CENTRE FOR ECONOMIC POLICY RESEARCH

CEPR PRESS

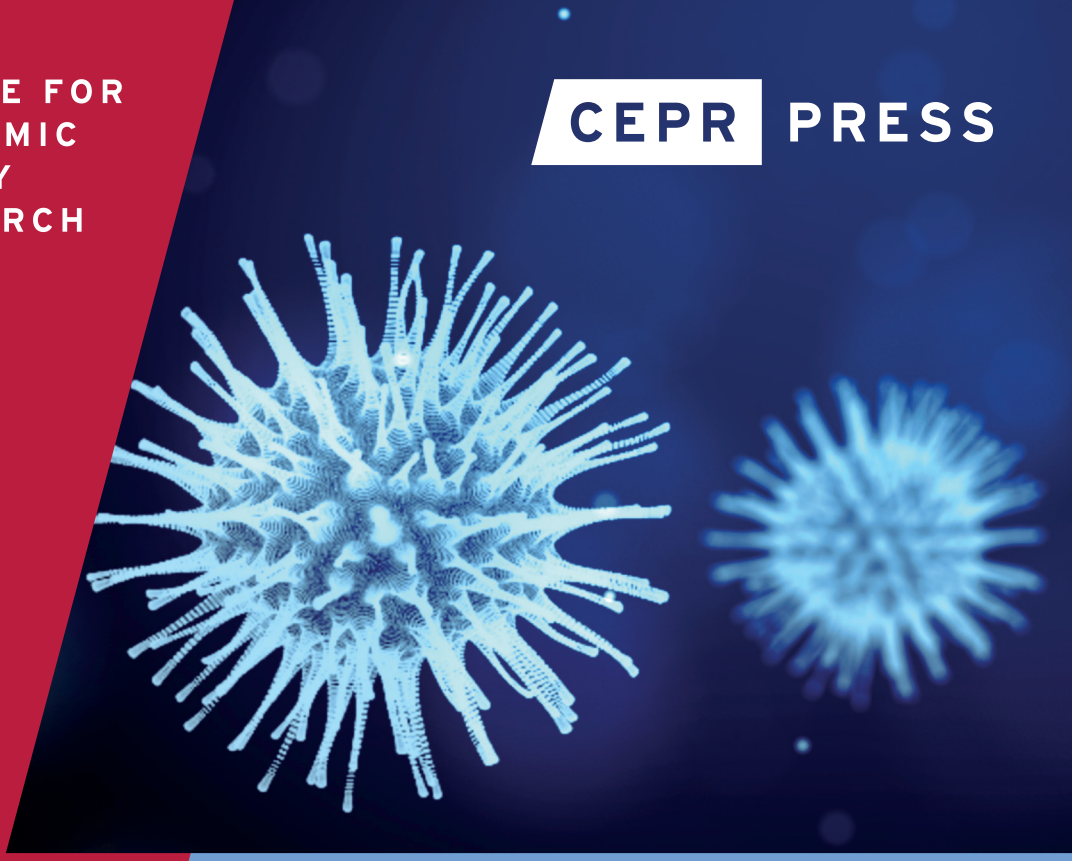

COVID ECONOMICS

VETTED AND REAL-TIME PAPERS

ISSUE 9
24 APRIL 2020

Covid Economics
Vetted and Real-Time Papers

Covid Economics, Vetted and Real-Time Papers, from CEPR, brings together formal investigations on the economic issues emanating from the Covid outbreak, based on explicit theory and/or empirical evidence, to improve the knowledge base.

Founder: Beatrice Weder di Mauro, President of CEPR
Editor: Charles Wyplosz, Graduate Institute Geneva and CEPR

Contact: Submissions should be made at https://portal.cepr.org/call-papers-covid-economics-real-time-journal-cej. Other queries should be sent to covidecon@cepr.org.

Editorial Board

Ethics

Covid Economics will publish high quality analyses of economic aspects of the health crisis. However, the pandemic also raises a number of complex ethical issues. Economists tend to think about trade-offs, in this case lives vs. costs, patient selection at a time of scarcity, and more. In the spirit of academic freedom, neither the Editors of *Covid Economics* nor CEPR take a stand on these issues and therefore do not bear any responsibility for views expressed in the journal's articles.

Covid Economics
Vetted and Real-Time Papers

Issue 9, 24 April 2020

Contents

CEPR PRESS

Covid Economics **Issue 9, 24 April 2020**

Covid-19 and the macroeconomic effects of costly disasters[1]

Sydney C. Ludvigson,[2] Sai Ma[3] and Serena Ng[4]

Date submitted: 18 April 2020; **Date accepted:** 19 April 2020; **Date revised:** 14 September 2020

Covid Economics 9, 24 April 2020: 1-24

The outbreak of COVID-19 has significantly disrupted the economy. This paper 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 COVID-19. A costly disaster series is constructed over the sample 1980:1-2020:04 and the dynamic impact of a disaster shock on economic activity and on uncertainty is studied using a VAR. While past natural disasters are local in nature and come and go quickly, COVID-19 is a global, multi-period event. We therefore study the dynamic responses to a sequence of large disaster shocks. Even in a fairly conservative case where COVID-19 is a 5-month shock with its magnitude calibrated by the cost of March 2020 Coronavirus relief packages, the shock is forecast to lead to a cumulative loss in industrial production of 20% and in service sector employment of nearly 39% or 55 million jobs over the next 12 months. For each month that a shock of a given magnitude is prolonged from the base case, heightened macro uncertainty persists for another month.

1 The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Board or the Federal Reserve System. We are grateful to Emanuel Moench and seminar participants at the Federal Reserve Board, the Federal Reserve Bank of New York, and The ECB Research Conference, September 2020. for helpful comments. Ng acknowledges support from the National Science Foundation under grant SES-1558623.

2 Professor of Economics , New York University.
3 Economist, Federal Reserve Board.
4 Edwin W. Rickert Professor of Economics, Columbia University.

Covid Economics 9, 24 April 2020: 1-24

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. The econometric identification of these shocks can be quite challenging, however, making it difficult to empirically distinguish cause and effect. By contrast, economic fluctuations driven by natural disasters such as earthquakes and tsunamis are often more cleanly exogenous shocks, facilitating the identification of their dynamic causal effects. But these types of "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.

Figure 1 shows the responses of US flight departures, initial claims for unemployment insurance, and macro uncertainty from Jurado, Ludvigson, and Ng (2015) to Hurricane Katrina in 2005:08. The number of flight departures dropped immediately in response to Katrina's landfall and both initial claims and macro uncertainty rose sharply. But the impact on initial claims was highly transitory, while the peak effects on macro uncertainty and flight departures slowly build.

A global pandemic is likewise a natural disaster that functions as an exogenous shock with potentially grave economic consequences. But unlike a conventional natural disaster shock, the novel coronavirus 2019 (COVID-19) shock is a multi-period event that simultaneously disrupts supply, demand, and productivity channels, that is almost perfectly synchronized within and across countries, and that has cataclysmic health, social, and economic implications not just for the foreseeable few weeks after the crisis, but for a long time period.

The ability to design policies to mitigate the economic impact of COVID-19 requires reference estimates of the effects of the shock. This paper provides some preliminary estimates of these effects. Our analysis has two ingredients. The first is the construction of a costly disaster (CD) time series from historical data to measure the pecuniary costs of previous disasters. The second is an analysis of the dynamic impact of a costly disaster shock on different measures of economic activity and on a measure of uncertainty. We then design different profiles for the shock to engineer the dynamic effects of a natural disaster interpreted as a large, multi-period, constraint on the ability to produce and consume, as would be characteristic of a pandemic.

We find that the macroeconomic impact of COVID-19 is larger than any catastrophic event that has occurred in the past four decades. Although the CD series has short memory, the effects on economic activity are more persistent. Even under a fairly favorable scenario where

Figure 1: Responses to Hurricane Katrina

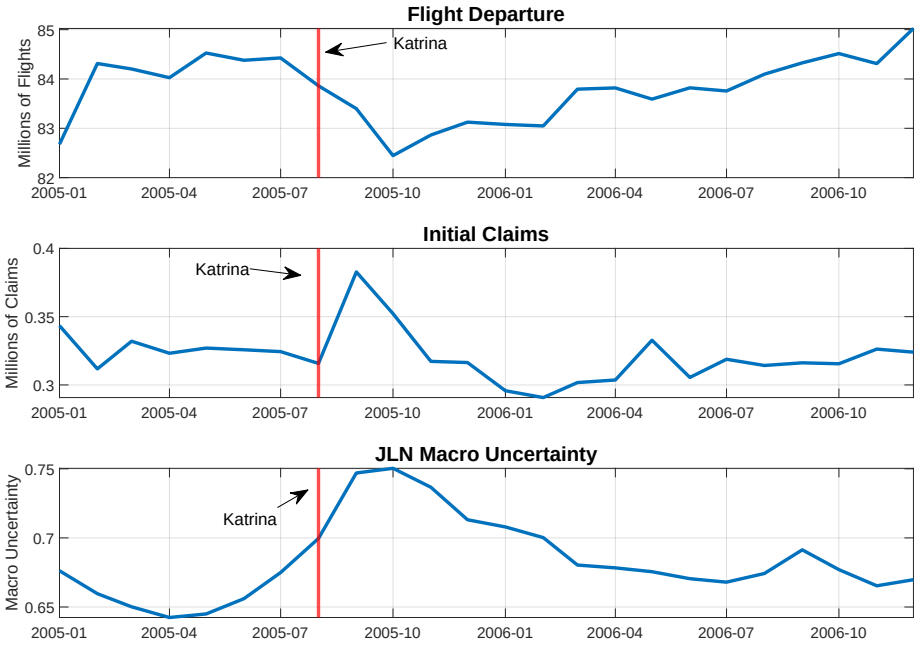

Note: The figure plots number of flight departures in the US, initial claims and JLN macro uncertainty during 2005:01 to 2006:12. The vertical red line indicates the month of Katrina landfall in 2005:08.

the shock persists for only five months and where the initial magnitude is calibrated by the cost of Coronavirus relief packages passed in March of 2020, the estimates suggest that there will be a peak loss in industrial production of 12% and in service sector employment of 5.28% respectively. This translates into a cumulative ten-month loss in industrial production of 20.5%, an employment loss of nearly 39% (or 55 million jobs), and five months of elevated macroeconomic uncertainty. Estimates that allow for nonlinear effects give more pessimistic predictions entailing steeper and longer losses. To the best of our knowledge, this paper is one of the very few time-series analyses of natural disasters on aggregate economic activity, and the first such study of COVID-19.

2 Data and Methodology

Our analysis uses monthly data on disasters affecting the U.S. over the last forty years taken from two sources. The first is from the National Oceanic and Atmospheric Administration (NOAA), which identifies 258 costly natural events ranging from wildfires, hurricanes, flooding, to earthquakes, droughts, tornadoes, freezes, and winter storms spanning the period 1980:1-2020:04 for $T = 482$ 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 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 lives 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.

It is more challenging to measure the dollar cost of the COVID-19 shock. Ideally, one would measure the total dollar cost of mandatory stay-at-home orders across the United States. Although firm-level insurance against losses attributable to business closures exists, these policies cover only short-term closures due to idiosyncratic incidents such as fire and flooding–they do

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

Covid Economics 9, 24 April 2020: 1-24

Figure 2: Time Series of Disaster Series: 1980:1-2020:04

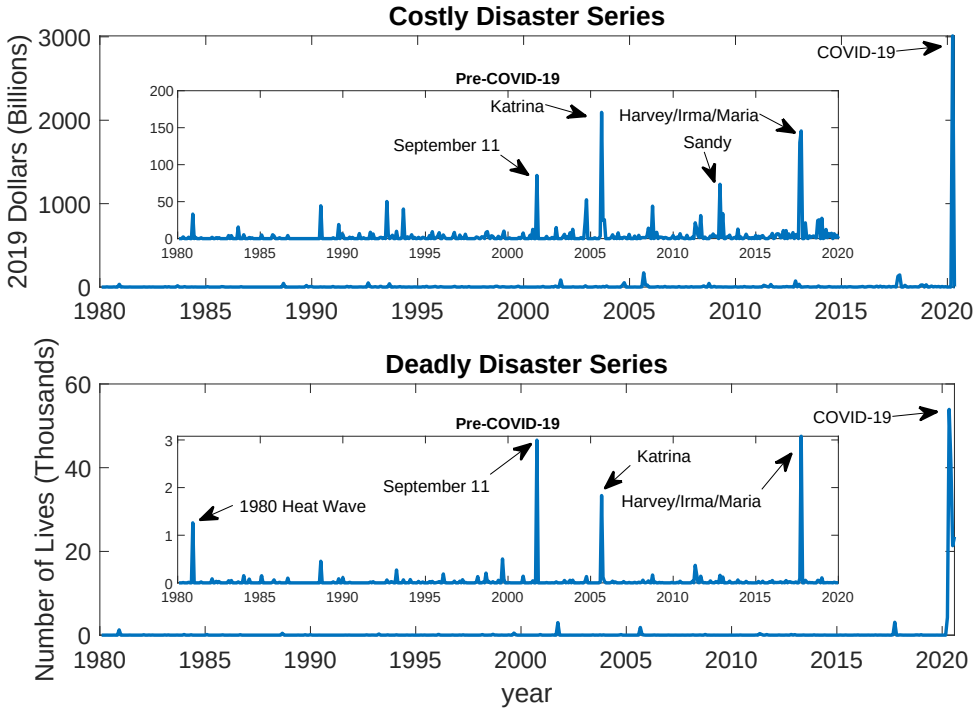

Note: The figure plots the Costly and Deadly Disaster series. The sample spans 1980:01 to 2020:04.

not cover losses due to pandemics or legally mandatory shut-downs. We therefore instead use the dollar value of the Coronavirus relief packages that were passed by U.S. Congress and signed into law in March 2020 as a crude estimate of the dollar cost of COVID-19.[2] These packages total 3.01 trillion dollars, authorized in four separate measures.[3] This dollar cost dwarfs any of those associated with previous U.S. natural disasters in our dataset. The nonlinearities implied by outlier shocks are partially addressed in the penultimate section of the paper.

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

[2]We also use a more conservative estimate from the American Property Casualty Insurance Association, as discussed below.

[3]Source: https://www.npr.org/2020/05/15/854774681/congress-has-approved-3-trillion-for-coronavirus-relief-so-far-heres-a-breakdown

The packages include 26 billion for testing, 217 billion for state and local governments, 312 billion for public health, 513 billion for all businesses in the form of tax breaks meant to help all businesses, 532 billion for large corporations in the form of loans, 784 billion for individuals, and 871 billions for small businesses in the form of forgivable loans under certain conditions.

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 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 deadly disaster series that tallies the number lives lost for all 259 events considered in the analysis.[4]

Figure 2 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. There are four prior 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. Both disaster series are evidently heavy-tailed, and we will return to this point below.[5] Because the size of the increases in both our calibrated COVID-19 CD shock and COVID-19 deaths dwarfs the previous disasters, the latter are shown in inset on the figure, where the COVID-19 values appear on the far right.

Because both the CD and the DD values for COVID-19 represent extreme observations even relative to any of the previous U.S. natural disasters in our dataset, we estimate the parameters of all of our empirical specifications on pre-COVID-19 data, from 1980:01 to 2020:02. Where we use the COVID-19 observations on CD and DD, instead, is in engineering shock profiles that may be deemed more appropriate for the COVID-19 disaster. Specifically, we use these data to help calibrate the size and duration of the COVID-19 event as they pertain to shocks to CD. We use the deaths data in the penultimate section of the paper in a nonlinear specification designed to capture the effects of extreme disasters that are also deadly.

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

[4]Source: `https://en.wikipedia.org/wiki/Casualties_of_the_September_11_attacks`

[5]We also considered CD scaled by real GDP (in 2019 dollars). The VAR analysis using scaled series delivers quantitatively similar results. It's worth noting that 1992 Hurricane Andrew and 1988 Drought costed more, scaled by 1992 and 1988 real GDP, than 2012 Hurricane Sandy.

Covid Economics 9, 24 April 2020: 1-24

Covid Economics 9, 24 April 2020: 1-24

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 thinking about the size of the COVID-19 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

$$X_t = \begin{bmatrix} \text{CD}_t \\ Y_t \\ U_t \end{bmatrix} = \begin{bmatrix} \text{Costly Disaster} \\ \log \text{(Real Activity)} \\ \text{Uncertainty} \end{bmatrix}, \tag{1}$$

where CD is our costly disaster series just described, U_t is a measure of uncertainty, and Y is one of four measures of real activity that will be discussed below. The long-run trends of all three variables are removed using the methodology in Müller and Watson (2017) before the VAR estimation.[6]

We estimate the VAR using monthly data from 1980:01 to 2020:02, and thus exclude the extremely high value of CD during the COVID-19. 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 $\text{diag}(B) = 1$. For identification, B is assumed to be lower triangular. That is, 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 their very nature, exogenous. The resulting structural VAR (SVAR) has a structural moving average representation taking the form

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

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).

[6] Our results remain robust if we instead include a long-run trend in the VAR estimation.

The goal of the exercise is to trace out the effect of COVID-19 on itself, on economic activity Y over time, and on 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 monthly 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 and is available from 2000 onwards. 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 studied 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 Responses to a One σ One Period Shock

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. In what follows, the blue line depicts the median response and the dotted lines refer to 68 percent confidence bands. 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.

The top left panel of Figure 3 is based on the measure of macro uncertainty in Jurado, Ludvigson, and Ng (2015) (JLN). It shows that the impact of a one-standard deviation positive 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 3 shows that JLN uncertainty rises following a positive CD shock, and that the heightened uncertainty persists for three months. The bottom panel replaces the JLN measure of macro uncertainty by the measure of financial uncertainty developed in Ludvigson, Ma, and Ng (2019) (LMN). 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

Covid Economics 9, 24 April 2020: 1-24

Figure 3: Dynamic Response of CD and U to a σ Shock

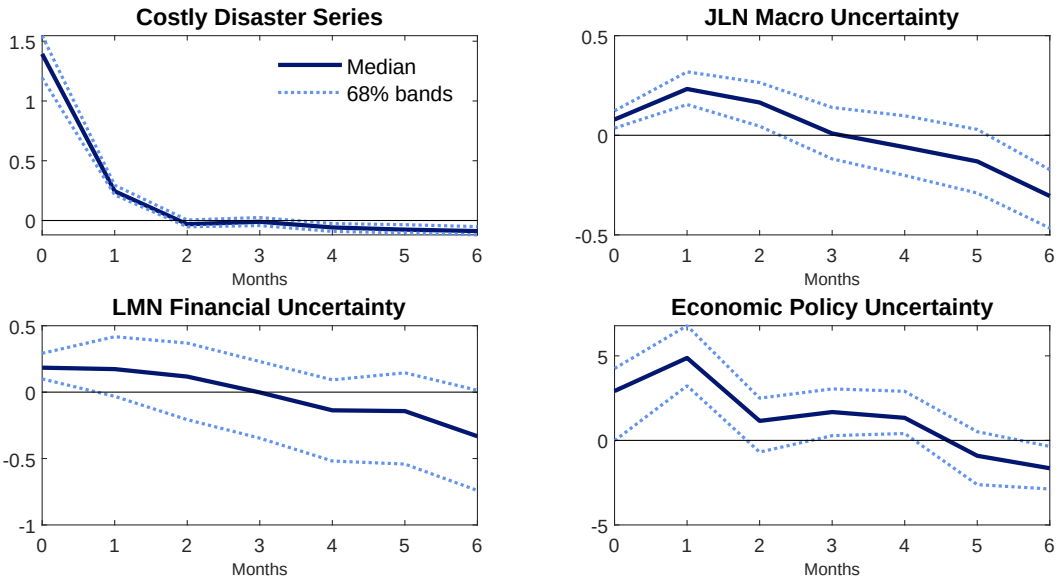

Note: The figure plots the dynamic responses to a positive one-standard deviation CD shock. The posterior distributions of all VAR parameters are estimated using Bayesian estimation with flat priors and the 68% confidence bands are reported in dotted lines. The sample spans 1980:01 to 2020:02.

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.

Next, we consider the effect of a one standard deviation CD shock on four measures of Y, all using JLN macro uncertainty in the VAR. The left top panel of Figure 4 shows that monthly IP immediately drops by 0.05% on impact but becomes statistically insignificant after two months. As seen from Figure 3, two months 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 estimated effect of CD on IP may be attributable to the fact that natural disasters have not had much direct impact on regions of the U.S. where the bulk of industrial production takes place. The top right panel shows that a CD shock triggers a statistically significant rise in unemployment claims IC for about two months with a statistically significant decline in claims (i.e. a rebound in employment) thereafter.

Covid Economics 9, 24 April 2020: 1-24

The bottom left panel of Figure 4 shows that a CD shock leads to an immediate and statistically significant drop in the number of employed workers in the service industry, ESI. Unlike results using IP and IC as Y, the 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 the most 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 4: Dynamic Response of Real Activities to a σ Shock

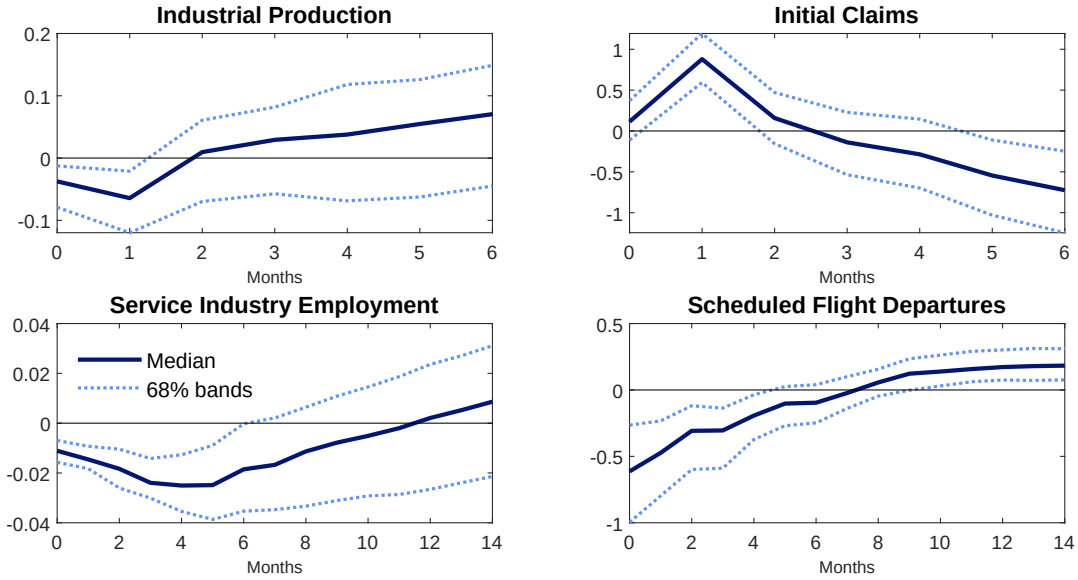

Note: The figure plots the dynamic responses to a positive one-standard deviation CD shock. The posterior distributions of all VAR parameters are estimated using Bayesian estimation with flat priors and the 68% confidence bands are reported in dotted lines. The sample spans 1980:01 to 2020:02.

Taken together, this baseline estimation using pre-COVID-19 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.

COVID-19 differs from historical disasters in several dimensions. The initial impact of

Covid Economics 9, 24 April 2020: 1-24

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. The historical disasters were also short-lived, and with the exception of a drought that lasted over a year, they have an average duration of only one month. 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 COVID-19. COVID-19 is a global pandemic and the effects traverse across states and countries. In April 2020, 91% of the world population live in countries with restricted travel.[7] By contrast, the most disastrous events in our CD disaster series in terms of loss of life were Katrina and 9/11, but the number of deaths due to COVID-19 far exceeds the deaths due to Katrina and 9/11 combined. Moreover, five months into the pandemic, the crisis had yet to reach its peak, and there is a good deal of uncertainty as to whether normalcy will return by the end of 2020. 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. Past disasters created destruction in physical capital, while COVID-19 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. Taken together, these considerations suggest that the dynamic effects of CD need to be altered to reflect shock profiles commensurate with our understanding of COVID-19, which means shocks that last longer than one period, and much larger than one standard deviation.

Covid Economics 9, 24 April 2020: 1-24

4 Effects of Prolonged Shocks

This section addresses the problem that COVID-19 is not a one-shot shock. Ideally, the duration of the shock is the life of the virus which is not only unobservable, but potentially endogenous. To the extent that a COVID-19 shock can be thought of as an economic shock that 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.

Let $\mathbb{X}^t$ collect all information in X at time t and at all lags. From the moving-average representation of the SVAR given in (2), we see that if there are two consecutive shocks of one

[7] 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/

standard deviation, the dynamic response of X_{t+h} is

$$\mathbb{E}\Big[X_{t+h}\big|e_{1t}=\sigma,e_{1t-1}=\sigma;\mathbb{X}^t\Big]-\mathbb{E}\Big[X_{t+h}\big|e_{1t}=0,e_{1t-1}=0;\mathbb{X}^t\Big]=\Psi_h d_1+\Psi_{h+1}d_1,$$

where d_1 is the first column of the matrix $B\Sigma^{1/2}$. 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 $.5\Psi_h d_1+2\Psi_{h-1}d_1$. 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 now interpreted as a constraint on economic activity and so only need the first column of Ψ_h for $h=1,...H$.

