu. Loss (10σ)	-1.06%	10.95%	-3.72%	-38.50%
Cumu. Loss (60σ)	-6.38%	65.69%	-22.30%	-231.01%
Cumu. Loss (100σ)	-10.63%	109.48%	-37.16%	-385.01%
	2-month Shock			
Shock Size	Industrial Production	Initial Claims	Service Employment	Flights
1σ	-0.12%	1.10%	-0.08%	-1.28%
10σ	-1.21%	10.97%	-0.84%	-12.78%
30σ	-3.64%	32.90%	-2.51%	-38.33%
60σ	-7.28%	65.79%	-5.03%	-76.67%
100σ	-12.13%	109.65%	-8.38%	-127.78%
Cumu. Loss (10σ)	-2.44%	20.06%	-7.26%	-80.69%
Cumu. Loss (60σ)	-14.62%	120.36%	-43.54%	-484.13%
Cumu. Loss (100σ)	-24.36%	200.59%	-72.56%	-806.88%
	3-month Shock			
Shock Size	Industrial Production	Initial Claims	Service Employment	Flights
1σ	-0.14%	1.10%	-0.13%	-1.89%
10σ	-1.42%	10.97%	-0.84%	-12.78%
30σ	-4.26%	31.92%	-3.83%	-56.66%
60σ	-8.53%	63.85%	-7.66%	-113.32%
100σ	-14.22%	109.65%	-12.76%	-188.87%
Cumu. Loss (10σ)	-3.74%	28.31%	-10.53%	-122.70%
Cumu. Loss (60σ)	-22.44%	169.44%	-63.18%	-736.22%
Cumu. Loss (100σ)	-37.40%	283.11%	-105.31%	-1228.03%

Note: This table shows maximum negative dynamic response of real activity from the nonlinear local projection of $\mathbf{X}_t = (CD_t, Y_t, U_{Mt})'$ for one-month, two-month, and three-month shocks. The size of the shock is indicated in the first column. The "cumu. loss" is the sum of all negative (positive for IC) responses within 12 months. The sample spans 1980 Jan. to 2019. Dec.

6 Conclusion

Based on the monthly data on costly disasters affecting the U.S. over the last forty years, we provide some preliminary estimates of the macroeconomic impact of COVID19. We find that even in a fairly conservative scenario without nonlinearities, large multiple-period shocks like COVID19 can create a 12.75% drop in IP, a loss in service employment of 17%, sustained reductions in air traffic, while macroeconomic uncertainty lingers for up to five months.

There are, of course, caveats to the analysis. First, COVID19 is different from past disasters in many ways, and the historical data may well over- or under-estimate the effects. As mentioned above, the disasters in history have not led to serious disruptions in industrial production. The relatively small loses found for IP must be interpreted in this light. Second, we have focused the dynamic responses under one year because the longer horizon results are not very well determined. This could be a consequence of the short-memory nature of disaster shocks. Furthermore, to the extent that the CD series is heavy-tailed, it is fair to question whether standard Bayesian sampling procedures or frequentist asymptotic inference based on normal errors are appropriate. Nonetheless, consideration of different profiles suggests that the total magnitude of the shock is a more important determinant of losses than the duration or magnitude in any single period. Our framework can be used to obtain rough and ready estimates of the cost of COVID19 under different assumptions.

References

Auerbach, A. J., and Y. Gorodnichenko (2012): "Measuring the output responses to fiscal policy," *American Economic Journal: Economic Policy*, 4(2), 1–27.

Baker, S. R., N. Bloom, and S. J. Davis (2016): "Measuring economic policy uncertainty," *The Quarterly Journal of Economics*, 131(4), 1593–1636.

Box, G. E., and G. C. Tiao (1975): "Intervention Analysis with Applications to Economic and Environmental Problems," *Journal of the American Statistical association*, 70(349), 70–79.

Fox, A. J. (1972): "Outliers in Time Series," *Journal of the Royal Statistical Society Series B*, 34:3, 350–363.

Gascon, C. (2020): "COVID-19: Which Workers Face the Highest Unemployment Risk?," *St. Louis Fed On the Economy*, https://www.stlouisfed.org/on-the-economy/2020/march/covid-19-workers-highest-unemployment-risk.

Jorda, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161–182.

Jurado, K., S. C. Ludvigson, and S. Ng (2015): "Measuring Uncertainty," *The American Economic Review*, 105(3), 117–1216.

Ludvigson, S. C., S. Ma, and S. Ng (2019): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?," *American Economic Journal: Macroeconomics,* forthcoming, http://www.econ.nyu.edu/user/ludvigsons/ucc.pdf.