Figure 5: Dynamic Response of CD and U to Multi-period one σ Shock

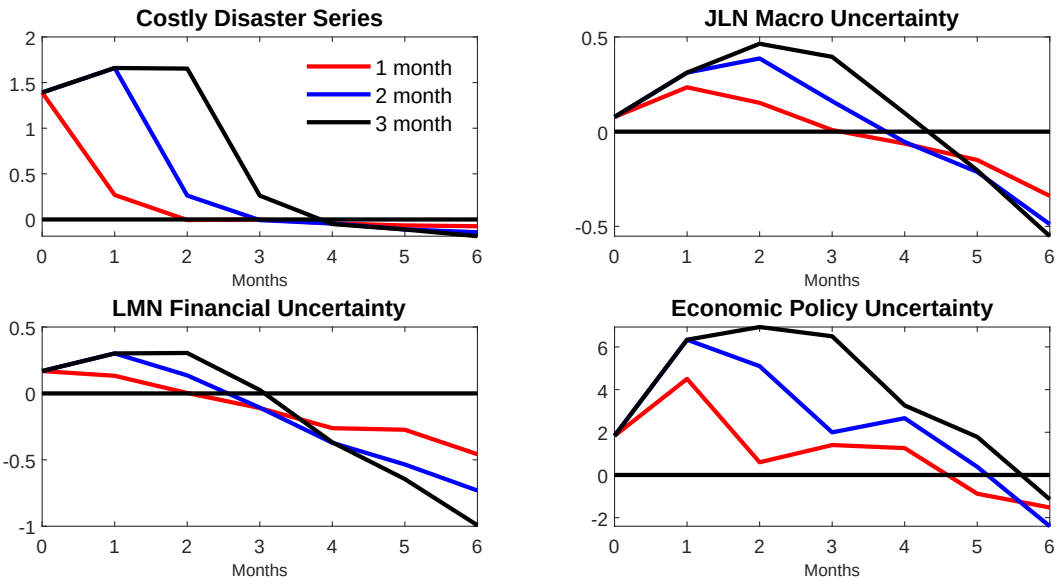

Note: The figure plots the dynamic responses to multi-period consecutive positive one-standard deviation CD shocks. The sample spans 1980:01 to 2020:02.

Figure 5 reports the response of CD and U, similar to Figure 3, except that there are now three consecutive one-standard deviation shocks. To avoid clutter, the confidence bands are not plotted as their significance can be inferred from Figure 3. The red line is the same as the one period shock reported in Figure 3 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

Covid Economics 9, 24 April 2020: 1-24

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 6: Dynamic Response of Real Activities to Multiperiod one σ Shock

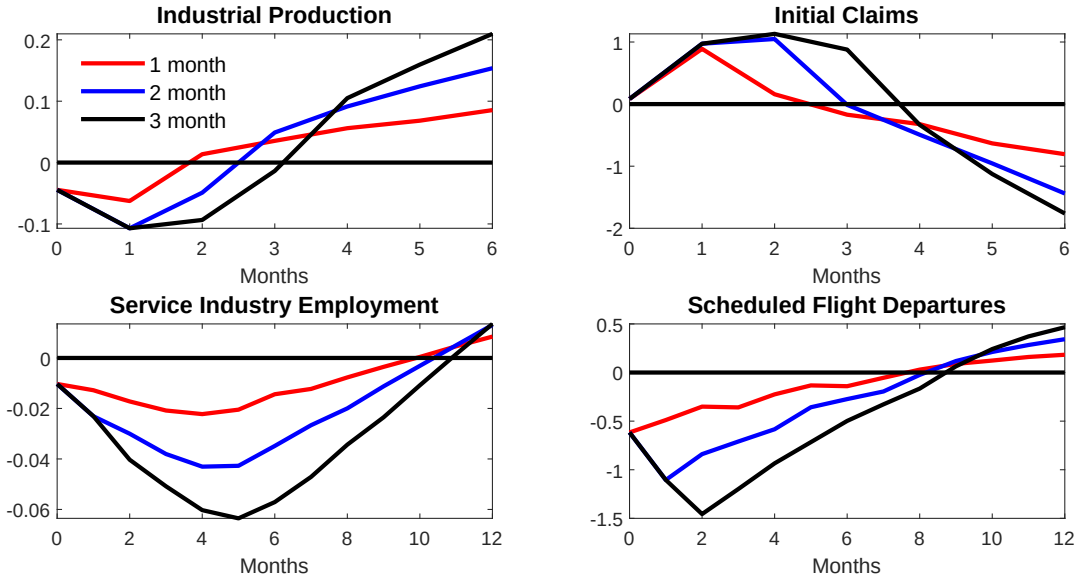

Note: The figure plots the dynamic responses to multi-period consecutive positive one-standard deviation CD shocks. The sample spans 1980:01 to 2020:02

Figure 6 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 4 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, presumably because initial claims can only be filed once, and the losses are front loaded, 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 quarters. Similarly, instead of an immediate recovery, multi-period shocks reduce scheduled flight departures by two more months before a slow recovery begins.

Covid Economics 9, 24 April 2020: 1-24

5 Results for Multiperiod Multi-σ Shocks

We now engineer the shock profile to reflect our understanding of the COVID-19 disaster. For this, we consider dynamic responses to multi-period large shocks. To get a sense of the magnitude of COVID-19, note that 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, as of August 2020, COVID-19 has already resulted in 159,000 deaths in the US, which has more fatalities than the Korean War (92,134) and has exceeded the number of deaths due to the Vietnam War (153,303).[8]

Thus, for the magnitude of the shock, our baseline profile of COVID-19 is based on the fact that Hurricane Katrina was a 11σ shock and the magnitude of the CD series for 2020:03 based on the March relief package is 17.5 times larger than the cost of Katrina. We therefore take 192σ (11 times 17.5) as the benchmark magnitude of COVID-19. We also consider a more conservative profile based on an estimated cost of business closure provided to us by American Property Casualty Insurance Association, which results in a one-trillion dollar cost during the peak of COVID-19.[9] This translates into a cost of COVID-19 that is 5.9 times larger than that of Katrina, and therefore we take 65σ as the magnitude of COVID-19 for this case.

As for the duration, we calibrate the shock profiles by using the fraction of states that are listed as "not reopening" weighted by their GDP contributions as of 2019:Q4. Table 1 reports the fraction of GDP (2019:Q4) earned in states that are categorized as reopened/reopening versus those that are not. As of July 31, 52.4% of 2019:Q4 GDP was earned in states that are not reopening. Some of these states are pausing or reversing previous reopening plans because of the surge of new COVID-19 positive cases in late June and early July. Therefore, we first calibrate the size of shock in July to be 52.4% of the size of shock in March (192σ) or 100σ shock. If we assume that the shock was zero in the interim months, then the five-month shock profile from March 2020 to July 2020, is a (192,0,0,0,100) standard deviation shock profile. As an alternative profile, we also consider a five-month $(192,0,88,79,100)\sigma$ shock profile. This alternative profile is based on the fraction states that were not reopening weighted by their

[8] Source: `https://en.wikipedia.org/wiki/United_States_military_casualties_of_war`.

[9] These preliminary estimates were calculated by the American Property Casualty Insurance Association (APCIA), in the framework of looking at Business Interruption type of coverages (which do not normally cover pandemics). So they do not directly reflect assumptions about total revenue and/or total operating expenses, which would result in larger numbers. According to APCIA, the main component driving these estimates are payroll and benefits. These estimates consider potential insurance type costs per month if all businesses in the US are closed and all receive compensation for the relevant costs – lost profit, payroll/benefits, and partial/additional expenses.

Covid Economics 9, 24 April 2020: 1-24

GDP contributions from May to July.

Table 1. State-level Reopening Summary Statistics

Snapshot	Fraction of 2019 Q4 GDP Earned in States	
	Reopening	Not Reopening
As of April 30	12.30%	87.70%
As of May 31	53.90%	46.10%
As of June 30	59.03%	40.97%
As of July 31	47.56%	52.44%

Note: This table report the fraction of 2019 real GDP earned in states that are "reopening" and "not reopening". The source of the data is from the New York Times (NYT, link: https://www.nytimes.com/interactive/2020/us/states-reopen-map-coronavirus.html). "Reopening" states include all those where every major sector has reopened or is in the process of reopening, albeit possibly under restrictions such as social distancing; these are categorized as "reopened" or "reopening" by the NYT. "Not reopening" includes all states that are assigned to one of the following NYT categories: "regional opening," "shutdown," "pausing," and "reversing." The state-level 2019 GDP estimates are obtained from Bureau of Economic Analysis.

A large shock shifts up the dynamic responses relative to a one-standard-deviation shock presented in Figure 4, while a multi-period shock shifts the dynamic responses to the right as shown in Figure 6. It is of interest to ask how the dynamic responses would change if the disruption is spread over more periods. Figure 7 plots the dynamic responses of a $(192,0,0,0,100)\sigma$ shock profile in dark blue. Plotted next in dotted blue is a five-month $(192,0,88,79,100)\sigma$ alternative shock profile.

The picture that emerges from Figure 7 is that cumulative losses are primarily determined by the total magnitude of the shock rather than the magnitude in any one period. But the longer the duration holding the shock size each period fixed, the larger are the losses and the slower the recovery. The losses for ESI and SFD are particularly steep and persistent.

We report in Table 2 the maximum response in a 12-month period, where the location of the maximum can be inferred from Figure 6. Table 2 also reports the cumulative loss over the months with negative responses.[10] These maximum and cumulative losses are reported for four different shock profiles. The first two shock profiles, $(192,0,0,0,100)\sigma$ and $(192,0,88,79,100)\sigma$, have initial magnitudes that are calibrated based on the Coronavirus relief package passed in March 2020. The next two shock profiles, $(65,0,0,0,34)\sigma$ and $(65,0,30,27,34)\sigma$, have initial

[10]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.

Covid Economics 9, 24 April 2020: 1-24

Figure 7: Dynamic Response to Two Shock Profiles

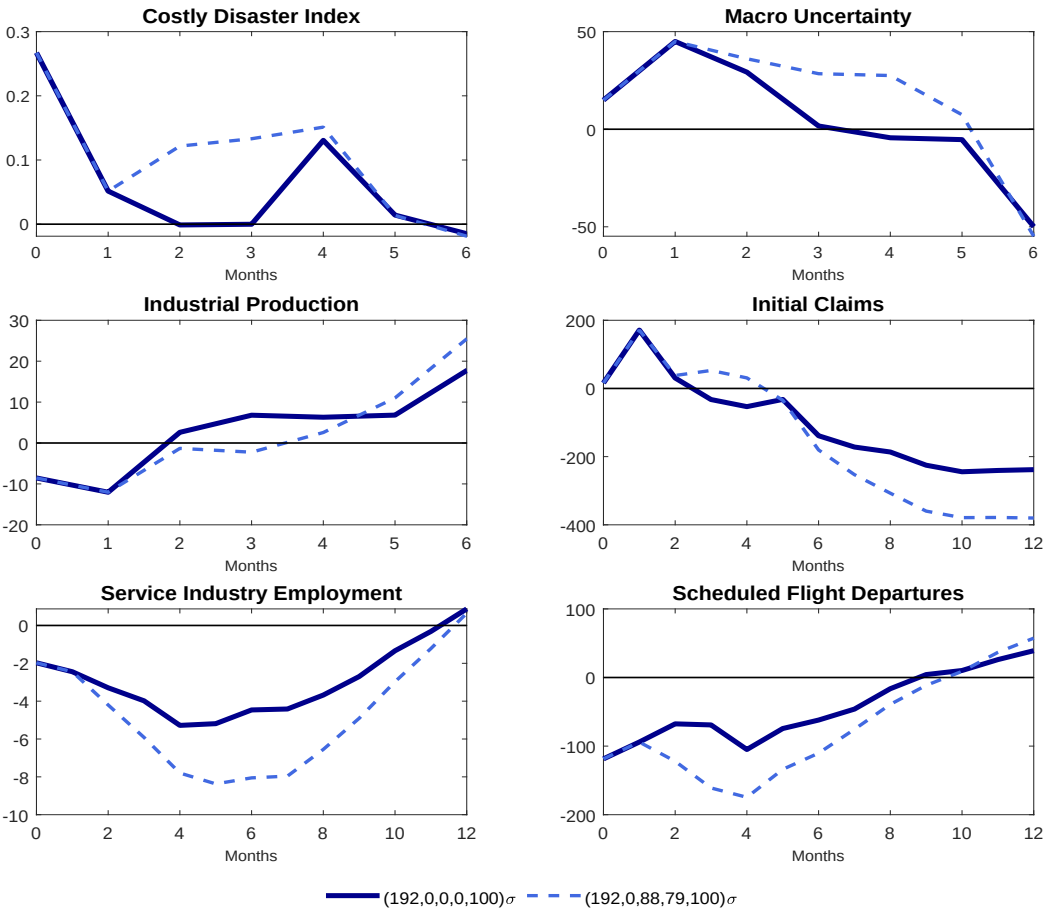

Note: The figure plots the dynamic responses to different disaster shock profiles. The sample spans 1980:01 to 2020:02

magnitudes that are calibrated based on the APCIA insurance cost. Then the size of the subsequent shocks are calibrated based on the fraction of states that are not reopening, as defined above.

Table 2: Max Negative and Cumulative Effects of COVID-19 Shock Profiles

Shock Profiles	Industrial Prod.	Initial Claims	Service Emp.	Flights
	Calibration based on Relief Pacakage			
(192,0,0,0,100)σ *max*	−12.04%	171.12%	−5.28%	−118.63%
Cumulative Losses	−20.58%	217.78%	−39.07%	−653.16%
(192,0,88,79,100)σ *max*	−12.04%	171.12%	−8.37%	−174.61%
Cumulative Losses	−24.15%	308.82%	−62.37%	−1040.70%
	Calibration based on APCIA insurance cost			
(65,0,0,0,34)σ *max*	−4.07%	57.93%	−1.79%	−40.16%
Cumulative Losses	−6.97%	73.73%	−13.23%	−221.12%
(65,0,30,27,34)σ *max*	−4.07%	57.93%	−2.83%	−59.11%
Cumulative Losses	−8.17%	104.55%	−21.11%	−352.33%

Note: Rows labeled *max* show the maximum negative dynamic response from VAR $X_t = (CD_t, Y_t, U_{Mt})'$ for different shock profiles. "Cumulative loss" is the sum of all negative (positive for IC) responses within 12 months. The sample spans 1980:01 to 2020:02.

Table 2 shows that our first shock profile (192,0,0,0,100)σ will lead to a maximum drop in industrial production of 12.04% occurring after one month, a 5.28% maximum loss in service sector employment (over 7 million jobs) occurring after four months, and a 118.63% reduction in scheduled flights after two months. The reduction in service sector employment implications are not trivial because over 75% of workers (or over 140 million) are employed in the service sector. The implied cumulative reduction of 39%, or loss of nearly 55 million service sector jobs before the onset of recovery is staggering. These numbers reach a cumulative reduction of 62%, or a loss of 88 million service jobs, for the (192,0,88,79,100)σ profile.

6 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 severe disasters, where the degree of severity is measured by the number of deaths. Let S_t be an observable variable. We estimate a series of single equation regressions,

Covid Economics 9, 24 April 2020: 1-24

Covid Economics 9, 24 April 2020: 1-24

one for each h, to obtain the dynamic response at lag $h \geq 1$:[11]

$$X_{it+h} = \alpha_0 + \beta^h(L)'X_{t-1}(L) + S_{t-1}\left(\delta_0^h + \delta_1^{h\prime}X_{t-1}\right) + \varepsilon_{it+h}, \tag{3}$$

where $S_t = \frac{\exp(\gamma \mathrm{DD}_t)}{1+\exp(\gamma \mathrm{DD}_t)}$ is a logistic function in the number of deaths in our deadly disaster series, DD, normalized to be mean zero and variance one. By construction, S_t is bounded between zero and one and downweights extreme observations to a degree that depends on the parameter γ. S_t is close to one in the most deadly disasters (when DD is large) and close to zero in the least deadly disasters. To form a benchmark for comparison, we consider using the values S_t would take during three different disasters: using deaths during the month of hurricane Katrina in August of 2005, the month of September 11 of 2001, and using the average monthly deaths from March to July of 2020 due to COVID-19. The latter corresponds to DD $= 40,000$. If γ is sufficiently large, S_t is approximately the same value (unity) for all these three deadly events, so a high-γ calibration would be unable to distinguish them. (See Figure A1 for a plot of S_t under different values of γ). Therefore we choose $\gamma = 0.25$ for the baseline estimation, resulting in $S_t = 0.85$, 0.95, and 1 for Katrina, September 11, and COVID-19, respectively.

As for the VAR, we estimate equation (3) using pre-COVID-19 data. We orthogonalize the residuals in (3) using a Cholesky decomposition, i.e., let $\Omega = \mathbb{E}\left(\varepsilon_t\varepsilon_t'\right) = PP'$, where P is the lower triangular Cholesky factor when the variables ordered as in (1). We then generate the dynamic impulse responses $\widehat{\beta}^h P_1 + S_{DD_j}\widehat{\delta}_1^h P_1$, where P_1 is the first column of P, and where S_{DD_j} is a value for S_t in one of the three disasters just discussed. Figure 8 plots the dynamic responses to a one-period, one standard deviation shock. The red line reports the responses using the COVID-19 calibrated value $S_{DD_j} = 1$ while the black line reports the responses using $S_{DD_j} = 0.95$ from Katrina. For comparison, the blue line shows the dynamic responses from the linear VAR reported in Figure 4. For IP and IC, the responses of the nonlinear models are similar to the linear model (in blue). Both responses peak almost immediately after the shock but IP recovered much more slowly for nonlinear models. For SFD and ESI, the negative responses in the red line are larger and more persistent. The sensitivity of to these results to the choice of γ are shown in Figure (A2). Naturally, real variables have larger negative responses to CD shock when γ is smaller, since smaller values give greater weight to extreme values of DD_t.

Table 3 summarizes the maximum and cumulative negative responses for the shock profiles studied above, but this time using the nonlinear model under the COVID-19 calibrated value of $S_{DD_j} = 1$.[12] Compared to estimates from linear VAR reported in Table 2, the maximum impact of the disaster shock is much larger for all measures of activity, particularly so when

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

[12] The estimated CD shock for Katrina from the regression of CD series on the RHS variables in equation (3)

Figure 8: Dynamic Response of Real Activities to a σ Shock: Non-linear Model

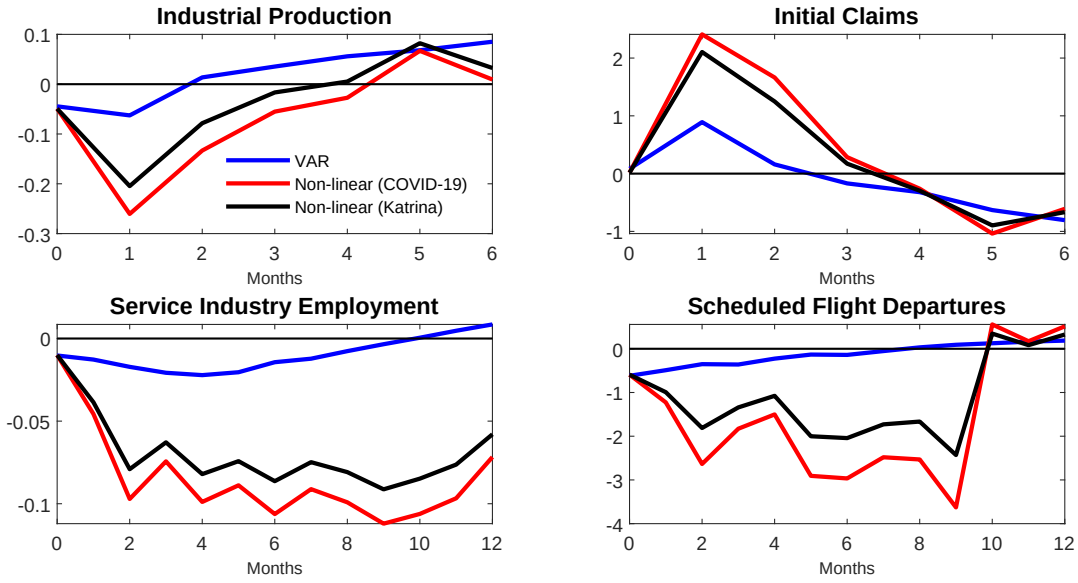

Note: The figure plots the dynamic responses to a positive one-standard deviation CD shock from the nonlinear model. The red lines show the dynamic responses using the COVID-19 calibrated value for S_t, which corresponds to DD = 40,000 average monthly deaths. The black lines use the value for S_t when DD is equal to the number of deaths in the month of Hurricane Katrina. The blue line reports dynamic responses estimated by the linear VAR. The dynamic responses for the nonlinear model are estimated via local projection. The sample spans 1980:01 to 2020:02.

Table 3: Max Neg & Cumulative Effects of COVID-19: Nonlinear Model

Shock Profiles	Industrial Prod.	Initial Claims	Service Emp.	Flights
	Calibration based on Relief Pacakage			
(192,0,0,0,100)σ *max*	−50.02%	462.69%	−31.80%	−742.59%
Cumulative Losses	−142.67%	887.03%	−286.09%	−4479.70%
(192,0,88,79,100)σ *max*	−50.02%	462.69%	−49.84%	−1105.80%
Cumulative Losses	−251.52%	1582.50%	−445.16%	−7380.10%
	Calibration based on APCIA insurance cost			
(65,0,0,0,34)σ *max*	−16.93%	156.64%	−10.76%	−251.40%
Cumulative Losses	−48.30%	300.30%	−96.86%	−1513.20%
(65,0,30,27,34)σ *max*	−16.93%	156.64%	−16.87%	−374.36%
Cumulative Losses	−85.15%	535.75%	−150.71%	−2498.50%

Note: Rows labeled *max* show the maximum negative dynamic response from nonlinear model for $X_t = (CD_t, Y_t, U_{Mt})'$ for different shock profiles, using COVID-19 calibrated value for S_t corresponding to DD = 40,000 average monthly deaths. "Cumulative loss" is the sum of all negative (positive for IC) responses within 12 months from the nonlinear model under the same calibration. The sample spans 1980:01 to 2020:02.

the shock extends more than one period. The first profile of (192,0,0,0,100)σ shocks now leads to a maximum one-month reduction in IP of 50.02% after one month, a 742% reduction of scheduled flights after nine months, and service employment loss of 31.80% after nine months, which is roughly 45 million jobs. The cumulative 12-month losses are also much larger than the VAR estimates imply, generating a cumulative drop of 143% in IP and 286% in service sector employment. For the (192,0,88,79,100)σ profile, the numbers are even larger, with a cumulative decline in service sector employment of 445%. This more pessimistic scenario may have seemed inconceivable at the beginning of 2020, but between March to July 2020 alone there were 55 million new unemployment insurance claims in the US.

7 Conclusion

From monthly data on costly disasters affecting the U.S. over the last forty years, we provide some preliminary estimates of the macroeconomic impact of COVID-19 over the next 12 months from the start of the crisis in February/March of 2020. We find that even in a fairly conservative scenario without nonlinearities, large multiple-period shocks like COVID-19 could plausibly create a 12% monthly drop in IP, a cumulative loss of more than 55 millions jobs in the service industry, sustained reductions in air traffic, and heightened macroeconomic uncertainty several months. The nonlinear model suggests even more pessimistic outcomes.

is 11.4σ above its mean. Therefore, we continue to take 192σ as the benchmark magnitude of COVID-19 for the nonlinear model.

Covid Economics 9, 24 April 2020: 1-24

There are, of course, caveats to the analysis. First, COVID-19 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 smaller losses found for industrial production must be interpreted in this light. Second, we have focused the dynamic responses under one year because the longer horizon results are not well determined. This could be a consequence of the short-memory nature of disaster shocks. Third, 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, the different profiles all suggest steep declines in economic activity, with the longer the duration of the shock, the larger the cumulative losses.

Covid Economics 9, 24 April 2020: 1-24

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.

Müller, U. K., and M. W. Watson (2017): "Low-Frequency Econometrics," in *Advances in Economics and Econometrics: Volume 2: Eleventh World Congress*, vol. 2, p. 53. Cambridge University Press.

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

Figure A1: Logistic Transformation of Deadly Disaster Series

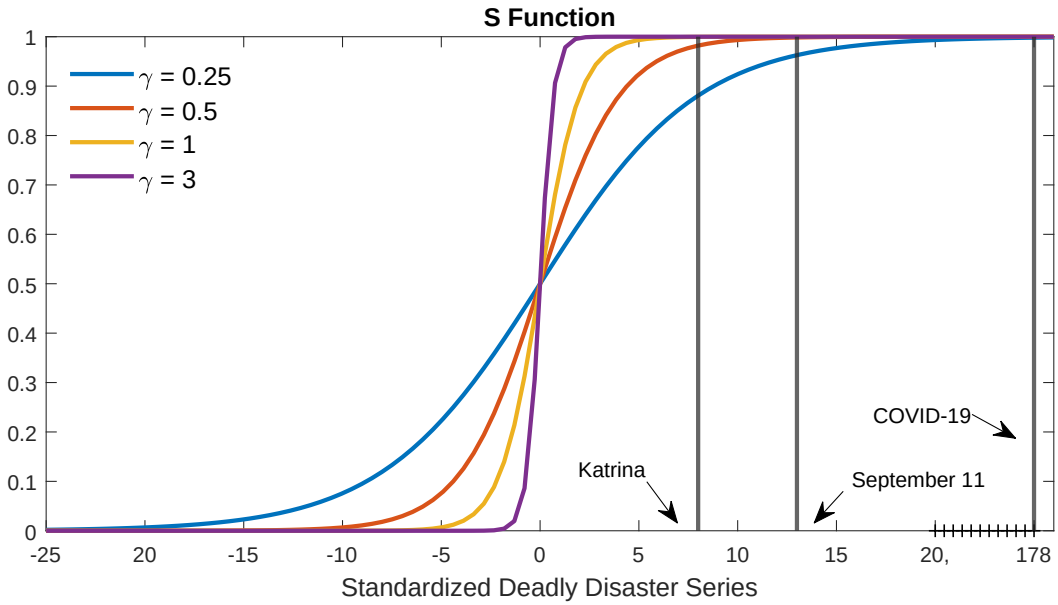

Note: The figure plots the function S over standardized DD series under different values of γ. The vertical lines indicate the values of standardized DD from Katrina, September 11, and COVID-19 calibrated value corresponding to 40,000 average monthly deaths.

Covid Economics 9, 24 April 2020: 1-24

Covid Economics 9, 24 April 2020: 1-24

Figure A2: Dynamic Response from Non-linear Model, Sensitivity Checks

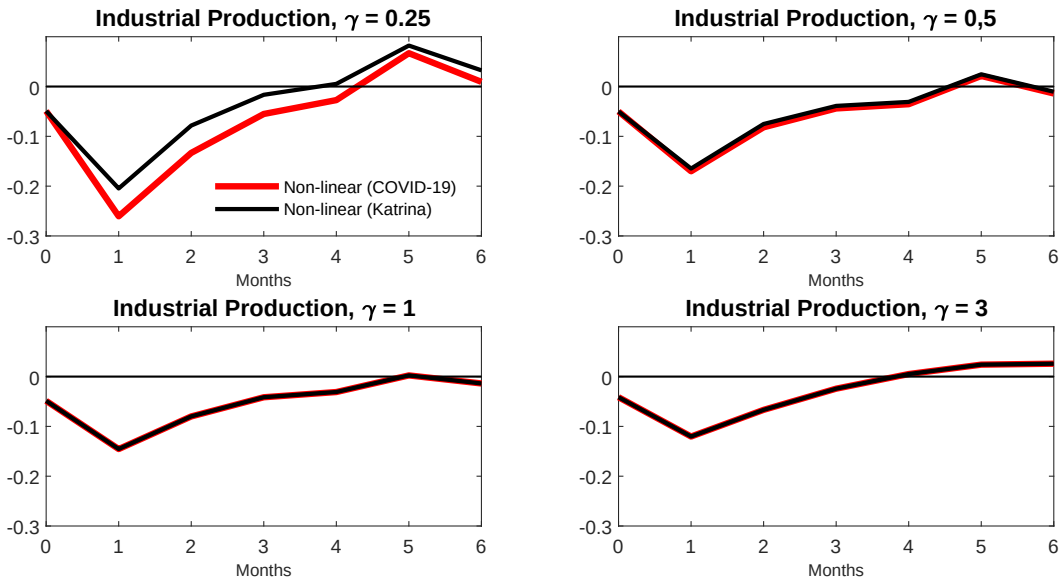

Note: The figure plots the dynamic responses to a positive one-standard deviation CD shock from the non-linear model with different values of γ specified in the subtitle. The red line is the dynamic responses using COVID-19 calibrated value for S corresponding to DD = 40,000 average monthly deaths and the black line uses value of DD series from Katrina. The dynamic responses for the non-linear model are estimated via local projection. The sample spans 1980:01 to 2020:02.

Table A1. List of States that are Reopening

Snapshot	
As of April 30 (10 States)	AK, CO, GA, MN, MS, MT, OK, SC, SD, TN
As of May 31 (38 States)	AL, AK,AZ,AR,CO,CT,DC,FL,GA,HI,ID,IN IA,KS,KY,LA,MD,MA,MN,MS,MO,MT,NV NH,NC,ND,OH,OK,RI,SC,SD,TX,UT,VT,VA WV,WI,WY
As of June 30 (37 States)	AL, AK, CO,CT,DC,FL,GA,HI,IN, IA,KS, KY ME,MD,MA,MN,MS,MO,MT,NE,NH,NJ,NY ND,OH,OK,PA,RI,SC,SD,TN,UT,VT,VA,WV WI,WY
As of July 31 (29 States)	AK,DC,GA,HI,IL,IA,KS,KY,ME,MD,MA,MN MO,MT,NE,NH,NY, ND,OH,OK,PA,RI,SD,TN UT,VT,VA,WV,WI

Note: This table lists the states that were reopened or reopening. The source of the data is from the New York Times (link: https://www.nytimes.com/interactive/2020/us/states-reopen-map-coronavirus.html). "Reopening states" include all states that are either "reopening" or "reopened" as of the date specified in the first column.

Table A2. Glossary

	Reopening
Reopened	States have reopened every major sector, though businesses are almost universally under restrictions, such as allowing fewer customers, requiring workers and customers to wear masks, and enforcing social distancing.
Reopening	States are reopening in stages, allowing some sectors to open ahead of others.
	Not Reopening
Regional Reopening	Governors are allowing regions that meet criteria for slowing the outbreak to open ahead of others. The hardest-hit areas remain under stricter lockdowns
Pausing	Sates have reopened some sectors, but paused or delayed plans to reopen further after seeing a rise in coronavirus cases.
Reversing	Some states have moved to close certain sectors statewide or in certain counties after seeing a surge in cases.
Shutdown	States remain on lockdown, with shutdown orders firmly in place.

Covid Economics 9, 24 April 2020: 1-24

Covid Economics Issue 9, 24 April 2020

Internal and external effects of social distancing in a pandemics[1]

Maryam Farboodi,[2] Gregor Jarosch[3] and Robert Shimer[4]

Date submitted: 18 April 2020; **Date accepted:** 19 April 2020

We use a conventional dynamic economic model to integrate individual optimization, equilibrium interactions, and policy analysis into the canonical epidemiological model. Our tractable framework allows us to represent both equilibrium and optimal allocations as a set of differential equations that can jointly be solved with the epidemiological model in a unified fashion. Quantitatively, the laissez-faire equilibrium accounts for the decline in social activity we measure in US micro-data from SafeGraph. Relative to that, we highlight three key features of the optimal policy: it imposes immediate, discontinuous social distancing; it keeps social distancing in place for a long time or until treatment is found; and it is never extremely restrictive, keeping the effective reproduction number mildly above the share of the population susceptible to the disease.

Covid Economics 9, 24 April 2020: 25-61

1 We thank SafeGraph for making their data freely available to the research community.
2 Jon D. Gruber Career Development Professor and Assistant Professor of Finance, MIT Sloan School of Management.
3 Assistant Professor of Economics and Public Affairs, Princeton University.
4 Alvin H. Baum Professor in Economics and the College, University of Chicago.

Covid Economics 9, 24 April 2020: 25-61

1 Introduction

A key parameter in workhorse models of disease transmission is the rate at which sick people infect susceptible people. A large set of policy measures such as compulsory social distancing aim at reducing this rate. In this paper, we model the rate of transmission as reflecting the choice of rational individuals who weigh the cost of getting infected against the benefits derived from social activity. These benefits capture both social and economic returns from physical human interaction.

To motivate the exercise, we use micro-data from SafeGraph to show that individuals across the United States substantially reduced their exposure to others long before state and local governments imposed the first "shelter-in-place" restrictions in response to the Covid-19 pandemic. We then show how to integrate such optimizing behavior into an otherwise standard epidemiological model. Specifically, we use optimal control theory to derive two ordinary differential equations (ODEs) which capture individual optimality. Together with two standard differential equations from epidemiological models, the resulting system of four differential equations fully summarizes the model and can easily be solved.

The model is consistent with the observation that social activity fell drastically before there were any legally-mandated restrictions on movement. It is thus a natural laboratory for evaluating social distancing policies. To do so, we also show how to characterize the symmetric Pareto optimal allocation. Optimal policy chooses a time path for the amount of social activity, recognizing the health consequences of a high level of activity. Optimal policy can likewise be described by the solution to a simple system of four ODEs. Comparison with the laissez-faire benchmark elucidates the external effects in disease transmission. Moreover, we show how perfect altruism eliminates the gap between equilibrium and optimum.

The internal benefits of social distancing come from the fact that someone is less likely to get sick if they engage in more social distancing. This is reflected in individual behavior. The external benefit comes from the fact that they are less likely to get other people sick, particularly other strangers. While individuals internalize the cost of social distancing, optimal policy also recognizes that individuals may ignore the external benefit of reducing the risk of transmitting illness. Moreover, optimal policy internalizes the effect of an additional sick person on the quality of health care available to inframarginal individuals.

We then turn quantitative. Since the model is very parsimonious, the calibration targets various epidemiological findings such as the initial growth rate of the disease, the duration of infectivity, and the fatality rate.

Our most important findings are the following: First, the laissez-faire equilibrium reduction in social activity due to the internalized cost of infection is strong. It delays the peak

outbreak and leads to a substantial reduction in the number of expected fatalities, relative to a no-response benchmark. However, individuals only reduce activity once the risk of infection becomes non-negligible.

Second, and in contrast, social distancing optimally starts as soon as the disease emerges, discontinuously suppressing social activity. This discrete drop in activity delays the pandemic and hence buys time. Because of the hope for a cure, this strictly reduces the expected number of deaths and yields a welfare gain.

Third, optimal social distancing is persistent. Absent a cure, social activity remains depressed for years. This is the flip side of delay. Because the pool of susceptible individuals drains only slowly, activity needs to remain suppressed for a long time unless a cure is found. Nevertheless, asymptotically social activity returns to normal, with infections stopping only because of herd immunity as the product of the disease's basic reproduction number R_0 and the share of people who are susceptible falls below one.

Fourth, social distancing is never too restrictive. At any point in time, the *effective reproduction number* for a disease is the expected number of people that an infected person infects. In contrast to the basic reproduction number, it accounts for the current level of social activity and the fraction of people who are susceptible. Importantly, optimal policy keeps the effective reproduction number above the fraction of people who are susceptible, although for a long time only mildly so. That is, social activity is such that, if almost everyone were susceptible to the disease, the disease would grow over time. That means that optimal social activity lets infections grow until the susceptible population is sufficiently small that the number of infected people starts to shrink. As the stock of infected individuals falls, the optimal ratio of the effective reproduction number to the fraction of susceptible people grows until it eventually converges to the basic reproduction number.

To understand why social distancing is never too restrictive, first observe that social activity optimally returns to its pre-pandemic level in the long run, even if a cure is never found. To understand why, suppose to the contrary that social distancing is permanently imposed, suppressing social activity below the first-best (disease-free world) level. That means that a small increase in social activity has a first-order impact on welfare. Of course, there is a cost to increasing social activity: it will lead to an increase in infections. However, since the number of infected people must converge to zero in the long run, by waiting long enough to increase social activity, the number of additional infections can be made arbitrarily small while the benefit from a marginal increase in social activity remains positive.

Now consider the role of social distancing in the short run. With a low initial infection rate, pushing the effective reproduction number below the share of susceptible people implies that the total number of individuals who get infected over any time interval will be small.

Covid Economics 9, 24 April 2020: 25-61

That means that the health status of the population—the share of susceptible and infected people—will barely change. It follows that if it is optimal to keep the effective reproduction number below the share of susceptible people initially, it will be optimal to do so at any later date. But we have just argued that this cannot be the case. Thus the effective reproduction number must always be bigger than the share of susceptible people. In our calibrated model, it is mildly so for a sustained period. We note that this argument is predicated on the assumption that the initial infection rate is small, since otherwise a period of strong social distancing may have a big effect on the health status of the population. Indeed, we verify that if the initial infection rate is large enough, an optimal policy may temporarily push the effective reproduction number below the share of susceptible people.

We then revisit our micro-data and contrast it quantitatively with the model under our baseline calibration. The model captures both the drop in social activity prior to any government intervention and the pace of the contraction surprisingly well. We corroborate this with additional aggregated data from Sweden that display similar patterns.

We then consider several robustness exercises with respect to our calibration. An important takeaway is that even large parameter changes matter little for the shape of equilibrium or optimal social activity. Optimal policy is fairly insensitive even to large parameter changes and is well summarized by the points discussed above: Immediate social distancing that ends only slowly but is not overly restrictive. This is reassuring given the large current amount of parameter uncertainty. These robustness exercises also document that our basic framework naturally accommodates a rich set of extensions. We therefore conclude with an additional set of proposed extension that we believe are of first order.

2 Related Work

Our basic approach builds on the susceptible-infected-recovered (SIR) model (Kermack and McKendrick, 1927). There is a rapidly growing body of work that uses this epidemiological model, together with standard economic models, to understand the interplay between disease transmission and economic activity.

The basic epidemiology model is reviewed in Atkeson (2020), who analyzes the optimal lock down policy in an economic environment. Following a similar approach, Eichenbaum, Rebelo and Trabandt (2020) study the two-way interaction of disease dynamics and economic activity in a macroeconomic SIR model. While theirs is a substantially richer environment, we obtain a larger degree of analytical tractability and treat equilibrium and optimum in a unified fashion that simply adds two ODEs to the SIR model. A further key difference is that, in their model, disease transmission and its externalities happens through consumption

Covid Economics 9, 24 April 2020: 25-61

Covid Economics 9, 24 April 2020: 25-61

and the government acts through a tax on consumption. In our setup, disease transmission happens through general activity, both economic and social and the government acts through restricting social activity. This allows us to directly map to newly available data on social activity, for example the SafeGraph data on foot traffic we discuss below.

Other papers that focus on the individual response to a pandemic include Garibaldi, Moen and Pissarides (2020), who use tools from search and matching and Krueger, Uhlig and Xie (2020), who focus on the shift in the sectoral composition of consumption as a force that mitigates the economic fallout of the pandemic.

Alvarez, Argente and Lippi (2020) study a planning problem similar to ours where the planner directly controls the amount of activity and trades of the losses from restrictions against the health benefits. However, they exogenously fix the amount of activity absent policy intervention, while we explicitly model the choice of social activity by individuals. This allows us to contrast the optimal amount of social contacts with the laissez-faire one, taking seriously that even absent mobility restrictions the negative effects of the epidemic are partially internalized.

Some other recent papers focusing on equilibrium and optimal policy include Jones, Philippon and Venkateswaran (2020), Hall, Jones and Klenow (2020), Dewatripont, Goldman, Muraille and Platteau (2020), Piguillem and Shi (2020), Barro, Ursua and Weng (2020), Glover, Heathcote, Krueger and Rios-Rull (2020), Keppo, Kudlyak, Quercioli, Smith and Wilson (2020), and Kaplan, Moll and Violante (2020).

The paper is also related to an older literature on social externalities, including Diamond and Maskin (1979) and Kremer and Morcom (1998). In particular, Diamond and Maskin (1979) introduce the distinction between a *quadratic* and a *linear matching technology*. With quadratic matching additional social activity by others raises the likelihood of social contact and thus disease transmission for all individuals. E.g., with more individuals in parks, restaurants and public transit any given trip to a park/restaurant/subway visit is more likely to lead to disease. Such a quadratic matching function has a search externality that, traditionally, is viewed as positive (Diamond, 1982), but that turns negative in an age of disease. It stands in contrast to a linear search technology where an individual's social contacts merely depend on her own social activity and not on those of others. We believe that such a technology applies to cases where social activity and the associated risk of disease transmission is explicitly sought out. We therefore argue that the quadratic technology is appropriate to model the dynamics of Covid-19, while a linear technology might be the right tool to model an epidemic like HIV.

Greenwood, Kircher, Santos and Tertilt (2019) study the HIV epidemic in Malawi using a computational choice-theoretic equilibrium model of sexual behavior. They model the

individual effort to find a partner in different markets, which are associated with different degrees of risk. And indeed, the risk of infection at any given market depends on the entire distribution of health types visiting the market in line with the arguments just made.

Budish (2020) treats disease containment as an economic constraint and discusses policies that maximize welfare subject to the containment constraint. In turn, our framework treats the system of differential equations governing the evolution of the sickness as the relevant constraint. A policy-maker maximizes welfare subject to that constraint fully taking into account the damage caused by the disease. As a consequence, the policy-maker may well choose policies that let the epidemic spread if the cost of containment are too high. Indeed, we find that this is always optimal, albeit at a slow pace.

Covid Economics 9, 24 April 2020: 25-61

3 Declining Social Activity, Early and Everywhere

In this section, we use newly available micro-data that document substantial behavioral changes across the United States even before any policy measures were taken.

We work with micro-data data from SafeGraph.[1] Among other things, SafeGraph provides highly disaggregated and detailed high-frequency information on individual travel in the United States. The population sample is a panel of opt-in, anonymized smartphone devices, and is well balanced across US demographics and space.

In early April 2020, SafeGraph made two datasets freely available to researchers.[2] Their first "Covid-19 Response Dataset," named "Weekly Patterns," registers GPS-identified visits to Points of Interest (POI) (primarily businesses) with exact known location in the United States at hourly frequency in a balanced panel. The data is currently available covering the period March 1 to April 11, 2020. The dataset is large. On March 1, the dataset recorded approximately 32.1 million individual visits to approximately 3.9 million POI.

The second dataset, named "Social Distancing Metrics," uses information from individual cell devices that can be assigned to a home address (using their night-time location) to measure individual foot traffic and its response to the outbreak. The dataset goes back to January 1, 2020 and currently, runs until April 9 and is likewise large. On March 1, the dataset contains information from over 20 million devices across 220,000 census block groups with at least 5 devices. Among other things, the data measures for each census

[1]Attribution: SafeGraph, a data company that aggregates anonymized location data from numerous applications in order to provide insights about physical places. To enhance privacy, SafeGraph excludes census block group information if fewer than five devices visited an establishment in a month from a given census block group.

[2]For detailed information, see https://docs.safegraph.com/docs/weekly-patterns and https://docs.safegraph.com/docs/social-distancing-metrics.

block group the median number of minutes a device dwells at its home location (variable median_home_dwell_time). In addition, it also measures the number of devices that spend the entire day-of-week at the home location (variable completely_home_device_count).

We construct our measures at the state level. We use the first dataset to count the total daily number of visits, for each state, to POIs. We proceed identically for our other two measures. We subtract the median minutes spent at home from $24 \times 60 = 1440$ and take a daily state-wide average. We similarly construct the state-wide fraction of all devices that leave the house at least once during any day. We express all three variables relative to a baseline week (dividing by the corresponding day during the first week of March).

This gives us, for each state, three different measures of the decline of social activity that naturally map to the model. Figure 1 reports our result for the first variable, visits to POI. We plot, at any given date, the median value (across states) of the decline relative to baseline, along with the max and min and 10th and 90th percentile.

The figure shows a remarkably uniform contraction of social activity across US states beginning in the second week of March, leveling off at some 50 percent relative to baseline towards the end of the month.

The figure also depicts the fraction of the US population subject to official stay-at-home or shelter-in-place orders. The figure shows that social activity began contracting 10 days before the first significant orders were put into place around March 20.

We complement this with the other two social distancing metrics we have available, days spent entirely at home and daily dwell time at home. Their decline is depicted in Figure 2. These two variables display a somewhat smaller decline of 20-30% relative to baseline. The basic pattern remains the same: Social activity starting contracting substantially, rapidly, and long before the first lockdown measures. This also happened across the board in the United States.

This offers a direct and readily available measure of the extent of individual social activity which is both the key endogenous variable in our model, as well as the key driver of the pandemic. Since we model social activity rather than, say, consumption, our model can directly connect with this high-frequency data. We will later confront the quantitative properties of our model with this evidence and argue that it offers a close account of the decline in social activity in the US in March 2020.

4 Model

The basic epidemiological framework is a continuous time SIR model with a possibility of death, i.e. a SIRD model. Individuals susceptible to the disease may become infected

Covid Economics 9, 24 April 2020: 25-61

Covid Economics 9, 24 April 2020: 25-61

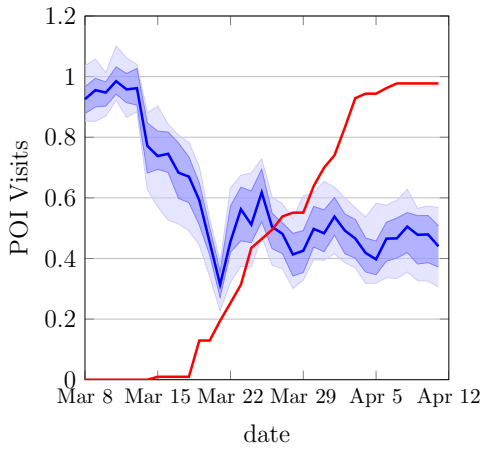

Notes: Visits to Points of Interest in SafeGraph's "Weekly Patterns" Covid-19 Response Dataset. We sum daily visits within each state and express them relative to the same day during baseline week (first week of March 2020). We plot the median (min, max, .. across states) decline relative to baseline at any given day. The solid blue line is the median. The dark share is $10\% - 90\%$ interval, and the light shade is min-max interval. The solid red line is the percent of the population subject to stay-at-home or shelter-in-place orders. The population subject to these orders is based on authors' own calculations using https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html.

Figure 1: Declining Activity, Early and Everywhere

Covid Economics 9, 24 April 2020: 25-61

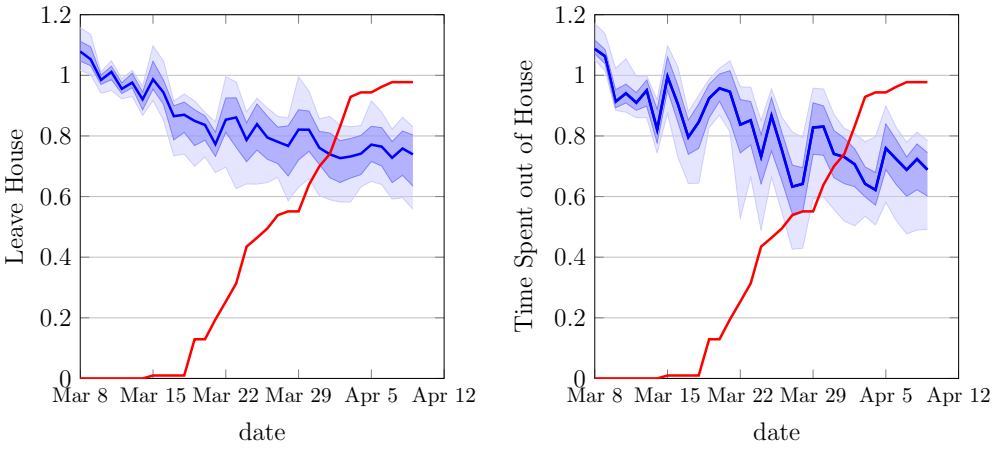

Notes: Social activity based on SafeGraph's "Social Distancing Metrics" Covid-19 Response Dataset. Left Panel: Fraction of devices that leave assigned "home" at least once during any day. Right Panel: Dwell time at home, median device. Measures at the census block group, we take state-wide weighted averages and express them relative to the same day during baseline week (first week of March 2020). We plot the median (min, max, .. across states) decline relative to baseline at any given day. The solid blue line is the median. The dark share is $10\% - 90\%$ interval, and the light shade is min-max interval. The solid red line is the population under lockdown. Population under lockdown is based on authors' own calculations using https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html.

Figure 2: Declining Activity, Early and Everywhere II

Covid Economics 9, 24 April 2020: 25-61

through contact with other infected individuals. The infected stochastically recover or die. Individuals do not know if they are susceptible or infected, but in our baseline model, they can tell when they have recovered from the disease. In our baseline model, individuals are otherwise homogeneous.

At each time $t \geq 0$, a measure 1 of individuals are in one of the four states, susceptible (s), infected (i), recovered (r), or deceased (d). Let $N_j(t)$, $j \in \{s, i, r, d\}$, denote the measure of individuals in each state. Thus $N_s(t)$ also gives the fraction of the population that has not gotten infected. Assume that $N_i(0) > 0$, so there is a seed of infection.

A susceptible individual can get infected by meeting an infected individual. However, individuals are unaware if they are infected. Infected individuals recover at rate $(1-\pi(N_i(t)))\gamma$ and die at rate $\pi(N_i(t))\gamma$, where $\pi(N_i(t)) \in [0, 1]$ is the infected fatality rate, the fraction of infected individuals who eventually die from the disease. We allow this to depend on the number of infected people, reflecting the possibility that the disease overwhelms the hospital system. A recovered individual knows that he is recovered. We assume that recovering from the disease confers lifetime immunity, so a recovered individual no longer transmits the disease and can no longer become sick. We lump the risk and cost of death (i.e. the lost value of life) and the cost of disease together in a single function $\kappa(N_i(t))$, the expected cost that an infected individual pays when he exits the infected state. Again, we allow for this cost to depend on N_i to capture that treatment quality (and hence π) might deteriorate when the health care system becomes overrun.

Individuals discount the future at rate ρ, and a cure for the disease is found at rate δ. For simplicity we assume that a cure, once found, is perfect and immediately wipes out the disease.

We assume that all living individuals choose their level of *social activity* a, and get utility $u(a)$ from social activity a. Assume u is single-peaked and its maximum is attained at a finite value $a^* > 0$. Without loss of generality, normalize $a^* = 1$ and $u(a^*) = 0$. The first normalization keeps the notation the same as the basic SIR model in the absence of a behavioral response to the outbreak. The latter allows us to use u as a measure of the utility loss from social distancing.

Disease transmission depends on the rate at which individuals have social interactions. Let $A_j(t)$ denote the aggregate amount of social activity by all individuals of type $j \in \{s, i, r\}$. We assume throughout that all individuals of a given type choose the same level of social activity, although when we study equilibrium, we consider the deviation of a single individual to a different level of activity. The rate that an individual of type j has *social interaction* with an individual of type j' is $A_j(t)N_j(t)A_{j'}(t)N_{j'}(t)$, the product of the level of social activity of the two types. In particular, the rate that susceptible individuals get sick is

Covid Economics 9, 24 April 2020: 25-61

$\beta A_s(t)N_s(t)A_i(t)N_i(t)$, where $\beta > 0$ captures the ease of transmitting the disease. Formally, this is a quadratic matching technology with random search.[3]

Together the assumptions that preferences u depend on social activity while disease transmission depends on social interactions, is central to our view of social distancing. It captures the idea that individuals value social activity (going to a restaurant, going for a walk, going to the office) and, absent health issues, are indifferent about whether other people are also engaging in social activity.[4] On the other hand, if an individual goes for a walk and doesn't encounter anybody, they cannot get sick. Thus interactions are critical for disease transmission.