Plagborg-Møller, M., and C. K. Wolf (2019): "Local projections and VARs estimate the same impulse responses," *Unpublished paper: Department of Economics, Princeton University*, p. 1.

Smith, A., and R. Katz (2013): "U.S. Billion-Dollar Weather and Climate Disasters: Data Sources, Trends, Accuracy and Biases," *National Hazards*, 67, 387–410.

Appendix: Local Projection versus VAR

We can alternatively compute the impulse response to a disaster shock using local projection (LP) method as in Jorda (2005),

$$Y_{t+h} = \alpha_0 + \sum_{l=1}^{L} \beta_{yl}^{h} Y_{t-l} + \sum_{l=1}^{L} \beta_{ul}^{h} U_{t-l} + \sum_{l=1}^{L} \beta_{dl}^{h} \mathrm{CD}_{t-l} + e_{t+h}, \text{ for } h = 1, 2...H \tag{2}$$

where Y_t is the log of real activity, and U_t is the macro uncertainty, CD_t is the costly disaster series as in the VAR. One advantage of using the local projection method is that it can readily incorporate higher-order or interaction terms and entails VAR results in the linear setting.

Figure 8 plots the VAR(6) median response (in blue), impulse responses using the LP method and associated 68 percent error bands (in dashed lines) when we use initial claims as a measure of real activity. Figure 8 demonstrates that, the local projection method yields very similar dynamic responses to VAR median responses, albeit non-identical.[10] When $h > 3$, the response of initial claims are slightly lower using local projection but both methods provide a similar picture that costly disaster drives up both initial claims and macro uncertainty in the near term.

Figure 8: Dynamic Response to Costly Disaster Shock: LP v.s. VAR

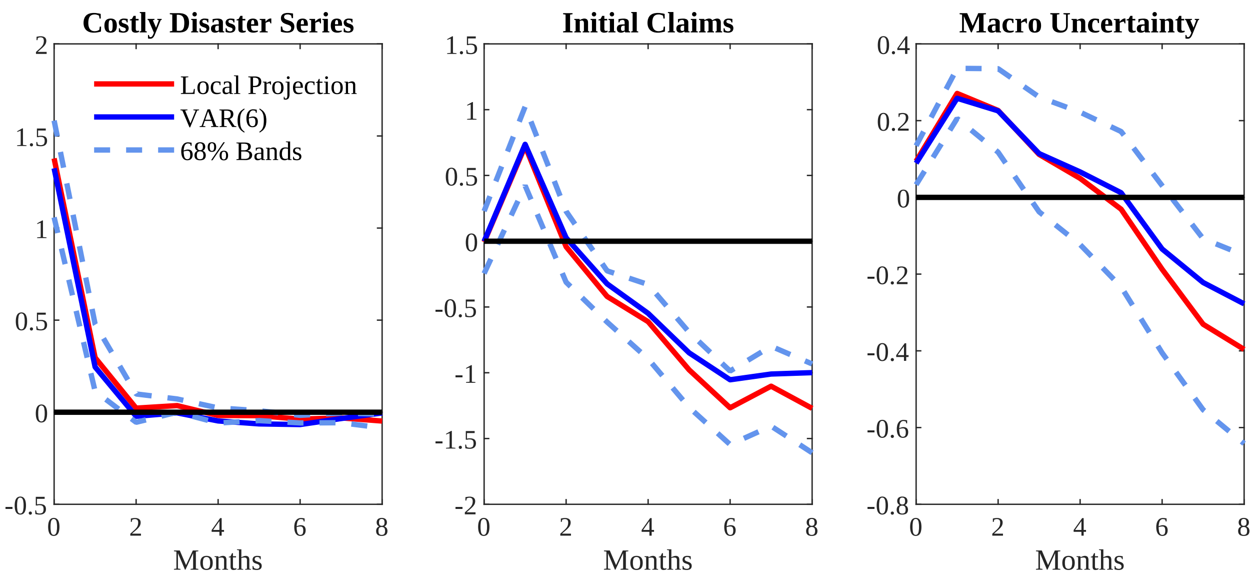

Note: The figure plots the impulse responses to costly disaster shock. The posterior distributions of all parameters are estimated using Bayesian estimation with flat priors and the 68% confidence bands are reported in dotted lines. The sample spans 1980 Jan. to 2019 Dec.

[10]This is due to the finite lag length. Plagborg-Møller and Wolf (2019) shows that VAR(∞) coincides with direct LP forecasts.