The disease transmission function captures the idea that if a particular group chooses little social activity $A_j(t)$, one is unlikely to make social contact with them. It also captures the idea that the amount of social interactions depends not only on a particular group's choice, but also on everyone else choices. This captures, for instance, that even an individual frequently going to a restaurant or to the office has few social interactions if nobody else is there. As a consequence, disease transmission displays a negative externality: increasing social activity increases other people's social interactions, putting them at a higher risk of infection.

Under our modeling assumptions, the laws of motion describing the aggregate state are given by

$$N_s'(t) = -\beta A_s(t)N_s(t)A_i(t)N_i(t) \tag{1}$$

$$N_i'(t) = \beta A_s(t)N_s(t)A_i(t)N_i(t) - \gamma N_i(t) \tag{2}$$

$$N_r'(t) = (1 - \pi(N_i(t)))\gamma N_i(t) \tag{3}$$

$$N_d'(t) = \pi(N_i(t))\gamma N_i(t) \tag{4}$$

with each of the $N_j(0)$ given. If $A_j(t) = 1$ for $j \in \{s, i\}$ and $\pi(\cdot) = 0$, this is the standard SIR model (Kermack and McKendrick, 1927). Allowing $\pi \neq 0$ gives us the SIRD model. We stress that there are not only social interactions between the susceptible and the infected,

[3]An alternative would be a linear technology, in which case type j individuals interact with type j' individuals at rate $\frac{A_j(t)N_j(t)A_{j'}'(t)N_{j'}'(t)}{\sum_{j''} A_j''(t)N_j''(t)}$. This formulation makes sense if type j individuals desire to interact with somebody at rate $A_j(t)$. With a linear technology, changing social activity by other people changes the distribution of social interactions without changing the level. Such an alternative might be the appropriate modeling choice for other epidemics such as HIV (Kremer and Morcom, 1998). However, for Covid-19 it seems unlikely that additional social activity by non-infected individuals would reduce the infection risk of the susceptible, all else equal. We highlight, however, that this assumption is crucial for certain outcomes, for instance the optimal social activity of the recovered, which one would want to boost with a linear matching technology.

[4]For an environment where interactions are critical, see Diamond (1982). One can imagine reasons why the marginal utility of social activity is increasing or decreasing in the social activity of others.

Covid Economics 9, 24 April 2020: 25-61

but also between all other groups. For instance, the recovered can optimally choose $A_r(t)$ without fear of infection; and susceptible individuals meet one another without consequences. Our social matching function allows these interactions to happen, but they do not affect the number of interactions where the disease gets transmitted.

Note that only $A_s(t)$ and $A_i(t)$ affect disease transmission. Also note our assumption that individuals do not know whether they are susceptible or infected. We thus impose the measurability restriction that $A_s(t) = A_i(t)$ and use $A(t)$ to denote this common level of social activity. The expected number of others in contact with an infected individual during the time she is infected is $\beta A(t)^2/\gamma$. In a population in which individuals do not change their behavior in response to the disease $A(t) = a^* = 1$, e.g. when the disease first emerged. Therefore the *basic reproduction number* R_0, defined as the expected number of others infected by an infected individual in a population where everyone is susceptible, is $R_0 = \beta/\gamma$. We also define the *effective reproduction number*, the expected number of others infected by an infected individual, given the current level of social activity $A(t)$ and the current fraction of susceptible people $N_s(t)$. This number is $R_e(t) = \beta A(t)^2 N_s(t)/\gamma$. Reductions in social activity drive the ratio of the effective reproduction number to the fraction of susceptible people, $R_e(t)/N_s(t)$ below the basic reproduction number R_0, which is a primitive of the environment.

5 Laissez-faire Equilibrium

In this section, we consider the problem of an individual who is either susceptible or infected choosing his own rate of social activity, taking the number and social activity of other infected people as given. We then impose the equilibrium restriction that individual outcomes must coincide with the aggregate.

An individual has rational beliefs about his own probabilities of being susceptible, infected, and recovered, which we denote by $n_s(t)$, $n_i(t)$, and $n_r(t)$, respectively. The individual knows when he is recovered but cannot distinguish between the susceptible and infected states. He thus chooses two time paths for social activity, $a(t)$ when he is susceptible or infected and $a_r(t)$ when he is recovered. Finally, the individual discounts future utility at rate ρ and recognizes that the problem ends at rate δ when a cure is found. Putting this together, the individual solves

$$\max_{\{a(t),a_r(t)\}} \int_0^\infty e^{-(\rho+\delta)t}\big((n_s(t)+n_i(t))u(a(t)) + n_r(t)u(a_r(t)) - \gamma n_i(t)\kappa(N_i(t))\big)dt \qquad (5)$$

subject to

$$\begin{aligned} n_s'(t) &= -\beta a(t) n_s(t) A(t) N_i(t) \\ n_i'(t) &= \beta a(t) n_s(t) A(t) N_i(t) - \gamma n_i(t) \\ n_r'(t) &= (1 - \pi(N_i(t)))\gamma n_i(t) \end{aligned}$$

with $n_s(0) = N_s(0)$, $n_i(0) = N_i(0)$, and $n_r(0) = N_r(0)$ given. The laws of motions reflect that the individual takes the time path of everyone else's choice $A(t)$ and the associated aggregate infection rate $N_i(t)$ as given. However, the individual's past choices of $a(t)$ affects his probability of being in each of the states.

To solve the individual's problem, first note that $a_r(t)$ affects the objective but none of the constraints. It is thus optimal to set $a_r(t) = a^* = 1$ and then $u(a_r(t)) = 0$ for all t. Dropping this control variable and the unnecessary third constraint, write the current value Hamiltonian as

$$\begin{aligned} H(n_s(t), n_i(t), a(t), \lambda_s(t), \lambda_i(t)) = (n_s(t) + n_i(t))u(a(t)) - \gamma n_i(t)\kappa(N_i(t)) \\ - \lambda_s(t)\beta a(t) n_s(t) A(t) N_i(t) + \lambda_i(t)\big(\beta a(t) n_s(t) A(t) N_i(t) - \gamma n_i(t)\big), \end{aligned}$$

where $\lambda_s(t)$ and $\lambda_i(t)$ are the co-state variables associated with the two remaining constraints.

There are three necessary first order conditions for optimal control. First, the derivative of the Hamiltonian with respect to the control variable $a(t)$ is zero:

$$(n_s(t) + n_i(t))u'(a(t)) = (\lambda_s(t) - \lambda_i(t))\beta n_s(t) A(t) N_i(t). \tag{6}$$

This static first order condition balances the returns from social activity and the risk of getting infected. Second and third, the derivatives with respect to the state variables $n_s(t)$ and $n_i(t)$ are equal to minus the time derivative of the costate, with a correction for discounting:

$$(\rho + \delta)\lambda_s(t) - \lambda_s'(t) = u(a(t)) + (\lambda_i(t) - \lambda_s(t))\beta a(t) A(t) N_i(t), \tag{7}$$

$$(\rho + \delta)\lambda_i(t) - \lambda_i'(t) = u(a(t)) - \gamma\big(\kappa(N_i(t)) + \lambda_i(t)\big). \tag{8}$$

There are two more necessary conditions for optimality, the transversality conditions

$$\lim_{t\to\infty} e^{-(\rho+\delta)t}\lambda_s(t) n_s(t) = \lim_{t\to\infty} e^{-(\rho+\delta)t}\lambda_i(t) n_i(t) = 0. \tag{9}$$

Equilibrium requires that individual and aggregate behaviors are consistent at every point in time, $n_s(t) = N_s(t)$, $n_i(t) = N_i(t)$, and $a(t) = A(t)$ for all $t > 0$. Imposing those

Covid Economics 9, 24 April 2020: 25-61

restrictions gives us a system of four differential equations and one static equation. Together with initial conditions for $N_s(0)$ and $N_i(0)$ and the transversality conditions (9), these fully summarize the model:

$$\begin{aligned}
N_s'(t) &= -\beta A(t)^2 N_s(t) N_i(t) \\
N_i'(t) &= \beta A(t)^2 N_s(t) N_i(t) - \gamma N_i(t) \\
(\rho+\delta)\lambda_s(t) - \lambda_s'(t) &= u(A(t)) + (\lambda_i(t) - \lambda_s(t))\beta A(t)^2 N_i(t), \\
(\rho+\delta)\lambda_i(t) - \lambda_i'(t) &= u(A(t)) - \gamma\big(\kappa(N_i(t)) + \lambda_i(t)\big), \\
(N_s(t) + N_i(t))u'(A(t)) &= (\lambda_s(t) - \lambda_i(t))\beta A(t) N_s(t) N_i(t).
\end{aligned}$$

The first two differential equations impose $A_s(t) = A_i(t) = A(t)$ on the aggregate relationships (1) and (2). The last three equations correspond to equations (7), (8), and (6) with $a(t) = A(t)$, $n_i(t) = N_i(t)$, and $n_s(t) = N_s(t)$.

We solve this model through a backward shooting algorithm. Fix $\lambda_s(T)$ and $\lambda_i(T)$ at their asymptotic values at some faraway date T, but make $N_i(T)$ slightly positive. Then search for the value of $N_s(T)$ that achieves a desired initial condition for N_s and N_i at a much earlier date. In practice, we can find the equilibrium in a few seconds.

6 Social Optimum

We now solve the problem faced by a benevolent social planner who dictates the time path of social activity $A(t)$ and $A_r(t)$. The planner, like the individual, recognizes that a reduction in contacts lowers utility directly, but she also recognizes the externalities associated with illness. We show that this gives rise to a system of four ODEs which very closely resemble the ODEs shaping equilibrium we just discussed.

6.1 Planner's Problem

The planner solves

$$\max_{\{A(t),A_r(t)\}} \int_0^\infty e^{-(\rho+\delta)t}\big((N_s(t)+N_i(t))u(A(t)) + N_r(t)u(A_r(t)) - \gamma N_i(t)\kappa(N_i(t))\big)dt \quad (10)$$

subject to equations (1)–(3). As in equilibrium, it is optimal to set $A_r(t) = a^* = 1$ and so $u(A_r(t)) = 0$ for all t, since this does not affect the evolution of the state variables. Then

Covid Economics 9, 24 April 2020: 25-61

the Hamiltonian is

$$H(N_s(t), N_i(t), A(t), A_s(t), A_i(t)) = (N_s(t) + N_i(t))u(A(t)) - \gamma N_i(t)\kappa(N_i(t)) \\ - \mu_s(t)\beta A(t)^2 N_s(t) N_i(t) + \mu_i(t)\big(\beta A(t)^2 N_s(t) N_i(t) - \gamma N_i(t)\big).$$

The necessary first order condition with respect to the control A is

$$(N_s(t) + N_i(t))u'(A(t)) = 2(\mu_s(t) - \mu_i(t))\beta A(t) N_i(t) N_s(t), \tag{11}$$

while the necessary costate equations are

$$(\rho + \delta)\mu_s(t) - \mu_s'(t) = u(A(t)) + (\mu_i(t) - \mu_s(t))\beta A(t)^2 N_i(t), \tag{12}$$

$$(\rho + \delta)\mu_i(t) - \mu_i'(t) = u(A(t)) - \gamma\big(\kappa(N_i(t)) + N_i(t)\kappa'(N_i(t)) + \mu_i(t)\big) \\ + (\mu_i(t) - \mu_s(t))\beta A(t)^2 N_s(t). \tag{13}$$

Finally, the planner also has necessary tranversality conditions

$$\lim_{t\to\infty} e^{-(\rho+\delta)t}\mu_s(t)N_s(t) = \lim_{t\to\infty} e^{-(\rho+\delta)t}\mu_i(t)N_i(t) = 0. \tag{14}$$

There are a few key differences between the first order conditions in the equilibrium and optimum problem. First, the planner recognizes that raising $A(t)$ increases meetings at rate proportional to $2A(t)$, while in equilibrium raising $a(t)$ increases meetings at rate proportional to $A(t)$. This creates an extra factor of 2 in equation (11) compared to equation (6). Second, the planner recognizes the health care externality, that the cost of being sick may depend on how many people are sick $N_i(t)$. This is the additional term involving the derivative of the cost function κ in equation (13) compared to equation (8). Third, the planner recognizes that sick people get other people sick, while in equilibrium individuals do not care about this outcome. This shows up as the last term in equation (13).

As usual, the aggregate state still satisfies

$$N_s'(t) = -\beta A(t)^2 N_s(t) N_i(t)$$
$$N_i'(t) = \beta A(t)^2 N_s(t) N_i(t) - \gamma N_i(t)$$

Thus we again have a system of four differential equations and one static equation. We again solve it using a backward shooting algorithm, searching for the terminal value of $N_s(T)$ for given $N_i(T)$.

Before we put numbers into the model, we briefly discuss how a model with altruism can

Covid Economics 9, 24 April 2020: 25-61

encompass both laissez-faire and optimum.

6.2 Perfect and Imperfect Altruism

In the laissez-faire equilibrium, people only care about their own health. In reality, diseases are often transmitted to family and friends, and so it seems plausible that many people would like to reduce the risk of transmitting the disease, not just the chance of getting it. We capture this through a model of imperfect altruism, indexed by an altruism parameter $\alpha \in [0,1]$. At one extreme, $\alpha = 0$, individuals would not socially distance if they knew they were sick. This is the laissez-faire equilibrium. At the other extreme, $\alpha = 1$, individuals fully internalize the cost of making others sick, except for any possible congestion in the health care system. We show that this is the social optimum when κ is constant.

To illustrate this, we modify the individual objective function to assume individual are concerned about making others sick:

$$\max_{\{a(t),a_r(t)\}} \int_0^\infty e^{-(\rho+\delta)t}\big((n_s(t)+n_i(t))u(a(t)) + n_r(t)u(a_r(t)) \\ - \gamma n_i(t)\kappa(N_i(t)) + \alpha\beta a(t)n_i(t)A(t)N_s(t)(\lambda_i(t)-\lambda_s(t))\big)dt.$$

The new piece is the last term. When an individual infects a susceptible person, at rate $\beta a(t)n_i(t)A(t)N_s(t)$, she suffers a utility loss equal to a fraction α of the difference $\lambda_i(t) - \lambda_s(t)$, where again $\lambda_j(t)$ is the costate variable on $n_j(t)$, $j \in \{s,i\}$. In words, this difference represents the private cost of getting sick.

With this modification to the objective function, we can again write down the Hamiltonian and find the optimality and costate equations:

$$\begin{aligned}
\big(n_s(t)+n_i(t)\big)u'(a(t)) &= \beta A(t)\big(\lambda_s(t)-\lambda_i(t)\big)\big(n_s(t)N_i(t)+\alpha n_i(t)N_s(t)\big)\\
(\rho+\delta)\lambda_s(t) - \lambda_s'(t) &= u(a(t)) + \beta a(t)A(t)N_i(t)\big(\lambda_i(t)-\lambda_s(t)\big)\\
(\rho+\delta)\lambda_i(t) - \lambda_i'(t) &= u(a(t)) - \gamma\big(\kappa(N_i(t)) + \lambda_i(t)\big) + \alpha\beta a(t)A(t)N_s(t)\big(\lambda_i(t)-\lambda_s(t)\big).
\end{aligned}$$

As usual for equilibrium, we then impose the conditions $a(t) = A(t)$, $n_s(t) = N_s(t)$, and $n_i(t) = N_i(t)$. If $\alpha = 0$, this returns the equilibrium equations. If $\alpha = 1$ and κ is constant, this returns the equations describing the dynamics of the planner's solution. Intermediate values of α capture a degree of imperfect altruism.

In this formulation, we note one natural limit of altruism: We think it is unlikely that individuals view their own behavior as having an impact on the total number of infected people $N_i(t)$ and hence on the risk of death captured by κ. This is really an aggregate

Covid Economics 9, 24 April 2020: 25-61

outcome, in contrast to the possibility that one person's social activity makes someone else sick, which an individual is more likely to believe that she can control.

7 Quantitative Exercises

7.1 Calibration

We calibrate the model at a daily frequency to US data on the 2020 Covid-19 outbreak. We offer some robustness with regard to the most important choices made in this section in Section 8.

To begin with, we set $\rho = 0.05/365$ to capture a 5% annual discount rate. In addition, we set $\delta = 0.67/365$, which implies an expected duration at which a cure is found of 1.5 years. We highlight that this jointly implies a model of heavy discounting relative to a standard economic model.

Next, we set $\gamma = 1/7$ such that the expected length of sickness lasts 1 week. We recognize that the average disease lasts longer, but it appears that few people are infectious and asymptomatic for longer than a week. Lauer et al. (2020) report a median incubation period for COVID-19 of five days and that 98% of people who develop symptoms after an exposure do so within 11.5 days. Below, we offer some robustness with regard to this choice.

We calibrate β for the model to capture data on the doubling time at the onset of the Covid-19 outbreak. Specifically, we target an initial daily growth rate of the stock of infected equal to 30%, consistent with a doubling time of approximately 3 days.[5] From (2), we have that $\frac{N_i'(0)}{N_i(0)} = \beta - \gamma$, giving $\beta = 0.3 + \gamma = 0.443$. This implies a basic reproduction number of $R_0 = \frac{\beta}{\gamma} = 3.1$. Since there appears to be considerable uncertainty and disagreement about the value of R_0 we note that this strategy of backing out its value only relies on the expected duration of infectivity and aggregate data on doubling time in the early stages, two numbers that appear to be relatively well understood. However, some recent estimates suggest a higher value for R_0 and several authors work with a lower γ. We therefore offer a robustness exercise below where we use a higher value of R_0 and a correspondingly lower value of γ, while still hitting the initial 30% daily growth rate.

We work with the following period utility function,

$$u(a) = \log a - a + 1. \tag{15}$$

We think of the first part as the gross returns from social activity, in particular consumption,

[5] JHU's Center for Systems Science and Engineering reports a doubling time of 2–5 days in the US in the very early stages of the epidemic (Dong, Du and Gardner, 2020).

Covid Economics 9, 24 April 2020: 25-61

Covid Economics 9, 24 April 2020: 25-61

and of the second part as the cost associated with it. This gives an interior solution at $a^* = 1$ in a disease-free environment, with $u(1) = 0$.

We now turn to the cost of disease. We assume that the infection mortality rate for the disease is $\pi = 0.002$, independent of $N_i(t)$. This is lower than many estimates of the case mortality rate, but this smaller number is consistent with evidence that there are many undetected cases in real-world populations. Importantly, a constant death rate shuts down the health care externality which arises as health care deteriorates when many people are infected. As we show below, we still find that it is optimal to delay infections and typically optimal to avoid a high peak infection rate.

We think of the cost κ as equal to πv, where v is the value of a statistical life (VSL). We follow Greenstone and Nigam (2020) in assuming that $v =$\$10 Million for the US.[6] Roughly speaking, this number is based on evidence that a typical individual would pay \$10,000 to avoid a 0.1% probability of death. With a discount rate of ρ, this is equivalent to paying a constant stream of $\rho \times \$10,000$ to avoid this death risk, or equivalently \$1.37 per day. Compare this to US consumption per capita of approximately \$45,000 per year or \$123 per day, and we reach the conclusion that people would permanently give up over 1.1 percent of their consumption to avoid a 0.1 percent death risk.

To see how to map this into our model, ask someone with preferences (15) what fraction x of her consumption she would be willing to give up to avoid a 0.1 percent probability of death. If the answer is $x = 0.011$, then v solves

$$\frac{\log(1)}{\rho} - 0.001v = \frac{\log(1 - 0.011)}{\rho}$$

This implies $v = \frac{-1000 \log 0.989}{\rho} \approx 80,000$ in our model units. Multiplying these together gives $\kappa = \pi v = 160$.

We note that this value is in the same ballpark as the one chosen in several other recent paper on the outbreak. For instance, Alvarez, Argente and Lippi (2020) choose a five times higher fatality rate π but a far lower VSL of \$1.3 million. Similarly, Hall, Jones and Klenow (2020) work with a VSL 50% higher than Alvarez, Argente and Lippi (2020) (and so still far below our value) but a four-fold higher death rate which implies a very similar value of κ. Finally, Eichenbaum et al. (2020) pick a VSL of \$9.3 million and a fatality rate of 0.5 percent, which implies a two fold higher κ. We offer a robustness exercise below where we consider doubling κ to reflect a higher infected fatality rate or a higher VSL.

Finally, we assume that initially $N_i(0) = 10^{-6}$. Prior to date 0, we assume $A(t) = a^* = 1$,

[6] See Greenstone and Nigam (2020) for a useful discussion of this value and the use of VSL in calculations such as ours.

so the disease grew without any response in social activity. From a very low initial prevalence, this implies that approximately $N_i(0)/(R_0 - 1)$ individuals have recovered or died by date 0, leaving the remaining 0.9999985 individuals susceptible.

We summarize our calibration strategy in Table 1.

Parameters			
Parameter Description	Parameter	Value	Target
Conditional transmission prob.	β	$0.3 + \gamma$	Initial doubling time
Rate at which illness ends	γ	$1/7$	Duration until symptomatic
Cost of infection	κ	160	Death rate and VSL
Arrival rate of cure	δ	$0.67/365$	Exp. time until vaccine/cure
Discounting	ρ	$0.05/365$	Annual discount rate
Fraction initially affected	$N_i(0)$	10^{-6}	
Other			
Basic reproduction number	R_0	3.1	Implied by γ and β
Fraction initially susceptible	$N_s(0)$	0.9999985	no social distancing before $t = 0$

Notes: We calibrate the model at a daily frequency.

Table 1: Calibration

7.2 Results

We next turn to our quantitative results. As a simple benchmark, we begin with the basic SIRD model and then turn to laissez-faire equilibrium and the social optimum.

7.2.1 Basic SIRD Model

Figure 3 plots the dynamics of the pandemic in the SIRD model, that is without any behavioral response, $A(t) = a^* = 1$, under the assumption that a cure is not found. The top panel shows the share of people susceptible in levels. The bottom panel shows the share infected on a log scale. The pandemic unfolds rapidly even though only one out of one million individuals is initially infected. After several weeks, a sizable share of the population is infected. The infection rate peaks after seven weeks above 31 percent. As a consequence of the height of the peak, the benefits of herd immunity do not kick in before almost everyone is sick. By 14 weeks into the infection, only 5.3 percent of the population remains susceptible and 0.19% of the population has died, more than 600,000 people in a country the size of the United States. Although the pandemic ends quickly, the cost of the disease is substantial.

Covid Economics 9, 24 April 2020: 25-61

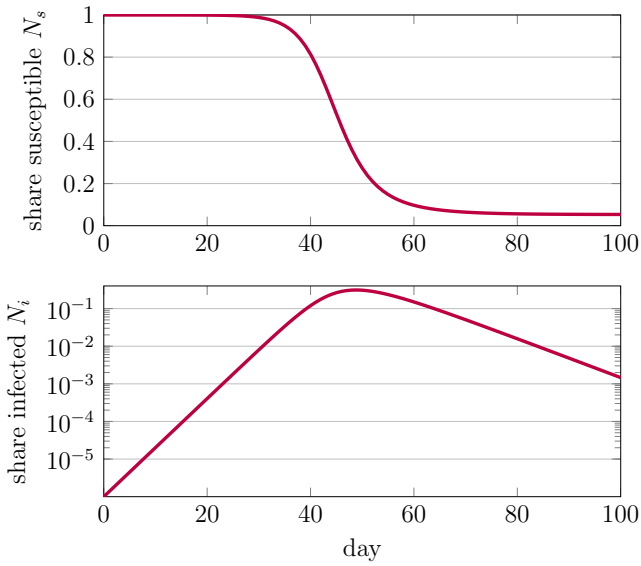

Notes: We set $A(t)$ to its optimal level absent disease, $a^* = 1$ and use equations (1) and (2).

Figure 3: Basic SIRD Model.

Measured in utility units, it is -136.6, equivalent to a permanent reduction in social activity to $a = 0.819$.

Of course, this model completely fails to capture the experience in places that did not institute any restrictions on social activity. For instance, Sweden hit 1040 total cases on March 15, 2020.[7] One month later, this number rose By mid April, this number had risen to about 11,927 confirmed cases despite the laissez-faire approach taken by the Swedish government. In contrast, our calibrated SIRD model predicts an increase by a factor of $e^{0.3\times30}$, or more than 8000-fold, during this period. Likewise, the SafeGraph micro-data document a remarkably uniform decline in individual social activity. The fact that this decline happened across the board in the US despite the large differences in policies also suggests that the basic SIRD model fails to capture a key aspect of this epidemic, namely that individual behavior responds to the risk of infection.

7.2.2 Laissez-Faire Equilibrium

We thus next turn to the disease dynamics in our laissez-faire equilibrium, which are depicted in Figure 4.[8] The top two panels again depict the share of susceptible and infected. The

[7]Retrieved from https://www.worldometers.info/coronavirus/country/sweden/.
[8]Note that all the figures are conditional on no cure having been found.

Covid Economics 9, 24 April 2020: 25-61

Covid Economics 9, 24 April 2020: 25-61

difference between laissez-faire and the basic SIRD model is stark. Despite the government not intervening at all, the peak infection rate is one tenth of the level in the SIRD model, 3.5 percent. In turn, the response of individual behavior substantially prolongs the epidemic, with the infection rate staying elevated for a much longer time, albeit at a lower level. This implies that the population reaches herd immunity at a far lower level of N_i, compared with the SIRD model. Taking into account the possibility of a cure, in expectation just under 50 percent of people eventually get sick. Thus the expected death rate is about half as high as in the model without a behavioral response.

The third panel shows that individuals reduce their social activity by as much as 40 percent. The fourth panel depicts the ratio of the effective reproduction number to the fraction of susceptible people, $R_e(t)/N_s(t)$ on a double log scale.[9] Recall that this is the number of newly infected individuals for each infected person which would prevail if everyone were susceptible for a given level of social activity $A(t)$. This falls substantially, driving down the doubling time for the disease, but remains strictly above 1. As a consequence, the fact that N_i eventually starts falling is a consequence of the stock of susceptible people becoming smaller. Putting this together, the total welfare loss in equilibrium is equivalent to a permanent reduction in social activity to $A = 0.854$.

We note that the dynamics of social activity, A, under laissez-faire mirrors the behavior of N_i. The reason is simply that there is little private incentives to lower social activity when the risk of individual infection is negligible. As a consequence, individuals do not restrict activity until infections are rampant.

Taken together, the internalized part of the disease risk substantially delays the "wave," lowers the peak infection rate, and slows infections to an extent that allows the population to achieve a higher asymptotic susceptible share. In fact, as we show next, the laissez-faire equilibrium dynamics are closer to the optimal dynamics than to the SIRD model. This suggests that explicitly modeling the internalized dimension of disease outbreak is of first order importance.

7.2.3 Optimal Policy

Figure 5 contrasts the laissez-faire with the optimal policy. Reflecting the external effects of social activity, the key property of the optimal policy is delay. While peak infection in the SIRD model occur after 49 days and the equilibrium behavioral response delays the peak until 62 days have lapsed, the optimal policy delays it for 341 days. Because infections increase more slowly, the peak level of infection is also far lower under the optimal policy,

[9] On a double log scale, the vertical distance between two points y_1 and y_2 is proportional $\log(\log y_1) - \log(\log y_2)$.

Covid Economics 9, 24 April 2020: 25-61

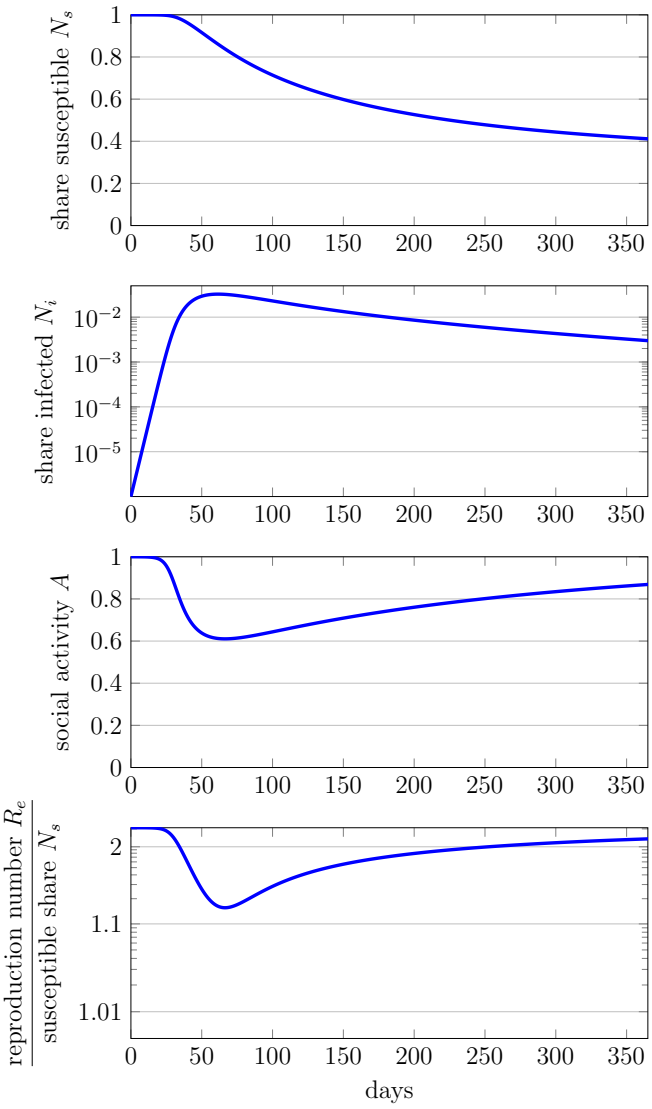

Notes: See Table 1 for calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 4: Laissez-Faire equilibrium.

Covid Economics 9, 24 April 2020: 25-61

at 0.28 percent. We stress that this strong desire to "flatten the curve" is true in a model without any explicit cost of peak-loading of infection rates, i.e. where κ is constant. A health care externality would make the case for flattening the curve even stronger.

An optimal policy buys time for a cure. Asymptotically, a bit more than $1-\gamma/\beta = 0.677$ of the population will eventually get sick, assuming a cure is never developed. This reflects the fact that activity will optimally asymptote back to a^*. Taking into account the possibility of a cure, however, reduces that to just 8.4 percent. Thus the expected death rate is one-sixth as large under the optimal policy as in equilibrium. The resulting welfare loss is equivalent to a permanent reduction in social activity to $A = 0.907$, significantly less than the equilibrium reduction to 0.854. The welfare cost is also well below the cost of permanently suppressing the disease by setting $R_e(t) = N_s(t)$, i.e. by setting $A = \sqrt{\gamma/\beta} = 0.568$ until a cure is found. The welfare cost of this policy is equivalent to permanently reducing social activity to $A = 0.870$.

The optimal policy achieves delay by acting preemptively. The third panel in Figure 5 shows the degree of social distancing in laissez-faire and optimum. At the outbreak of the disease, equilibrium behavior does not change because the risk of individual infection is negligible. Optimal policy, however, immediately curtails social activity. The planner recognizes that lowering the initial transmission rate buys time. In particular, the bottom panel shows that the ratio of the effective reproduction number to the share of susceptible individuals $R_e(t)/N_s(t)$ is far below its uncurtailed counterpart R_0 (which corresponds to the intercept of the laissez-faire time path). Even if a full outbreak is eventually inevitable, the social gains from immediate social distancing are of first order because of discounting and because of the hope for a cure. The cost, however, are of second order, since $u'(a^*) = 0$.[10]

Despite the initial reduction in social activity, a remarkable feature of the optimal policy is that social distancing is never extremely intense. The planner could, of course, push the effective reproduction number $R_e(t)$ below the share of susceptible people $N_s(t)$, ending the disease, but he never chooses to do so, instead choosing values slightly above $N_s(t)$. While the stock of infected people $N_i(t)$ eventually starts declining under the optimal policy, this is a consequence of the fact that the stock of susceptible people $N_s(t)$ falls substantially below 1, combined with a limited reduction in social activity $A(t)$.

To gain some intuition for this observation, we note that social activity eventually needs to return to its pre-pandemic level. The reason is that under any feasible policy, the share of infected individuals must converge to zero.[11] If $A(t)$ were converging to a number smaller

[10] In contrast, Alvarez, Argente and Lippi (2020) assume the economy without disease is in a corner and so reductions in social activity, lockdowns in their words, have a first order cost.

[11] Formally, we have that the sum of the number of recovered and deceased people evolve as $(N_r'(t) + N_d'(t)) = \gamma N_i(t)$ (equations 3 and 3). Since $N_r(t) + N_d(t)$ is bounded above by 1, their sum must converge,

Covid Economics 9, 24 April 2020: 25-61

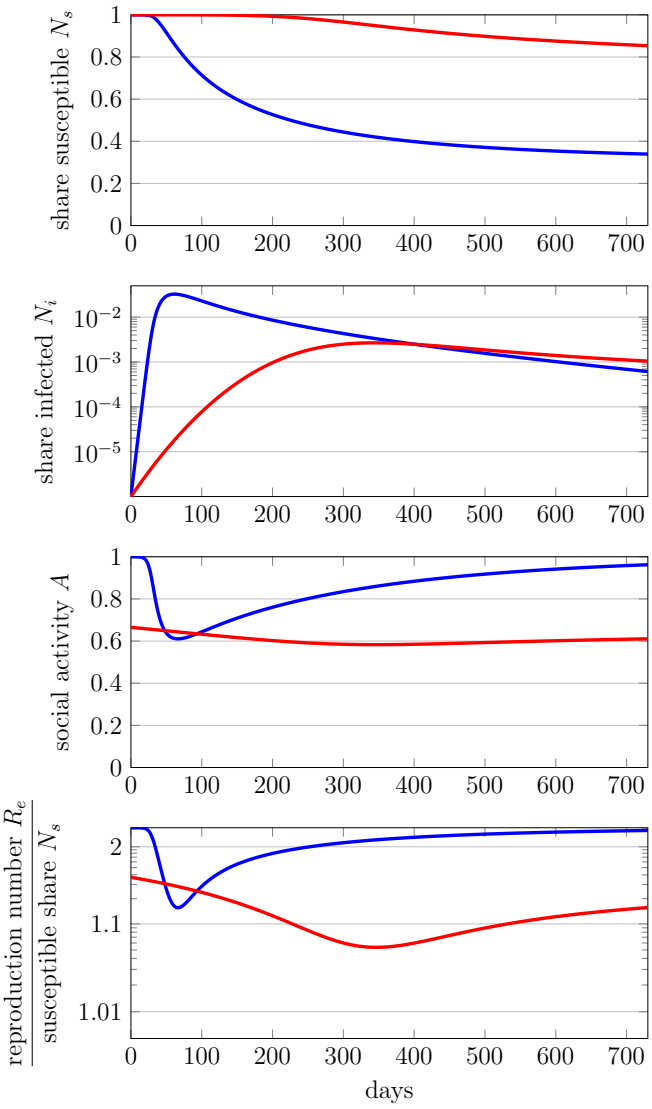

Notes: See Table 1 for calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 5: Optimal Policy vs Laissez-Faire.

Covid Economics 9, 24 April 2020: 25-61

than $a^* = 1$, there would be a first order gain from a temporary and small increase in A, while the cost would be negligible if $N_i(t)$ is sufficiently small. Thus long-runs social distancing cannot be optimal.

Now suppose we suppress $R_e(t)$ below $N_s(t)$ at some early time t. Doing this for a while will reduce the infection rate to a negligible share of the population. But since the number of infected people never reaches zero, any attempt to relax social distancing will quickly lead to a reemergence of the disease, quickly undoing the effect of keeping $R_e(t)$ below $N_s(t)$. Thus setting $R_e(t)$ below $N_s(t)$ only makes sense if the intent is to keep this policy in place forever. But we have just argued that this is not optimal.

Finally, we discuss the shape of the recovery. We note that, under laissez-faire, social activity is almost back to its pre-pandemic level after 2 years. This is not the case under the optimal solution, which curtails activity for decades or until a cure is found. That is the flip-side of delay: The optimal solution delays in the hope of finding a cure. If no cure is found, herd immunity only builds very slowly and so restrictions on social activity must persist far longer than under laissez-faire. These restrictions do disappear eventually, with $R_e(t)/N_s(t)$ converging to R_0 as the level of the susceptible population falls slightly below $1/R_0$.

Overall, an important observation is that the planner achieves a delay in infections over the first year of the pandemic without completely locking down the economy. The key instead is an early and long-lasting reduction in social activity that is moderate in magnitude.

7.3 Revisiting the Evidence on Change in Individual Behavior

We briefly re-visit the quantitative evidence from SafeGraph from section 3. We note that none of the targets we chose in calibrating our model were actually related to the response of social activity to the Covid-19 outbreak. We do not take a stance on whether the hastily implemented lockdowns and mobility restrictions are close to the social optimum. But we believe that the response in individual behavior witnessed prior to implementation of any policy measures should be picked up by the laissez-faire equilibrium.

Figures 1 and 2 show a decline of some 25-50% in terms of activity across our three different metrics. We note that these metrics all have an inherent cardinality and so their decline is quantitatively meaningful. This is well captured by our laissez-faire model which suggests a decline of individual activity by 40% as can be seen in Figure 5.

Furthermore, we note that the model also captures the pace of the decline surprisingly well. In the model, equilibrium social activity declines form 98% on day 20 to 63% on day

which requires $N_i(t)$ converges to zero.

Covid Economics 9, 24 April 2020: 25-61

51. The contraction in the SafeGuard data was even faster. For example, Figure 1 shows that POI visits fell from 98 percent on March 11 to 45 percent on March 20 as state and local governments started to issue stay-at-home and shelter-in-place orders. The difference might reflect a delayed understanding of the seriousness of the disease.

We complement this analysis with an exploration of aggregate data available from Google at https://www.google.com/covid19/mobility/ for Sweden. This data uses Google's location history to track mobility along various dimensions relative to a baseline. Sweden has arguably been the Western country with the least restrictions on mobility and social activity throughout March 2020. The Swedish data currently cover the time period from February 23 to April 5.

We restrict attention to the categories "Retail & recreation," "Transit Stations," and "Workplace". The data suggest a reduction of activity of about 25%, 39%, and 28% for these three categories, respectively. This suggests that Swedes indeed internalize parts of the risk associated with the disease and therefore reduced activity, in line with our model. Like the US micro-data presented in section 3, the data squares up well with the laissez-faire dynamics under our preferred model calibration, both in terms of the size and pace of the decline.

We also note that we could have directly targeted this type of data in our calibration, in particular to select a value for κ. That is, a natural approach would be to let the individual response in behavior "reveal" the perceived cost of infection instead of relying on controversial and noisy direct measures of the infection mortality rate, π, and the value of a statistical life, v. However, as we show below, the responsiveness of the time path of $A(t)$ to variation in $\kappa \equiv \pi v$ is limited and so we have opted with the direct measures. Nonetheless, we believe that the response of behavior absent policy intervention can be a useful source of information to reveal perceptions given the current level of uncertainty surrounding many key parameters.

8 Robustness

In this section, we show the laissez-faire and optimal dynamics given perturbations to various parameters. Given the large uncertainty around many of the key modeling parameters we consider fairly large changes. We find that our main findings are robust to alternative parameter choices. In particular, we find that a strong laissez-faire equilibrium reduction in social activity; an immediate and persistent optimal reduction in social activity that only disappears in the very long run; and with one notable exception, a limit in the extent of social distancing such that the effective reproduction number $R_e(t)$ stays above the share of

susceptible individuals $N_s(t)$.

8.1 Alternative Utility Function

In Figure 6, we modify the utility function to $u(a) = -\frac{1}{2}(1-a)^2$. This leaves $u(1) = u'(1) = 0$ and also leaves $u''(1)$ unchanged. However, it implies less curvature in the utility function and in particular that $u'(0)$ is finite. This modification reduces the marginal value of social activity, and so both equilibrium and optimal social activity fall with this calibration. Still, all of our takeaway messages hold with this calibration: optimal policy has an immediate and sustained reduction in social activity, although the extent of it is limited and so the effective reproduction number remains above the share of susceptible individuals.

8.2 Cost of Disease

This subsection shows the results for $\kappa = 320$, that is we double our baseline parameter value of the expected cost of infection. For instance, this can be viewed as capturing a death rate of $\pi = 0.004$ (compared with $\pi = 0.002$ in the baseline). We show the resulting dynamics for optimum and laissez-faire in Figure 7.

Doubling the cost of infection cuts the peak equilibrium infection rate by more than half, from 3.3 percent to 1.5 percent. The peak optimal infection rate declines by a larger percent, although from a very low level of 0.276 percent to a negligible 0.014 percent. The optimal policy delays a sizable outbreak for an extremely long time, effectively until a cure has been found with very high probability. It does so by initially reducing social activity and then continuing to suppress it to a level with $R_e(t)$ just above $N_s(t)$ for decades. While not visible in the figure, social activity is eventually allowed to return to normal and a fraction less than γ/β of individuals remain disease free. We note of course that all of this again is conditional on no cure being found, an exceedingly unlikely scenario.

8.3 Higher Duration of Infectivity and Higher R_0

Some authors use a considerably longer duration of infectivity $\frac{1}{\gamma}$. For instance, Hall, Jones and Klenow (2020) set $\gamma = \frac{1}{18}$. We follow them here but maintain the target of a 30 percent daily growth rate in a world without social distancing. We therefore adjust $\beta = 0.3 + \frac{1}{18}$. This gives a substantially higher basic reproduction number of $R_0 = 6.4$.

We report the corresponding results in Figure 8. This has qualitatively little impact on the results, but there are significant quantitative differences. Most noticeably, with a much higher basic reproduction number, social distancing is less effective at reducing the

Covid Economics 9, 24 April 2020: 25-61

Covid Economics 9, 24 April 2020: 25-61

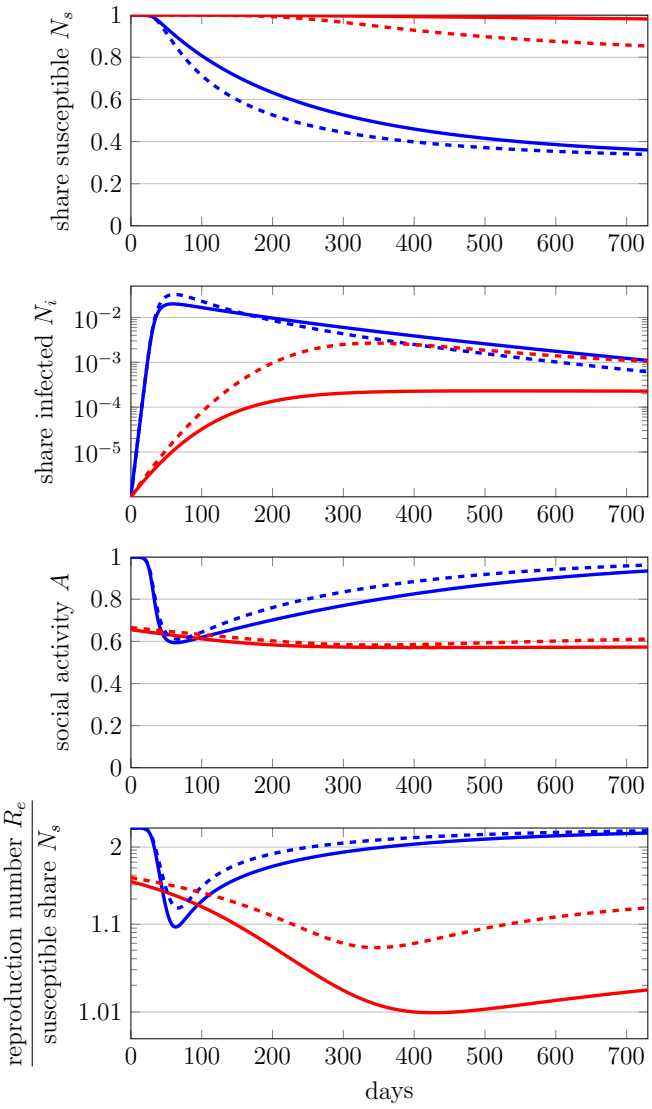

Notes: See Table 1 for calibration. Only change: We assume the utility from social activity is $u(a) = -\frac{1}{2}(1-a)^2$. Dashed lines show the baseline calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 6: Optimal Policy vs Laissez-Faire with quadratic utility.

Covid Economics 9, 24 April 2020: 25-61

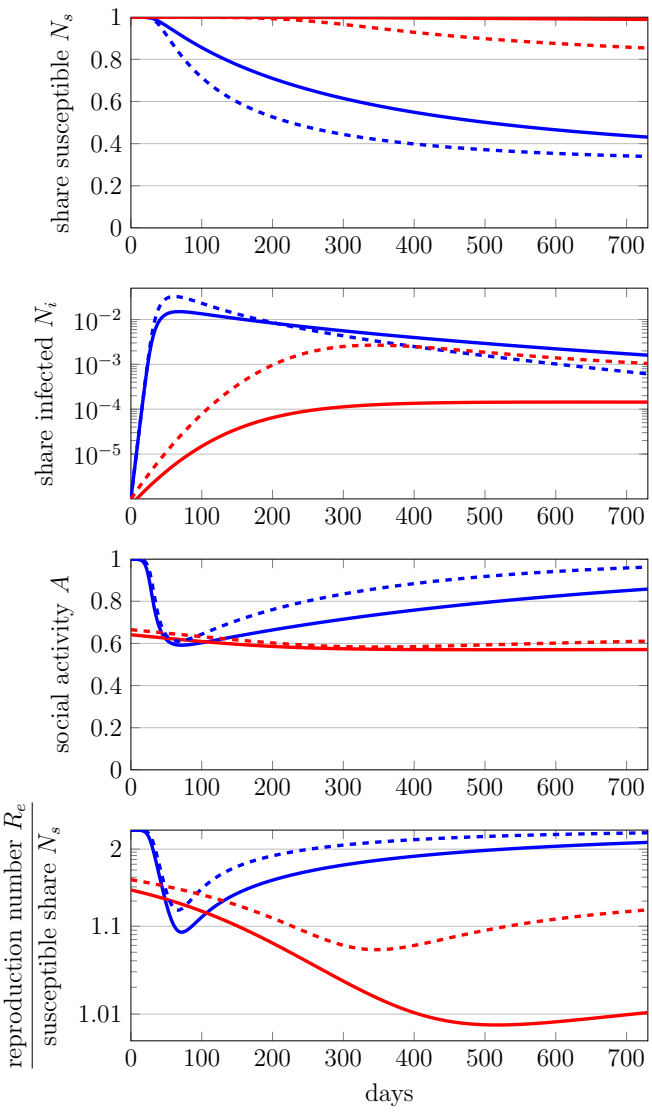

Notes: See Table 1 for calibration. Only change: We set $\kappa = 320$ instead of $\kappa = 160$. Dashed lines show the baseline calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 7: Optimal Policy vs Laissez-Faire with high cost of infection.

infection rate. The peak infection rate in equilibrium is substantially elevated. Even more noticeably, optimal policy now allows for a substantial wave of infections, since the cost of suppressing it is prohibitive. Conversely, both equilibrium and optimal policy see an even larger contraction of social activity. In either case, optimal policy still keeps $R_e(t) > N_s(t)$, as before.

8.4 Lower Cure Probability

We have so far assumed that the expected time until a cure is 1.5 years. Figure 9 shows what happens if we double this to 3 years. The laissez-faire equilibrium dynamics are effectively the same, reflecting the fact that individuals are insensitive to the discount rate, and difficulties in finding a cure are equivalent to a reduction in discounting. The dynamics of the disease, however, differ sharply under the optimal policy. In particular, the planner allows a substantial wave of infections to occur, so the expected number of sick (and hence fatalities) rises much more quickly. This reflects a reduction in the benefits of delay.

Still, the path of optimal policy does not change qualitatively. Immediate and long-lasting social distancing is optimal, only briefly interrupted by a more severe constraint on activity—still with $R_e(t) > N_s(t)$—at the peak of the infection. In particular, social distancing lasts even once it appears that the share of susceptible people has reached a plateau.

8.5 Stock of Initially Infected

We have experimented with the stock of the initially infected for several reasons. First, one might think that as the fraction of initially infected individuals becomes exceedingly small one of our key policy lessons—that the optimal policy immediately curtails social activity to buy time—no longer holds. This is indeed true in the limit: For $N_i(0)$ small enough the expected uncurtailed outbreak date is so far in the future that the (social) gains from social distancing must vanish. However, suppose we start with $N_i(0) = 1/7 \cdot 10^9$, i.e. patient zero, a natural lower bound on the initial seed. With $\beta - \gamma = 0.3$, in just 30 days $N_i(t)$ exceeds 10^{-6} if there is no social distancing. The delay motive we have appealed to above therefore remains quantitatively powerful.

Second, we have also substantially increased $N_i(0)$ up to 1 percent. This could for instance capture the situation in places where the Covid-19 appeared first, or it could capture the situation in places where the initial policy response was botched or individuals did not respond because of, say, false information. We plot the corresponding time paths for equilibrium and optimum in Figure 10.

Covid Economics 9, 24 April 2020: 25-61

Covid Economics 9, 24 April 2020: 25-61

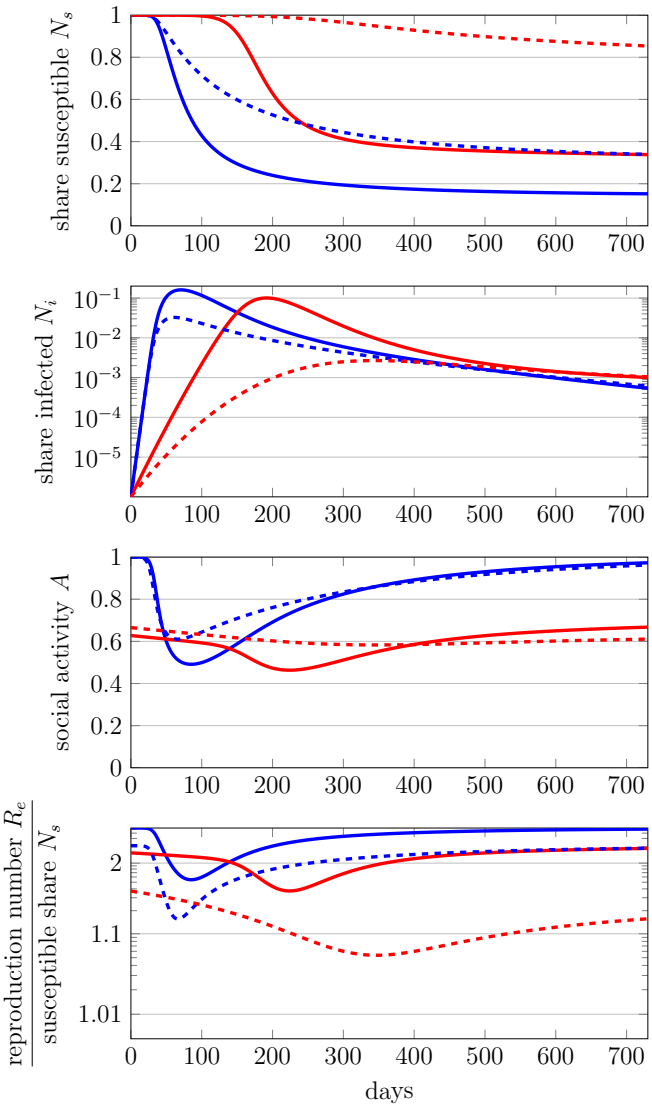

Notes: See Table 1 for calibration. Only change: We set $\gamma = 0.056$ (instead of $\gamma = 0.143$) and $\beta = 0.356$ such that $R_0 = 6.4$ (instead of 3.1). Dashed lines show the baseline calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 8: Optimal Policy vs Laissez-Faire with long duration of infectivity.

Covid Economics 9, 24 April 2020: 25-61

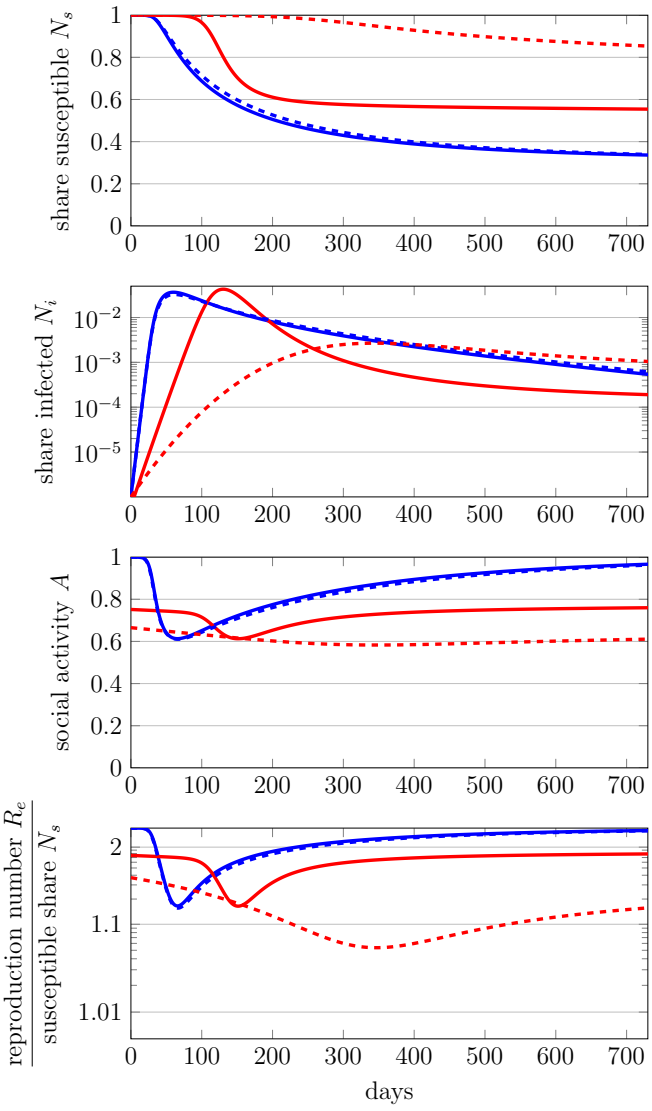

Notes: See Table 1 for calibration. Only change: We set $\delta = \frac{0.33}{365}$ (instead of $\frac{0.67}{365}$). Dashed lines show the baseline calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 9: Optimal Policy vs Laissez-Faire with low probability of cure.

The qualitative patterns are largely unchanged but the optimal policy acts more aggressively. An important observation, shown in the bottom panel of the figure, is that in this case optimal policy suppresses social activity for the first 67 days to such an extent that $R_e(t)$ falls below $N_s(t)$; this cannot be shown on the double log scale. As a consequence of the high initial stock of infections, the planner has less room to delay the wave of infections and so drives down the stock of infections from the outset. Nonetheless, optimal policy monotonically relaxes social distancing, ultimately relying on a reduction in the share of susceptible individuals for the continued decline in infections. The equilibrium time path, in turn, remains largely unchanged. Starting with more infections just shifts it ahead in time.

9 Next Steps

9.1 Heterogeneity

Several key dimensions of heterogeneity come to mind. In particular, a key feature of the Covid-19 pandemic is the case fatality rate by age: Individuals age 65 and older seem to have a far higher fatality rate than younger ones. At the same time, one might argue that they have a smaller cost associated with a reduced social activity since they are primarily retired.

We believe that our current framework can naturally and easily be extended to accommodate these dimensions of heterogeneity. With k types of people, we would need to solve a system of $2k$ differential equations in equilibrium. Assuming a utilitarian social planner that can separately choose the level of social activity for each type, there would be a corresponding system of $2k$ differential equations describing the optimal policy. We are currently working on that extension.

9.2 Antibody Tests

We have assumed that everyone knows when their infection ends. In reality, few people are currently tested, and so many recovered people cannot be sure whether they were sick. We can capture this by separating the recovered state into two categories, depending on whether the individual is aware they were sick. Someone who does not know they have recovered from Covid-19 must choose the same level of social activity as the susceptible and infected, given their lack of information. In this environment, antibody testing may be useful for letting recovered people know that they are recovered. This has the dual benefit of allowing those individuals to return to a normal level of social activity and of discouraging social activity by those who are still susceptible or infected.

Covid Economics 9, 24 April 2020: 25-61

Covid Economics 9, 24 April 2020: 25-61

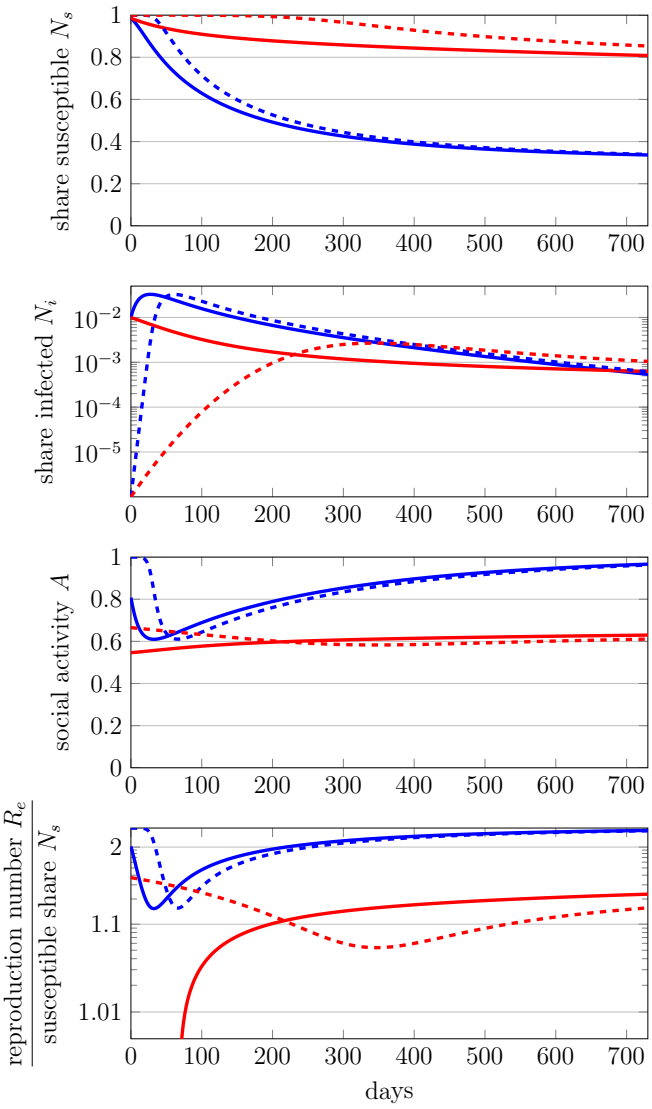

Notes: See Table 1 for calibration. Only change: We set $N_i(0) = 10^{-2}$ instead of $N_i(0) = 10^{-6}$. This also reduces the initial value of $N_s(0)$ to 0.985 instead of 0.9999985. Dashed lines show the baseline calibration. The second plot is drawn on a log scale and the fourth plot on a double log scale.

Figure 10: Optimal Policy vs Laissez-Faire with higher stock of initially infected.

9.3 Vaccines

We have so far assumed that at a random future date, a cure will end all costs associated with the disease. This is an unlikely scenario. More plausible is the gradual roll-out of a vaccine that shifts some people from susceptible to recovered without enduring an infection. To the extent the vaccine provides imperfect coverage, even a 100 percent vaccination rate will leave some people vulnerable to future infections, and potentially slow the recovery in social activity.

10 Conclusion

This paper uses standard dynamic optimal choice tools from economics to integrate privately optimal behavior and policy analysis into an epidemiological model. Our quantitative exercises reveal several robust patterns: Even in laissez-faire, individuals sharply reduce social activity due to risk of infection gradually. However, an optimal policy immediately curtails social activity immediately and delays the full outbreak to buy time. But even the optimal policy lets the disease eventually run its course: If no cure is found a large number of people eventually get infected. The framework we develop is general and tractable: Optimal behavior and policy are encoded in an additional set of differential equations that can jointly be solved with the epidemiological block. We therefore view the tools offered here as a natural building block to explore the role of various additional features of Covid-19 and their interplay with optimal policy.

Covid Economics 9, 24 April 2020: 25-61

Covid Economics 9, 24 April 2020: 25-61

References

Alvarez, Fernando, David Argente, and Francesco Lippi, "A Simple Planning Problem for COVID-19 Lockdown," March 2020.

Atkeson, Andrew, "What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios," Technical Report, National Bureau of Economic Research 2020.

Barro, Robert J, José F Ursua, and Joanna Weng, "The Coronavirus and the Great Influenza Epidemic-Lessons from the coronavirus potential effects on mortality and economic activity," 2020.

Budish, Eric, "$R < 1$ as an Economic Constraint: Can We "Expand the Frontier" in the Fight Against Covid-19?," April 2020.

Dewatripont, Mathias, Michel Goldman, Eric Muraille, and Jean-Philippe Platteau, "Rapid identification of workers immune to COVID-19 and virus-free: A priority to restart the economy," Technical Report, Discussion paper, Universit Libre de Bruxelles 2020.

Diamond, Peter A., "Aggregate demand management in search equilibrium," *Journal of Political Economy*, 1982, *90* (5), 881–894.

_ **and Eric Maskin**, "An Equilibrium Analysis of Search and Breach of Contract, I: Steady States," *The Bell Journal of Economics*, 1979, *10* (1), 282–316.

Dong, Ensheng, Hongru Du, and Lauren Gardner, "An interactive web-based dashboard to track COVID-19 in real time," *The Lancet infectious diseases*, 2020.

Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt, "The Macroeconomics of Epidemics," April 2020.

Garibaldi, Pietro, Espen R. Moen, and Christopher A. Pissarides, "Modelling contacts and transitions in the SIR epidemics model," Technical Report, CEPR Covid Economics Working Paper, Issue 5 2020.

Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and Jose Victor Rios-Rull, "Health versus Wealth: On the Distributional Effects of Controlling a Pandemic," Technical Report, Working Paper 2020.

Greenstone, Michael and Vishan Nigam, "Does Social Distancing Matter," March 2020.

Greenwood, Jeremy, Philipp Kircher, Cezar Santos, and Michèle Tertilt, "An equilibrium model of the African HIV/AIDS epidemic," *Econometrica*, 2019, *87* (4), 1081–1113.

Hall, Robert E., Charles I. Jones, and Peter J. Klenow, "Trading off Consumption and COVID-19 Deaths," 2020.

Jones, Callum, Thomas Philippon, and Venky Venkateswaran, "Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home," 2020.

Kaplan, Greg, Benjamin Moll, and Gianluca Violante, "Pandemics According to HANK," Technical Report, Working Paper 2020.

Keppo, Jussi, Marianna Kudlyak, Elena Quercioli, Lones Smith, and Andrea Wilson, "For Whom the Bell Tolls: Avoidance Behavior at Breakout in COVID19," Technical Report, Working Paper 2020.

Kermack, William Ogilvy and Anderson G McKendrick, "A contribution to the mathematical theory of epidemics," *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 1927, *115* (772), 700–721.

Kremer, Michael and Charles Morcom, "The effect of changing sexual activity on HIV prevalence," *Mathematical biosciences*, 1998, *151* (1), 99–122.

Krueger, Dirk, Harald Uhlig, and Taojun Xie, "Macroeconomic Dynamics and Reallocation in an Epidemic," Technical Report, CEPR Covid Economics Working Paper, Issue 5 2020.

Lauer, Stephen A, Kyra H Grantz, Qifang Bi, Forrest K Jones, Qulu Zheng, Hannah R Meredith, Andrew S Azman, Nicholas G Reich, and Justin Lessler, "The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application," *Annals of internal medicine*, 2020.

Piguillem, Facundo and Liyan Shi, "The optimal covid-19 quarantine and testing policies," Technical Report, Einaudi Institute for Economics and Finance (EIEF) 2020.

Covid Economics 9, 24 April 2020: 25-61

Covid Economics **Issue 9, 24 April 2020**

Using the eye of the storm to predict the wave of Covid-19 UI claims[1]

Daniel Aaronson,[2] Scott A. Brave,[3] R. Andrew Butters,[4] Daniel W. Sacks[5] and Boyoung Seo[6]

Date submitted: 16 April 2020; **Date accepted:** 18 April 2020

We leverage an event-study research design focused on the seven costliest hurricanes to hit the US mainland since 2004 to identify the elasticity of unemployment insurance filings with respect to search intensity. Applying our elasticity estimate to the state-level Google Trends indexes for the topic "unemployment," we show that out-of-sample forecasts made ahead of the official data releases for March 21 and 28 predicted to a large degree the extent of the Covid-19 related surge in the demand for unemployment insurance. In addition, we provide a robust assessment of the uncertainty surrounding these estimates and demonstrate their use within a broader forecasting framework for US economic activity.

Covid Economics 9, 24 April 2020: 62-79

1 The views herein are our own and do not necessarily represent those of the Federal Reserve Bank of Chicago or the Federal Reserve System. R. Andrew Butters, Daniel W. Sacks, and Boyoung Seo would like to thank the generous funding of the Kelley School of Business at Indiana University. The authors thank, without implicating, Timothy Slaper, and the Indiana Department of Workforce Development for helpful conversations and seminar participants at Indiana University, as well as Ross Cole and Michael Fogarty for excellent research assistance.

2 Vice President and Director of Microeconomic Research, Federal Reserve Bank of Chicago.
3 Senior Policy Economist, Federal Reserve Bank of Chicago.
4 Assistant Professor, Kelley School of Business, Indiana University.
5 Assistant Professor, Kelley School of Business, Indiana University.
6 Assistant Professor, Kelley School of Business, Indiana University.

Covid Economics 9, 24 April 2020: 62-79

1. Introduction

The advent of private sector "big data" has the potential to substantially alter the landscape of economic statistics and forecasting, as researchers use this data to develop statistical methods that can be used to improve the timing and corroborate the accuracy of official statistics.[1] One of the most broadly known examples of these new data are the Google Trends indexes, which track search patterns for words or phrases entered into the Google search engine.[2] We show how to make use of this timely information in order to calibrate the take-up for unemployment insurance (UI). Our key insight is to use the historical experience of US hurricanes. This approach is especially promising because these events generate very large and sharp spatial changes in new UI claims.[3] We apply this insight to the seven costliest hurricanes to make landfall in the mainland US since 2004 and show that the Google Trends data is highly predictive of subsequent UI take-up. Furthermore, using the same model, we demonstrate that the unprecedented take-up in UI during the first few weeks of the Covid-19 pandemic was largely predicted on the sole basis of Google Trends.

Extreme events, like the Covid-19 pandemic, are difficult to handle for traditional macroeconomic forecasting models that rely on standard sources of economic data that often take weeks, and in some cases months, to be released. Even initial UI claims, one of the rare economic series available at a weekly frequency, has a lag of 5-12 days. But the timeliness of weekly claims makes it an important and commonly used leading indicator of the business cycle.[4] As such, modern forecasting models that make use of data series observed at mixed frequencies rely heavily on the signal UI claims provides in the real-time data flow (see, for example, Brave et al. (2019)). Being able to have both an early and accurate indication of its movements, therefore, has value to economic researchers. To demonstrate, we show how our short-term forecasts of UI claims during the Covid-19 pandemic can be incorporated into traditional macroeconomic forecasting models.

2. Background on Forecasting UI Claims with Google Trends

Forecasters have increasingly become aware of the value of the Google Trends data. For example, Choi and Varian (2012) show that searches on the topic of "unemployment" are capable of predicting turning points in US initial UI claims.[5] A key limitation to Choi and Varian's analysis, however, is

[1]For examples of how firm level information can provide a leading indicator of construction spending and payroll employment, see Aaronson et al. (2016) and Cajner et al. (2019), respectively.

[2]See https://trends.google.com/trends/?geo=US.

[3]In this respect, our research design is similar to the "identification through heteroscedasticity" approach put forth in Rigobon (2003). The use of hurricanes as a source of identification is common in economics. See, for example, Gallagher and Hartley (2017), Deryugina (2017), Deryugina et al. (2018), and Ortega and Taşpinar (2018).

[4]For example, it is included in the Conference Boards Leading Economic Index and is closely followed by financial markets to gauge future changes in payrolls and unemployment. See also Gordon (2009).

[5]Similar results are reported by D'Amuri and Marcucci (2017) for the US; McLaren and Shanbhogue (2011) for the UK; Askitas and Zimmermann (2009) for Germany; Suhoy (2009) for Israel; and Tuhkuri (2016) for Finland.

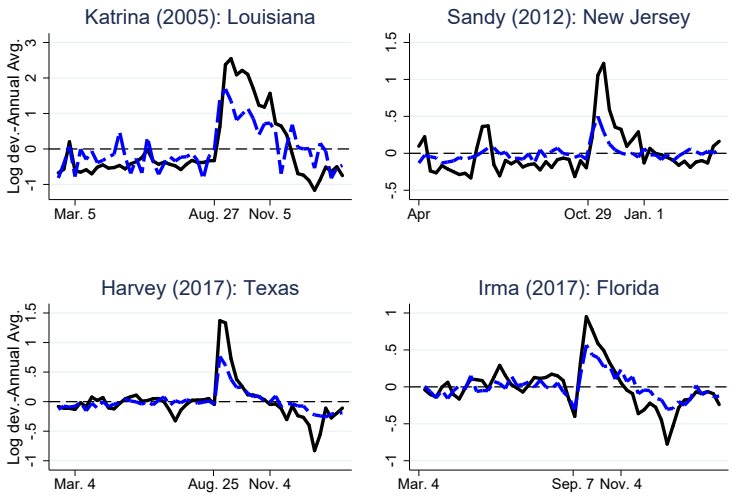

Figure 1 Selected Hurricane Event Studies

Notes: This figure displays the time series of the log share of initial unemployment insurance claims (black line) and the log ratio of the Google Trends measure (blue dashed line) for the "unemployment" topic less the annual average for the six months leading up to and following the landfall of each of our top four hurricane event studies.

that Google searches on unemployment reflect a catch-all term for labor markets, including not only rising demand for unemployment insurance during downturns but simultaneously potential news about future labor market conditions. This can lead the long-run association between initial UI claims and search intensity as measured by Google Trends to potentially be weak even if the two are very highly related during turning points in the business cycle.[6] We overcome this limitation by looking at the variation in both initial UI claims and the Google Trends unemployment index around the landfall of hurricanes.

Hurricanes have a tremendous short-run effect on local economies, resulting in temporary surges in new UI claims among impacted areas. In the most extreme cases, new UI claims increase by as much as 300 log points. They also generate a substantial spike in unemployment-related searches on Google (e.g., see figure 1, more detail on data below). Because the hurricane induces search activity that is driven by demand for unemployment insurance (rather than news, for instance, about a labor market recovery), the relationship between initial UI claims and the Google Trends index around

[6]To the extent there is measurement error in the Google Trends series, the estimated relationship might also suffer from attenuation bias.

Covid Economics 9, 24 April 2020: 62-79

hurricane landfall can be used to more reliably calibrate models of the take-up for unemployment insurance.

Using a research design built on this premise described in section 3, we show that the Google Trends unemployment topic index is highly predictive of subsequent unemployment insurance take-up after a hurricane's landfall. This is true in both an in-sample and out-of-sample sense, the latter of which is verified using a leave-one-out cross-validation procedure to quantify model uncertainty. Applying our elasticity estimates to state-level changes in the Google Trends indexes during the Covid-19 pandemic, we next show in section 4 that our model would have accurately predicted the historic surge in new UI claims for the US in March 2020. An extension of the Brave et al. (2019) real-time forecasting model, presented in section 4, documents the impact that this historic surge in new UI claims is projected to have on US gross domestic product (GDP), payroll employment, and the unemployment rate.

Covid Economics 9, 24 April 2020: 62-79

3. Data and Research Design

The primary data used for our analysis is the initial claims of unemployment insurance of US states and territories and their Google Trends indexes measuring unemployment-related search histories. In figure 2, we display the time series of the aggregate versions of both of these variables over the time period from January 2004–March 2020.

Given our research design, we focus our analysis over a set of "event-windows" that encompass the top ten costliest hurricanes (in 2020 dollars), filtering on those that made landfall in the mainland US (NOAA (2020); see table 1). This leaves us with seven hurricanes: Katrina in 2005, Harvey in 2017, Sandy in 2012, Irma in 2017, Ike in 2008, Wilma in 2005, and Rita in 2005.[7]

We build an event window for each hurricane that corresponds to the six months leading up to and following landfall. Then, for each hurricane, we identify the state most heavily affected and compare its Google search and initial unemployment insurance claim experience relative to the nation as a whole. The most affected states are Louisiana for Katrina, Texas for Harvey, New Jersey for Sandy, Florida for Irma, Texas for Ike, Florida for Wilma, and Texas for Rita. Table 1 summarizes each of the events in our sample.

For each event, we gather the weekly Google Trends unemployment index for both the affected states as well as the country as a whole and the weekly initial unemployment insurance claims filed for both the affected state and the entire US.[8] The Google Trends indexes provide a time series index of the volume of queries for a particular topic users enter into the Google search engine in a

[7]The cost estimate of Hurricane Maria, which did not make it to the mainland US but affected Puerto Rico and neighboring islands, would make it the third costliest hurricane.

[8]The initial unemployment insurance (UI) claims come from the FRED database maintained by the St. Louis FED and made publicly available at https://fred.stlouisfed.org/. For both the affected state and the US UI claims, we use non-seasonally adjusted data.

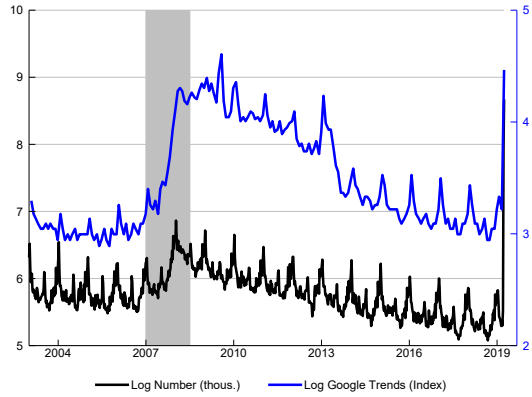

Figure 2 Initial Claims and Google Trends of Unemployment Searches

Notes: This figure displays the time series of initial unemployment insurance claims (black, non-seasonally adjusted (nsa)) and the Google Trends index (blue, nsa) for the "unemployment" topic over the time period of January 2004 to March 2020. US recessions currently defined by the National Bureau of Economic Research are shaded in gray.

Covid Economics 9, 24 April 2020: 62-79

given geographic area. The index is based on the overall share a search term or topic makes up of the total number of searches over that time period for the region in question. Google Trends then standardizes the index such that the maximum of the time series is normalized to 100. In an effort to capture many possible alternative search terms and/or combinations that are all likely to be related to "unemployment," we leverage Google Trends' "broad matched" unemployment "topic" index.

We measure the association between Google Trend search intensity and UI claims with the following equation:

$$\ln\left[\frac{\text{UIClaims}_{1ht}}{\text{UIClaims}_{0ht}}\right] = \alpha_h + \beta \ln\left[\frac{\text{GoogleTrends}_{1ht}}{\text{GoogleTrends}_{0ht}}\right] + \varepsilon_{ht} \quad (1)$$

where h indexes each of the 7 hurricane events, and t indexes the weeks within the year-long window for each event. Because Google Trends search activity is reported almost in real-time, while UI claims are reported with a lag, the estimates from this model could be used directly as a means of "nowcasting" UI claims.[9] We accommodate arbitrary seasonal and secular trends common across states by estimating the regression specification in terms of log-shares on log-ratios, where the subscript 1 references the affected state and the subscript 0 denotes our reference group, which in

[9]This process of predicting the present, or "nowcasting" as it was termed by Giannone et al. (2008), has become an important part of the daily workflow of many private sector analysts and economic researchers.

Event	Damage ($2020)	Event Date	Affected State	Control State	Placebo State	Start Date	End Date
Katrina	170B	8/27/2005	Louisiana	Arizona	New Jersey	2/2005	2/2006
Harvey	131B	8/25/2017	Texas	Michigan	New Jersey	2/2017	2/2018
Maria	95B	9/23/2017	-	-	-	-	-
Sandy	74B	10/29/2012	New Jersey	Washington	Texas	4/2012	4/2013
Irma	53B	9/7/2017	Florida	Missouri	New Jersey	3/2017	3/2018
Andrew	50B	8/23/1992	-	-	-	-	-
Ike	37B	9/7/2008	Texas	Michigan	New Jersey	3/2008	3/2009
Ivan	29B	9/13/2004	-	-	-	-	-
Wilma	26B	10/25/2005	Florida	Missouri	New Jersey	4/2005	4/2006
Rita	25B	9/25/2005	Texas	Michigan	New Jersey	3/2005	3/2006

Table 1 Summary of Hurricane Events

Notes: This table summarizes the top ten hurricanes used in our event study analysis, including the approximate mid-point of each hurricane's duration, the affected state used for each event, the matched control state with a comparable UI system, the placebo state, and the +/- 6 month window that comprises the event window. Of the top ten hurricanes, Maria, Andrew, and Ivan were excluded from our analysis because they either did not hit the mainland US or occurred before Google Trends data was available.

Covid Economics 9, 24 April 2020: 62-79

our baseline set of results will be the US as a whole. To accommodate differences in the average number of UI claims for the states of various population sizes across our events, we also include a set of event fixed-effects (α_h) in the above equation.

Hsu et al. (2018) document considerable heterogeneity across US states in the generosity of unemployment insurance. As generosity is likely to affect UI take-up, we also consider an alternative specification which instead matches each hurricane-affected state with the single US state that is simultaneously most like it in terms of UI generosity and the furthest from it in physical distance. The former ensures that take-up (and hence search activity) is likely to be similar in these states all else equal, while the latter ensures that the comparison state is not likely to suffer any major impact from the hurricane. In this instance of equation 1, the subscript 1 refers to the hurricane affected state and the subscript 0 refers to the comparison, or control state. The chosen comparison state best matching the criteria described above for each of the seven hurricanes is noted in table 1.[10] We also show that our research design is robust to a shorter-length event window, using (population) weighted regressions, and results in a null result for a "placebo" type event study.[11]

[10] The matching process is detailed in the supporting materials. To summarize, we rescaled the maximum UI benefit measure of Hsu et al. (2018) to be in standard deviation units from the cross-sectional average for our sample period in order to facilitate easily making comparisons across states on the dimension of UI generosity.

[11] For the "placebo" exercise, we assign each hurricane to one of the affected states in our study that is outside the region of the most affected state for that hurricane. We do this so as not to confound the multiple state effects of these hurricane events.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Google Trends Ratio	1.162	0.982	1.124	1.143	1.144	1.237	1.156	1.231
	(0.201)	(0.252)	(0.136)	(0.346)	(0.330)	(0.076)	(0.150)	(0.138)
Left Out Event		Katrina	Harvey	Sandy	Irma	Wilma	Ike	Rita
Adj. R-squared	0.509	0.309	0.488	0.516	0.495	0.554	0.523	0.558
Num. Observations	366	314	314	313	314	314	313	314

Table 2 Event Study Regression Results

Notes: This table reports regression results for equation 1, where we report the coefficient on the Google Trends Ratio term and its standard error, estimated by 10,000 event-level bootstrap and jackknife samples, in parentheses. In column 1, we report estimates using all of the hurricane events; and in each subsequent column, we report results from leaving each of the seven hurricane events out individually.

Covid Economics 9, 24 April 2020: 62-79

4. Results

Table 2 reports our estimates from equation 1.[12] In our baseline specification (column 1), the estimated coefficient on the Google Trends measure is both large and positive, with an elasticity of 1.16 for initial UI claims, as well as precisely estimated, with a standard error (estimated from 10,000 event-level bootstrap and jackknife samples) of 0.20. Given the precision of this estimate, we undertook an additional leave-one-out cross-validation of the regression to ensure that it was not overfit to a particular event. We report the estimates of the search elasticity from each of these additional specifications in columns 2-8 of table 2. The point estimates that resulted from this exercise were in the range of [0.98, 1.23], mostly in-line with the standard error of our overall estimate.

Using each leave-one-out estimate as the means to produce a one-week ahead forecast for each hurricane event in our sample, we are also able to provide an out-of-sample metric of the forecast accuracy of the model. Averaging across all of the events, the *out-of-sample* R-squared is 0.75.

To provide an assessment on the possibility that alternative explanations could have generated our results, we also conducted a series of robustness checks. To ensure our results are not contaminated by changes in state-level policies around the hurricane event, we ran our main regression specification under a more restrictive window of +/- 3 months around the hurricane event (see, table S1 in the supplementary materials). To explore the possibility that there might be some delay in how search results in UI claims, we also ran a version of the model with a lag of Google Trends intensity (see, table S2 in the supplementary materials). To control for the possibility that overall UI generosity

[12]For a scatter plot of UI claims share and Google Trends ratio net of event fixed-effects, see figure S2 in the supplementary materials.

might be driving the results, we also conducted our main regression approach using a matched control state (see, table S3 in the supplementary materials) as the reference group. Finally, we also ran a "placebo" test by assigning each hurricane event to an alternative state amongst our treated states that was unlikely to be impacted by the hurricane (e.g., New Jersey for Harvey; see, table S4 in the supplementary materials), in addition to re-running our main specification with population weights (see, table S5 in the supplementary materials).

In each case, we find only modest differences with our baseline. In particular, restricting to a shorter event window (1.23) and using population weights (1.01) each lead to modest changes in the point estimate, while using a state matched by UI generosity modestly lowers the point estimates across the range of leave-one-out exercises [0.72, 1.06]. And, while there appears to be some scope for delayed effects in how search leads to unemployment insurance claims, the implied overall impact from this specification (i.e., the sum of coefficients) indicates a modest increase in the overall elasticity. Finally, across the range of placebo tests, one cannot reject the null of no effect. Combined, these tests suggest significant predictive power for the Google Trends measure in capturing variation in initial UI claims in the states affected by the landfall of the hurricanes in our sample.

Given the appealing features of our regression model in predicting initial UI claims over our set of seven major hurricane events, we next report what our estimates would imply for the number of UI claims during the early days of the Covid-19 pandemic. With several states implementing shelter-in-place policies of their residents and many businesses shutting down in March 2020, the resulting rise in UI claims was expected to be large–but how large was not clear. To this point, private sector forecasts for initial UI claims in the early days of the pandemic using a variety of methods (including aggregating recent news report accounts) ranged from as low as 1 to as much as 4 million for the week ending March 21, 2020 (e.g., Yglesias (2020)).

Figure 3 reports our model estimates of initial UI claims on a non-seasonally adjusted basis for the weeks ending March 21, 2020 and March 28, 2020–which were subsequently reported at 8:30am (EST) on March 26, 2020 and April 2, 2020. While considerable variation exists across the individual states, across-the-board the model forecasts substantial increases in initial UI claims that would make it the largest single and two-week increases in US history. The model's forecasts (of log-changes) are also well calibrated in both weeks.[13] In both weeks, 80% of states fall within their 95% confidence prediction interval, with several of the states falling outside in both instances, but in opposite directions (e.g., California)–a further indication of some delay in the resulting claims from search (see, table S2 in the supplementary materials).

[13]In a cross sectional regression of log actual changes in UI claims on the predicted log change, we cannot reject the joint null hypothesis that the constant is zero and the coefficient on the predicted changes is one at the 1% significance level, for either week.

Covid Economics 9, 24 April 2020: 62-79

Covid Economics 9, 24 April 2020: 62-79

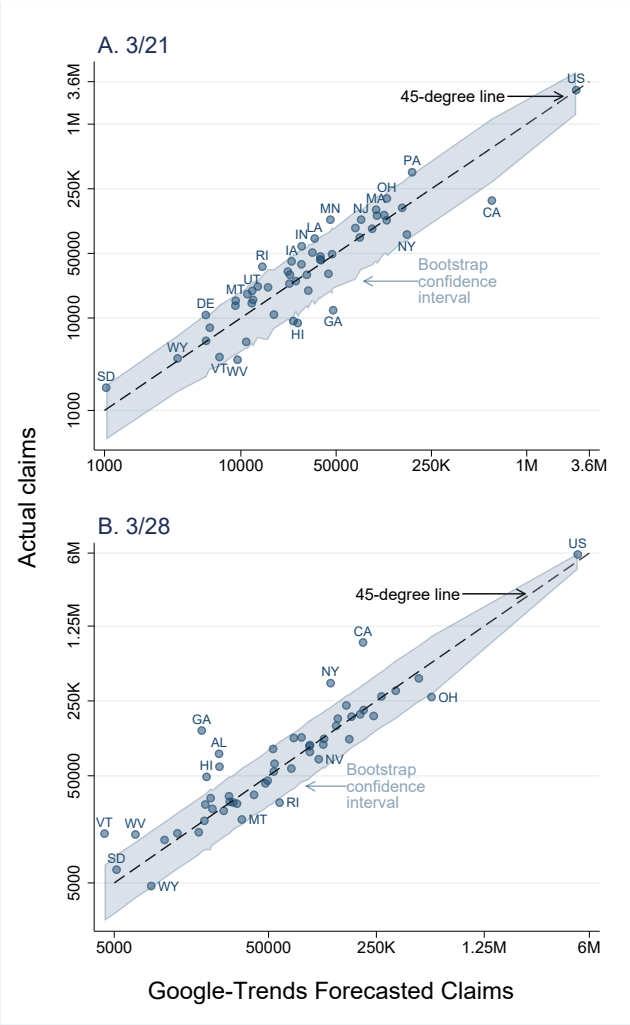

Figure 3 State and National Initial Unemployment Insurance Claims

Notes: This figure reports the forecasted values for each state from our model using Google Trends, as well as the reported initial unemployment insurance claims (final) number for the week ending March 21, 2020 (panel A) and the (advance) number for the week ending March 28, 2020 (panel B). The interpolated 95% confidence interval is also reported based on 10,000 event-level bootstrap and jackknife samples.

Using our point estimate from the baseline regression, we forecasted initial UI claims for the US as a whole to be 2.9 million (panel A), with a 95% prediction interval of 1.0 to 4.6 million, for the week ending March 21, 2020. The advance (and ultimately, final) release of initial UI claims was 2.9 million for the week ending March 21, falling on top of our estimates. Our projection for the week of March 28, 2020 predicted that new UI claims for the US as a whole would be between 4.2 and 6.0 million, with a mean prediction of 5.0 million (panel B).[14] The advance report of UI Claims for the week ending March 28, 2020 came in at 5.8 million. It is important to note that these forecasts represent an out-of-sample prediction, as opposed to some of the other recent approaches that have used early reports from individual states as the basis for the aggregate forecast (e.g., Goldsmith-Pinkham and Sojourner (2020)).[15]

To obtain an estimate of the broader labor market effects of the pandemic, we take our initial UI claims predictions for the last two weeks of March and combine them with the results for the first two weeks of March to arrive at a March estimate of monthly initial UI claims. Then, we "plug" this estimate into the real-time data flow as of March 30, 2020 for an extended version of the Brave et al. (2019) mixed-frequency BVAR (MF-BVAR) model of US economic activity.[16] The historic surge in new UI claims in March 2020 predicted by our model is interpreted as a highly negative shock to economic activity in the MF-BVAR. Figure 4 compares the model forecasts with versus without our UI claims plug; the level of GDP is 13% lower (blue bar), payrolls are 11% lower (orange bar), and the unemployment rate is 7.6 percentage points higher (gray bar) by the first quarter of 2021 when our plug is included. All three results represent substantial informational gains not found in the existing real-time data flow.

5. Conclusion

The Covid-19 pandemic is likely to test the limits of usefulness of traditional economic data and statistical models. The sheer magnitude of the shock to labor market activity is unprecedented in recent history. However, we have shown that recently developed high frequency "big data" sources like Google Trends can be useful predictors of labor market conditions even in such a very uncertain environment. Furthermore, while other methods used to-date to predict initial UI claims with this data in the current environment have relied primarily on the early reporting of results by states, our

[14] To obtain a prediction for the week of March 28, 2020, we used the two weeks of Google Trends data following March 7, 2020 along with that week's initial unemployment insurance (UI) claims number.

[15] A version of the Choi and Varian (2012) model estimated on data from 2014–2019 predicts UI claims to be 267 thousand and 1.7 million (non-seasonally adjusted basis) for the weeks ending March 21, 2020 and March 28, 2020, respectively. In contrast, the additional information in the early state reports led the Goldsmith-Pinkham and Sojourner (2020)'s model to predict UI claims to be 3.8 and 4.7 million (non-seasonally adjusted) for the same weeks.

[16] This version of the MF-BVAR contains 107 monthly and quarterly time series for US real economic activity, many of which are used by the Bureau of Economic Analysis as source data for US gross domestic product.

Covid Economics 9, 24 April 2020: 62-79

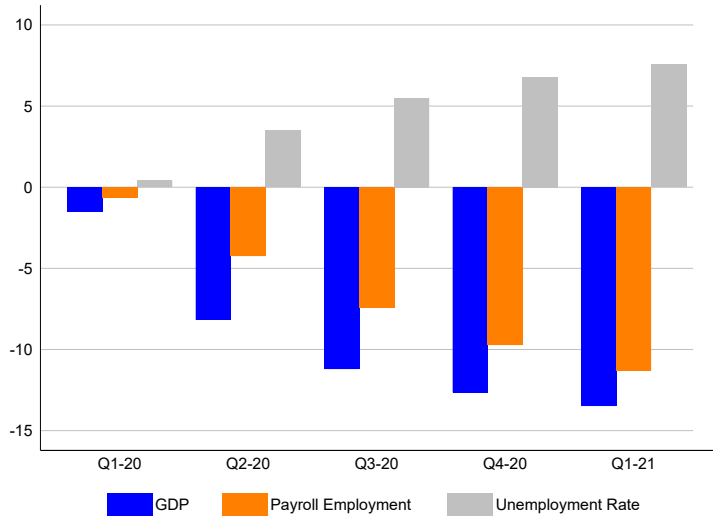

Figure 4 Impact of March 2020 UI Claims on Projections of Economic Activity

Notes: This figure displays the percentage difference (or percentage point difference for the unemployment rate) in each quarter between the projections for the level of US real GDP, payroll employment, and the unemployment rate with and without our estimates of March 2020 initial UI claims using data available as of March 30, 2020 from a 107-variable version of the Brave et al. (2019) mixed-frequency BVAR mode of US economic activity.

method is robust to the possibility that such reporting may at some point no longer be possible or feasible. This is made possible by the precision of our estimates afforded by studying past responses of UI claims and Google searches to hurricane landfalls in the US.

The pattern of initial UI claims and Google Trends searches for unemployment for states affected by the landfall of hurricanes appears to share very similar characteristics to the labor market response to the current pandemic.[17] Hurricanes, however, are by their very nature localized events whose impacts are limited based on a state's exposure to coastal areas. Covid-19 does not respect geographical boundaries. One potential risk of using our approach in the current environment is if the Google search behavior of users is radically different in local and national natural disasters. Furthermore, anecdotal reports suggest that the magnitude of the shock may also be impacting a person's actual ability to file a UI claim as well as his or her state's ability to process it, which could alter the timing of the relationship we identify.

[17]In other analysis examining UI claims, Bram and Deitz (2020) reach a similar conclusion.

Covid Economics 9, 24 April 2020: 62-79

References

Aaronson, D., S. Brave, and R. Cole (2016). Using private sector 'big data' as an economic indicator: The case of construction spending. *Chicago Fed Letter 366*.

Askitas, N. and K. F. Zimmermann (2009). Google econometrics and unemployment forecasting. *Applied Economics Quarterly 55*(2), 107–120.

Bram, J. and R. Deitz (April 10, 2020). The coronavirus crisis shock looks more like a natural disaster than a cyclical downturn. *Liberty Street Economics*.

Brave, S. A., R. A. Butters, and A. Justiniano (2019). Forecasting economic activity with mixed frequency BVARs. *International Journal of Forecasting 35*(4), 1692–1707.

Cajner, T., L. Crane, R. Decker, A. Hamins-Puertolas, and C. J. Kurz (2019, September). Tracking the labor market with 'big data'. FEDS Notes 2019-09-20, Board of Governors of the Federal Reserve System (US).

Choi, H. and H. Varian (2012). Predicting the present with Google Trends. *Economic Record 88*(s1), 2–9.

D'Amuri, F. and J. Marcucci (2017). The predictive power of Google searches in forecasting US unemployment. *International Journal of Forecasting 33*(4), 801–816.

Deryugina, T. (2017). The fiscal cost of hurricanes: Disaster aid versus social insurance. *American Economic Journal: Economic Policy 9*(3), 168–198.

Deryugina, T., L. Kawano, and S. Levitt (2018, April). The economic impact of Hurricane Katrina on its victims: Evidence from individual tax returns. *American Economic Journal: Applied Economics 10*(2), 202–33.

Gallagher, J. and D. Hartley (2017). Household finance after a natural disaster: The case of Hurricane Katrina. *American Economic Journal: Economic Policy 9*(3), 199–228.

Giannone, D., L. Reichlin, and D. Small (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics 55*(4), 665 – 676.

Goldsmith-Pinkham, P. and A. Sojourner (2020). The coronavirus crisis led to a record-breaking spike in weekly unemployment insurance claims. *Economic Policy Institute*. `https://www.epi.org/blog/coronavirus-record-breaking-spike-in-ui-claims/`.

Gordon, R. J. (2009). Green shoot or dead twig: Can unemployment claims predict the end of the American recession? *Vox*. `https://voxeu.org/article/us-recovery-may-2009-new-evidence-based-surprisingly-robust-linkage`.

Hsu, J. W., D. A. Matsa, and B. T. Melzer (2018, January). Unemployment insurance as a housing market stabilizer. *American Economic Review 108*(1), 49–81.

McLaren, N. and R. Shanbhogue (2011). Using internet search data as economic indicators. *Bank of England Quarterly Bulletin 51*(2), 134–140.

NOAA (2020). US billion-dollar weather and climate disasters. Technical report, National Centers for Environmental Information (NCEI).

Ortega, F. and S. Taşpinar (2018). Rising sea levels and sinking property values: Hurricane Sandy and New York's housing market. *Journal of Urban Economics 106*, 81–100.

Rigobon, R. (2003). Identification through heteroskedasticity. *The Review of Economics and Statistics 85*(4), 777–792.

Suhoy, T. (2009, July). Query Indices and a 2008 Downturn: Israeli Data. Bank of Israel Working Papers 2009.06, Bank of Israel.

Tuhkuri, J. (2016, March). Forecasting Unemployment with Google Searches. ETLA Working Papers 35, The Research Institute of the Finnish Economy.

Covid Economics 9, 24 April 2020: 62-79

Yglesias, M. (March 23, 2020). The coming unemployment catastrophe, in one chart. *Vox*. https://www.vox.com/2020/3/21/21188529/initial-unemployment-claims-goldman-sachs.

Covid Economics 9, 24 April 2020: 62-79

Supplementary Tables and Figures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Google Trends Ratio	1.232	1.207	1.191	1.209	1.202	1.283	1.219	1.295
	(0.192)	(0.289)	(0.168)	(0.215)	(0.346)	(0.268)	(0.197)	(0.070)
Left Out Event		Katrina	Harvey	Sandy	Irma	Wilma	Ike	Rita
Adj. R-squared	0.545	0.371	0.523	0.550	0.526	0.580	0.565	0.604
Num. Observations	203	174	174	174	174	174	174	174

Table S1 Three Month Window Event Study Regression Results

Notes: This table reports results for equation 1 using a +/-3 month window around the hurricane event. We report the coefficient on the Google Trends Ratio term and its standard error, estimated by 10,000 event-level bootstrap and jackknife samples, in parentheses. In column 1, we report estimates using all of the hurricane events; and in each subsequent column, we report results from leaving each of the seven hurricane events out individually.

Covid Economics 9, 24 April 2020: 62-79

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Google Trends Ratio	0.793	0.767	0.747	0.783	0.768	0.847	0.787	0.833
	(0.116)	(0.306)	(0.051)	(0.060)	(0.179)	(0.166)	(0.143)	(0.190)
Google Trends Ratio (t-1)	0.639	0.492	0.659	0.622	0.659	0.637	0.635	0.669
	(0.096)	(0.096)	(0.059)	(0.091)	(0.105)	(0.137)	(0.101)	(0.122)
Left Out Event		Katrina	Harvey	Sandy	Irma	Wilma	Ike	Rita
Sum of Coef.	1.43	1.26	1.41	1.40	1.43	1.48	1.42	1.50
SE	0.20	0.30	0.11	0.06	0.22	0.24	0.10	0.09
Adj. R-squared	0.61	0.38	0.60	0.62	0.61	0.64	0.63	0.67
Num. Observations	359	308	308	307	308	308	307	308

Table S2 Allowing for Delayed Claims in Event Study Regression Results

Notes: This table reports results from an alternative to our baseline specification where an additional lag of the Google Trends Ratio is included in the model. We report the coefficient on the contemporaneous and lagged value of the Google Trends Ratio term and their standard errors, estimated by 10,000 event-level bootstrap and jackknife samples, in parentheses. In column 1, we report estimates using all of the hurricane events; and in each subsequent column, we report results from leaving each of the seven hurricane events out individually. The sum of the lag coefficients and its associated standard error is reported for each specification as well.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Google Trends Ratio	0.915	0.720	0.862	0.904	0.903	1.066	0.902	0.972
	(0.204)	(0.339)	(0.118)	(0.196)	(0.323)	(0.257)	(0.174)	(0.218)
Left Out Event		Katrina	Harvey	Sandy	Irma	Wilma	Ike	Rita
Adj. R-squared	0.375	0.213	0.355	0.376	0.377	0.450	0.382	0.418
Num. Observations	366	314	314	313	314	314	313	314

Table S3 Alternative Control State Event Study Regression Results

Notes: This table reports results from an alternative to the baseline model that uses a control state matched by unemployment insurance generosity as the reference category for each hurricane event. We report the coefficient on the Google Trends Ratio term and its standard error, estimated by 10,000 event-level bootstrap and jackknife samples, in parentheses. In column 1, we report estimates using all of the hurricane events; and in each subsequent column, we report results from leaving each of the seven hurricane events out individually.

Covid Economics 9, 24 April 2020: 62-79

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Google Trends Ratio	0.075	0.117	0.059	0.037	0.055	0.105	0.090	0.084
	(0.086)	(0.061)	(0.026)	(0.052)	(0.031)	(0.147)	(0.098)	(0.104)
Left Out Event		Katrina	Harvey	Sandy	Irma	Wilma	Ike	Rita
Adj. R-squared	-0.014	-0.009	-0.016	-0.018	-0.016	-0.011	-0.011	-0.013
Num. Observations	366	314	314	313	314	314	313	314

Table S4 Placebo Event Study Regression Results

Notes: This table reports results from an alternative to the baseline model that uses a "placebo" state as the primary treated state (see table 1). We report the coefficient on the Google Trends Ratio term and its standard error, estimated by 10,000 event-level bootstrap and jackknife samples, in parentheses. In column 1, we report estimates using all of the hurricane events; and in each subsequent column, we report results from leaving each of the seven hurricane events out individually.

Covid Economics 9, 24 April 2020: 62-79

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Google Trends Ratio	1.036	0.949	0.898	1.024	0.982	1.190	1.026	1.221
	(0.261)	(0.354)	(0.267)	(0.263)	(0.290)	(0.247)	(0.264)	(0.224)
Left Out Event		Katrina	Harvey	Sandy	Irma	Wilma	Ike	Rita
Adj. R-squared	0.39	0.31	0.32	0.40	0.35	0.48	0.42	0.51
Num. Observations	366	314	314	313	314	314	313	314

Table S5 Population Weighted Event Study Regression Results

Notes: This table reports results from a weighted version of the baseline model where weights are based on population. We report the coefficient on the Google Trends Ratio term and its standard error clustered at the event in parentheses. In column 1, we report estimates using all of the hurricane events; and in each subsequent column, we report the results from leaving each of the seven hurricane events out individually.

Covid Economics 9, 24 April 2020: 62-79

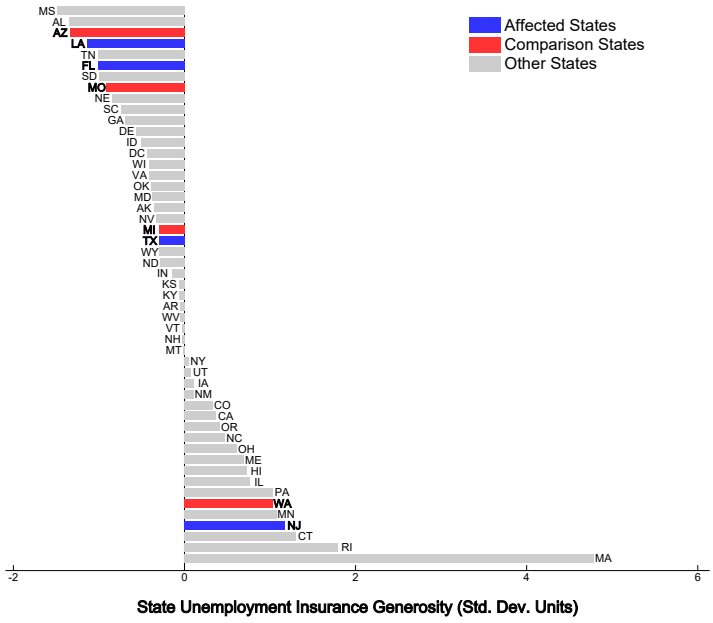

Figure S1 Matching on State Unemployment Insurance Generosity

Notes: This figure displays the 2004-2010 sample average of the Hsu et al. (2018) maximum benefit variable used to measure unemployment insurance generosity by state, measured in standard deviation units from the cross-sectional average. The states affected by hurricanes in our analysis are shaded in blue, and the matched control states are shaded in red. See table 1 for further details.

Covid Economics 9, 24 April 2020: 62-79

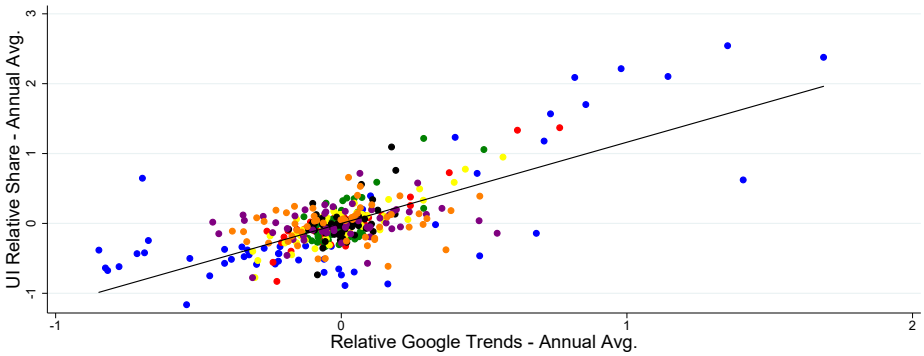

Figure S2 UI Claim Shares and Relative Google Trends

Notes: This figure displays the log share of initial unemployment insurance claims versus the log ratio of Google Trends for the "unemployment" topic for the seven hurricane-affected states relative to the US: Katrina (Louisiana, in blue), Harvey (Texas, in red), Sandy (New Jersey, in green), Irma (Florida, in yellow), Ike (Texas, in black), Wilma (Florida, in purple), and Rita (Texas, in orange).

Covid Economics **Issue 9, 24 April 2020**

Air passenger mobility, travel restrictions, and the transmission of the covid-19 pandemic between countries

Sekou Keita[1]

Date submitted: 17 April 2020; **Date accepted:** 17 April 2020

In this study, I discuss the role of international air traffic in spreading the new coronavirus COVID-19 around the world, with a focus on travel restrictions. I build on a sample of 34 mostly European countries reporting international flights to 154 destination countries. This dataset is combined with information on daily reported cases of COVID-19 infections in these countries. I find that more connected countries registered first infection cases significantly earlier than less connected countries. This effect was reinforced by direct flight connections to China. I also show that severe travel restrictions were implemented relatively late in most countries. For a group of 120 countries included in the sample of analysis, three out of four countries already had more than 50 confirmed cases when travel restrictions were implemented. In contrast, very early implementations of air travel restrictions were associated with a delayed onset of infections. As a takeaway for future outbreaks of infectious diseases, the results suggests that the early implementation of travel restrictions could be key in slowing down the spread of infections around the world. The design of a global emergency stop in international travel requires a high level of coordination at a multilateral level in order to preserve supply chains as much as possible.

1 Research Associate, Institute for Employment Research (IAB). I am grateful to Andrew Fallone, Andreas Hauptmann, and Ignat Stepanok for helpful comments. The usual disclaimers apply.

Covid Economics 9, 24 April 2020: 80-99

Covid Economics 9, 24 April 2020: 80-99

1 Introduction

In recent years, many studies have underlined the fact that increasing mobility of people due to globalization could accelerate the spread of new contagious diseases around the world (Hufnagel *et al.*, 2004; Hsu and Shih, 2010; Brockmann and Helbing, 2013: for example). In particular, previous studies have shown that air travel played a key role in spreading recent pandemics such as H1N1. As the number of persons circulating between countries has grown steadily with the progress of globalization, the potential for new contagious diseases to reach all countries of the world was already very high when the new virus COVID-19 was discovered in China in late 2019.

The new corona virus COVID-19 has largely validated these concerns by spreading at an unprecedented speed, reaching over 200 countries and territories within three months of the first documented human infection. Although it eventually reached almost all countries, the country-specific timing of the onset of the pandemic could play an important role in the ability of countries to cope with the crisis. A delayed start of growing number of cases may provide authorities with precious time to prepare for the crisis management. A later start of the pandemic in a country also increases the number of experiences in other countries to learn from.

The goal of this study is to discuss the potential role of travel restrictions in slowing down contagion around the world. I first verify that air traffic was indeed associated with spreading the new corona virus COVID-19 around the world. I use air traffic data from Eurostat between 34 reporting countries and 154 destination countries from July 2018 to June 2019 to measure the intensity of movement of people between two countries. I combine this dataset with information on daily reported cases of COVID-19 infections in 200 countries. I complement the dataset with information on the implementation of travel restrictions between pairs of countries, collected from multiple sources.

The number of air traffic passengers is particularly well suited to capture the intensity connections between countries because all aspects of globalization are facilitated by air traffic. The observation that air traffic brings geographically distant places closer together refers to the notion of effective distance (Brockmann and Helbing, 2013). For instance, trade and investments activities with geographically distant countries are eased by the potential for persons to quickly travel to meet business partners or customers. Likewise, the supply of international air traffic greases the wheel of international labor migration by allowing migrants to easily visit their country of origin when necessary. Finally, the industry of leisure consumption in foreign countries, in the form of

international tourism, has also greatly benefited from possibilities created by increasing air traffic supply.

The mechanism of transmission is straightforward: air travel enables the mobility of infected persons between countries within hours, greatly reducing the relevance of geographic distance. Moreover, the role of air traffic in accelerating the spread of infectious diseases is reinforced by the fact that passengers who are helthy at the time of boarding can get infected simply by travelling on the same plane as an infected person (Foxwell *et al.*, 2011). Several studies have already discussed the role of air traffic in spreading COVID-19 to or from specific locations, including China (Lau *et al.*, 2020), Brazil (da Silva Candido *et al.*, 2020), Pacific Islands (Craig *et al.*, 2020), African countries (Gilbert *et al.*, 2020), and Iran (Zhuang *et al.*, 2020).

Covid Economics 9, 24 April 2020: 80-99

Using data on mobility before the crisis ensures that the measure of mobility between countries is unaffected by travel restrictions implemented in the first months of 2020 as a consequence of the pandemic. The rationale is that a higher level of mobility of people between a country and the rest of the world is associated with a higher risk of importing the pandemic. Once some cases of COVID-19 infections have already been observed in a country, the virus can spread within the country without further contact with the rest of the world. For this reason, the focus of this analysis is on the very first cases of infections in a country.

The results of the analysis indicate that an additional 100,000 passengers per year in direct flight connections to China is associated with an acceleration of the first reported case of a COVID-19 infection in a country by 1.5 days. The results also show that the onset of the pandemic was closer between highly connected countries. An additional 10,000 direct flight connections per year between two countries is associated with a reduction in the gap in the number of days of the first reported case by 2 days. These effects are large when considering the fact that highly connected countries like Germany reported an average of more than 2,500 flights and 300,000 international air traffic passengers per week between July 2018 and June 2019, spread over more than 100 destination countries. Taken together, the results confirm the key role of the mobility of people in spreading the COVID-19 virus. Regarding the role of travel restrictions, I find that the vast majority of countries had already registered at least several dozen COVID-19 cases when travel restrictions were implemented. However, early implementations of travel restrictions were associated with delayed country-specific outbreaks of the pandemic.

Taken together, the results call for caution when discussions will begin to lift the current travel restrictions. As a lesson for future outbreaks of new infectious diseases, the results suggest that early severe travel restrictions could substantially slow down the spread of infections around the world. A global emergency stop in international travel should be designed in a way to minimize

effects on global supply chains too severely.

The remainder of the study is organized as follows: Section 2 presents the data and Section 3 presents the correlations between air travel connections and the country-specific onset timing of the COVID-19 pandemic. The role of travel restrictions implemented by countries around the world is discussed in Section 4. Finally, Section 5 draws some concluding remarks.

2 Data

I combine data from several sources for this analysis. First, I obtain information on air traffic between 34 reporting countries and 154 destination countries from Eurostat. Information on flights is reported at the airport level, but I aggregate it at the country level. In order to obtain a measure of connection intensity over a 12 month period, I keep observations from July 2018 to June 2019. The reason is that observations are not available for all countries after this date. I use this dataset to build two indicators: the number of flight connections between countries, and the number of passengers. While a round trip is counted as one flight connection, the number of passengers is the sum of passengers registered at departure and at arrival. Table 1 shows the number of destinations of direct flights, the number of flights, and the number of passengers carried for each of the 34 reporting countries.

Next, I obtain information on daily reported cases of COVID-19 infections for 200 countries up to March 31st, 2020, from the EU open data portal. For each country, I identify three dates corresponding to different stages of the onset of the pandemic: the day of the first reported case (D_{first}), the day of start of continuous growth in the number of reported cases (D_{start}), and the day the country reached 100 cases or more (D_{100}). The start of continuous growth, or country-specific pandemic starting day, is defined as the first day with a positive reported case followed by positive reported cases in the next three days. I then calculate the number of days between each of these dates and January 1st, 2020, to measure differences in the timing of the onset of the pandemic in a country ($D_c - January1st$, where $c \in \{first, start, 100\}$). I calculate the gap in the timing of the onset of the pandemic for pairs of countries as the absolute value in the difference in dates ($\mid Diff_{ij} \mid = D_{ci} - D_{cj}$, where $c \in \{first, start, 100\}$.

A drawback in data quality on COVID-19 infections is heterogeneity in testing policies. Indeed, the number of available tests varied greatly between countries. While some countries, such as South Korea, adopted a large-scale testing policy at an very early stage, others, like France, limited tests to persons with relevant symptoms. Such discrepancies in national testing rates could influence

Covid Economics 9, 24 April 2020: 80-99

the total number of confirmed cases reported in each country. Unfortunately, to the best of my knowledge, data on testing policies and number of performed tests is not yet available for most countries. Measurement errors due to varying testing policies across countries are more severe in country-level analyses, as in Section 3.1, than in bilateral analyses. The reason is that regressions based bilateral on bilateral data include country fixed effects as shown in equation 2. The country fixed effects absorb systematic differences in country-specific testing policies, which do not vary across partner countries.

Information on the implementation of travel restrictions is self-collected from multiple sources, included reports in newspapers such as the New York Times and Al Jazeera, and the International Air Transport Association (IATA). For each measure of timing of the onset of the pandemic, I generate a binary variable equal to one for a pair of countries when a travel restriction was in place between the countries on the date of first case, growth start, or 100 cases, and zero otherwise.

Finally, I obtain data on geographic proximity and the existence of common borders between countries from CEPII (Mayer and Zignago, 2011). I measure the bilateral distance between countries as the distance between the most populated cities, in multiples of 1,000 km.

3 Air travel connections and the COVID-19 pandemic

In this Section, I assess the affect of bilateral air travel connections on the country-specific timing of the onset of the crisis. The goal is not to predict the exact date of the outbreak of the pandemic in each country. Rather, the idea is to verify whether the strong statistical relationship between air travel and the spread of contagious diseases, which has been shown for previous pandemics such as H1N1, can also be observed for COVID-19.

3.1 Connections to China

I first assess whether more direct air traffic connections with China are associated with the speed of the onset of the pandemic in reporting countries, measured in number of days with respect to January 1st, 2020. Figure 1 suggests that, on average, countries with stronger air travel connections with China did indeed experience an earlier outbreak of the COVID-19 pandemic. However, Figure 1 also shows that the new COVID-19 virus also reached countries with no direct flight connections to China relatively quickly, indicating that effective distance to China was relatively short. In the strongly connected international travel system prevailing before the crisis, any country was actually at most one transit away from China.

Covid Economics 9, 24 April 2020: 80-99

CEPR PRESS

Covid Economics 9, 24 April 2020: 80-99

Table 1: Number of destinations and international air travel intensity, by reporting country

Reporting country	Number of destinations	Total flights	Total passengers
Austria	60	24.039	28.806
Belgium	53	19.990	27.043
Bulgaria	33	5.774	8.403
Croatia	32	6.250	7.330
Cyprus	34	5.627	8.855
Czech republic	40	13.591	17.259
Denmark	44	23.245	28.962
Estonia	24	3.298	2.666
Finland	37	13.737	17.276
France	115	87.502	127.992
Germany	103	138.951	191.376
Greece	41	20.165	32.199
Hungary	38	8.586	12.448
Iceland	15	4.168	7.108
Ireland	30	22.942	32.903
Italy	79	65.795	93.281
Latvia	26	5.759	5.097
Lithuania	26	4.352	5.108
Luxembourg	22	4.422	3.437
Malta	19	2.967	4.340
Montenegro	17	1.721	1.801
Netherlands	85	50.920	75.869
North Macedonia	22	1.597	2.194
Norway	32	16.293	20.323
Poland	43	23.226	29.992
Portugal	36	26.802	39.830
Romania	29	11.275	15.678
Slovakia	29	1.699	2.423
Slovenia	23	2.166	1.611
Spain	67	95.235	154.130
Sweden	44	19.142	24.148
Switzerland	61	37.847	48.239
Turkey	89	27.661	47.698
United Kingdom	101	139.980	225.102

Source: Authors' elaboration on data from Eurostat. The number of flight passengers is expressed in multiples of 100,000. The number of flights is expressed in multiples of 10,000.

Covid Economics 9, 24 April 2020: 80-99

Figure 1: Connections to China and timing of country-specific pandemic start

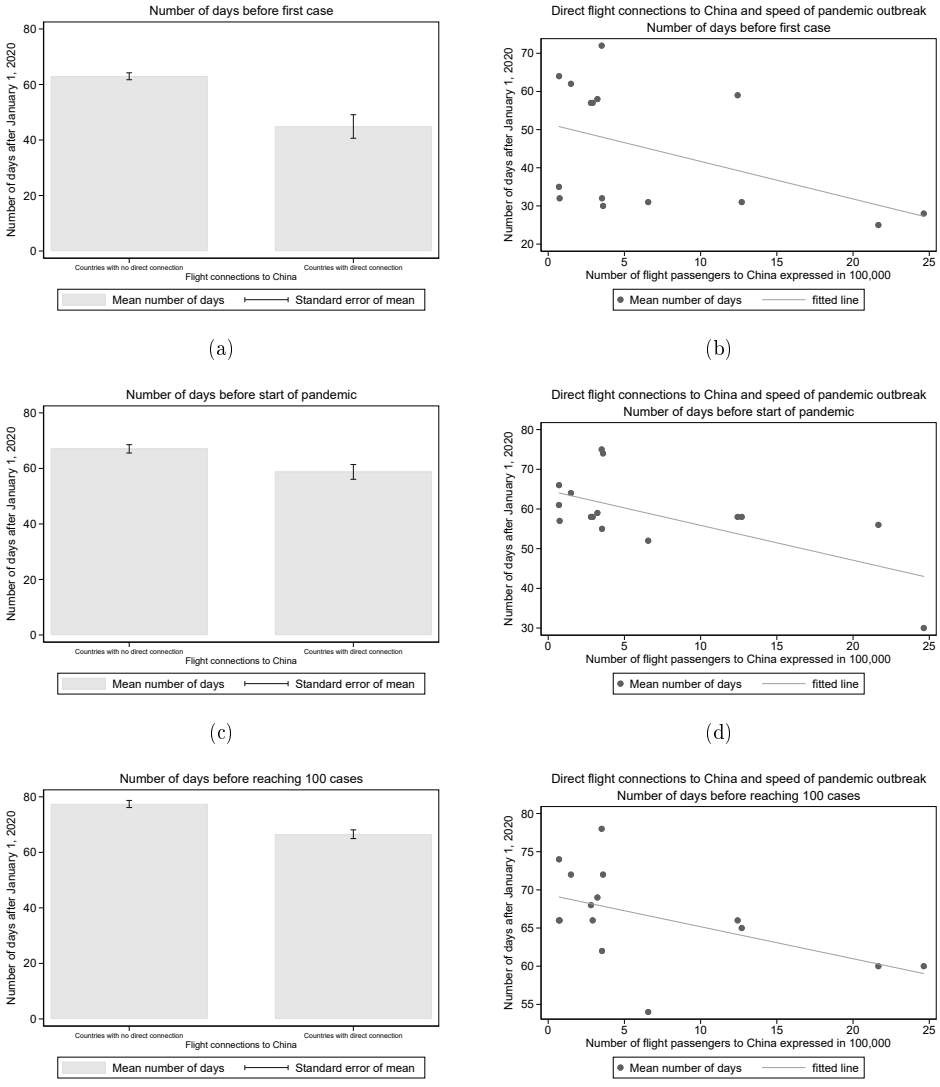

Source: Authors' elaboration on data from EU open data portal and self-collected information on traffic restriction implementations in countries. The country-specific pandemic starting day is defined as the first day with a positive reported case followed by positive reported cases in the next three days.

I estimate the following equation:

$$Day_i = \alpha + \beta_1 connect_i + \beta_2 dist_i + \epsilon_i \quad (1)$$

where Day_i is the number of days elapsed between January 1st, 2020, and the onset of the COVID-19 pandemic in country i. I use one of three dates to capture the onset of the pandemic in a country: the day of the first reported case, the day of start of continuous growth in the number of reported cases, or the day the country reached 100 cases or more. The variable of interest is $connect_i$, measuring either the number of flights between country i and China, or the number of flight passengers between country i and China. Geographical distance between country i and China is captured by $dist_i$ and ϵ_i is an error term.

The results of estimating equation 1 with ordinary least squares is shown in Table 2. Column (1) shows that, after controlling for geographic distance, an additional 100,000 passengers per year in direct flight connections with China is associated with an acceleration of the first reported case of a COVID-19 infection in a country by 1.49 days. The results in column (4) indicate that an additional 1,000 direct flights to China per year is associated with a 3.8-day advancement of the country-specific outbreak date. These correlation are in line with findings in previous studies showing that air travel played a key role in spreading infectious diseases around the world.

Table 2: Air traffic intensity with China and timing of country-specific pandemic outbreak

Dependent variable	Days until first case	Days until start	Days until 100 cases	Days until first case	Days until start	Days until 100 cases
	(1)	(2)	(3)	(4)	(5)	(6)
Air traffic with China (Passengers)	-1.489***	-0.913***	-0.732***			
	(0.244)	(0.328)	(0.118)			
Air traffic with China (Flights)				-3.895***	-2.372***	-1.922***
				(0.591)	(0.804)	(0.315)
Distance between most populated cities	-1.210	-2.844**	-2.404*	-0.967	-2.708**	-2.279
	(3.466)	(1.265)	(1.389)	(3.407)	(1.242)	(1.349)
Observations	34	34	34	34	34	34
R^2	0.399	0.465	0.417	0.421	0.481	0.437

Source: Authors' elaboration on data from Eurostat, CEPII, and EU open data portal. The number of flight passengers is expressed in multiples of 100,000. The number of flights is expressed in multiples of 1,000. Geographic distance is expressed in multiples of 1,000 km. Air traffic intensity is measured as the total number of flights passengers between a country and China from July 2018 to June 2019. The country-specific pandemic starting day is defined as the first day with a positive reported case followed by positive reported cases in the next three days. Robust standard errors are between parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Covid Economics 9, 24 April 2020: 80-99

3.2 Bilateral connections

In this Section, I assess the effect of air traffic connections on the gap in the onset of the pandemic between countries, measured in number of days. The rationale is that, in a network of connections, two highly connected countries with more mobility of people will experience closer outbreak dates of the COVID-19 pandemic than two less connected countries. I estimate the following equation:

$$Diff_{ij} = \alpha + \beta_1^{bil} connect_{ij} + \beta_2^{bil} dist_{ij} + \beta_3^{bil} border_{ij} + \phi_i + \varphi_j + \epsilon_{ij} \quad (2)$$

where $Diff_{ij}$ is the absolute value of the number of days elapsed between the onset of the COVID-19 pandemic in countries i and j. I use one of three dates to capture the onset of the pandemic in a country: the day of the first reported case, the day of start of continuous growth in the number of reported cases, or the day the country reached 100 cases or more. The variable of interest is $connect_{ij}$, measuring either the number of flights between country i and country j, or the number of flight passengers between the two countries. Country fixed effects ϕ_i and φ_j account for country-specific characteristics such as population size, quality of the health system, performance of the virus detection technologies, economic and political situations, and distance to Wuhan, China. The binary variable $border_{ij}$ takes the value 1 if the two countries share a common border, and 0 otherwise. Geographical distance between countries is captured by $dist_{ij}$ and ϵ_{ij} is an error term.

The result of estimating equation 2 is presented in Table 3. Column (1) shows that an additional 100,000 passengers per year in direct flight connections between two countries is associated with a reduction in the gap in the number of days of the first reported case by 0.67 days. The results in column (4) indicate that an additional 10,000 direct flights between two countries per year is associated with a 2.5-day reduction in the difference of outbreak days between countries. Strong effects of bilateral air traffic are also observed when using alternative variables to measure the bilateral gap in outbreak dates. In this specification, sharing a common border does not play a statistically significant role, except in reducing the gap of the date of the first case. It is worth noting that once air traffic intensity is controlled for, geographic distance is actually positively associated with the gap in outbreak dates.

Next, I assess whether the effect of bilateral air traffic in accelerating the spread of COVID-19 around the world is intensified when one of the partners has direct connections to China. For this specification, I restrict the destination countries to the 34 reporting countries. The reason is that the dataset does not allow us to observe whether non-reporting countries, such as Canada and the United States, have direct connections to China. Since this results in a smaller sample, I first verify that the baseline results can be replicated. Table A1 in the appendix shows that this is largely the case. The results in Table 4 show that, compared to two countries with no direct connection

Table 3: Bilateral air traffic intensity and timing of country-specific pandemic outbreak

Dependent variable: bilateral difference in	Days until first case	Days until start	Days until 100 cases	Days until first case	Days until start	Days until 100 cases
	(1)	(2)	(3)	(4)	(5)	(6)
Bilateral air traffic intensity (Passengers)	-0.669***	-1.444***	-0.528***			
	(0.133)	(0.220)	(0.133)			
Bilateral air traffic intensity (Flights)				-2.073***	-0.810***	-1.110***
				(0.299)	(0.162)	(0.184)
Common border	-2.344***	-1.522	0.107	-1.942***	-2.225	-0.176
	(0.644)	(1.495)	(0.663)	(0.644)	(1.482)	(0.656)
Distance between most populated cities	0.688***	-0.645	0.834**	0.546**	-0.538	0.906***
	(0.243)	(0.594)	(0.334)	(0.244)	(0.588)	(0.329)
Observations	1,389	1,389	1,389	1,389	1,389	1,389
R^2	0.726	0.361	0.730	0.730	0.365	0.730
Reporter FE	YES	YES	YES	YES	YES	YES
Partner FE	YES	YES	YES	YES	YES	0.725

Source: Authors' elaboration on data from Eurostat, CEPII, and EU open data portal. The number of flight passengers is expressed in multiples of 100,000. The number of flights is expressed in multiples of 10,000. Geographic distance is expressed in multiples of 1,000 km. Bilateral air traffic intensity is measured as the total number of flights passengers between two countries from July 2018 to June 2019. The country-specific pandemic starting day is defined as the first day with a positive reported case followed by positive reported cases in the next three days. Robust standard errors are between parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

to China, the existence of a direct air travel connection to China for at least one of the countries significantly brings the dates of the country-specific pandemic outbreak in the two countries closer together.

4 Travel restrictions

The central role of air traffic in spreading COVID-19 infections around the world raises the question of the effectiveness of travel restrictions. Indeed, one of the first responses to the pandemic were an unprecedented number of travel ban announcements. Given the fact that the pandemic ultimately reached almost all countries, the first impression is that travel restrictions were not very successful in slowing down the importation of COVID-19 infections in new countries.

Table 5 shows that a potential explanation could be that most travel restrictions were implemented relatively late, i.e. at a moment when infection cases had already grown out of control. For instance, for a group of 120 countries included in the sample of analysis, three out of four already had more than 50 reported cases when travel restrictions were implemented. The number of undocumented cases was probably much higher at that time, greatly reducing the effectiveness of travel restrictions.[1] Figure 2 shows that most travel restrictions were implemented in the second half of

[1] Nevertheless, coordinating the lifting of travel restrictions between countries might play a crucial role when the

Covid Economics 9, 24 April 2020: 80-99

Covid Economics 9, 24 April 2020: 80-99

Table 4: Bilateral air traffic intensity and differences in the timing of country-specific pandemic outbreak, connections to China

Dependent variable: bilateral difference in	Days until first case	Days until start	Days until 100 cases	Days until first case	Days until start	Days until 100 cases
	(1)	(2)	(3)	(4)	(5)	(6)
Bilateral air traffic intensity (Flights)	-1.698	5.523*	9.855**			
	(3.516)	(3.135)	(4.111)			
A least one direct connection to China	-0.451	-6.147**	-10.663***			
× Bilateral air traffic intensity (Flights)	(3.468)	(3.120)	(4.097)			
Bilateral air traffic intensity (Passengers)				-2.395	4.747*	8.151***
				(2.140)	(2.419)	(3.152)
A least one direct connection to China				1.045	-5.093**	-8.614***
× Bilateral air traffic intensity (Passengers)				(2.100)	(2.411)	(3.143)
A least one direct connection	11.689***	4.633***	8.485***	11.848***	4.865***	8.749***
	(1.597)	(0.917)	(0.837)	(1.554)	(0.911)	(0.828)
Common border	0.043	0.983	-0.818	-0.661	0.864	-0.950
	(1.704)	(0.643)	(0.642)	(1.706)	(0.640)	(0.636)
Distance between most populated cities	0.097	1.775***	1.703***	0.481	1.948***	1.906***
	(0.871)	(0.479)	(0.372)	(0.862)	(0.471)	(0.365)
Observations	784	784	784	784	784	784
R^2	0.483	0.765	0.634	0.478	0.762	0.626

Source: Authors' elaboration on data from Eurostat, CEPII, and EU open data portal. The number of flight passengers is expressed in multiples of 100,000. The number of flights is expressed in multiples of 10,000. Geographic distance is expressed in multiples of 1,000 km. Bilateral air traffic intensity is measured as the total number of flights passengers between two countries from July 2018 to June 2019. The country-specific pandemic starting day is defined as the first day with a positive reported case followed by positive reported cases in the next three days. Robust standard errors are between parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Covid Economics 9, 24 April 2020: 80-99

March 2020, when COVID-19 infection numbers were already relatively high in most countries.

Table 5: Number of infection cases at the time of travel restriction implementation

Number of infection cases	Share in percent	Cumulated
Less than 50 cases	26.96	26.96
50 to 199 cases	26.09	53.04
200 to 999 cases	23.48	76.52
1000 cases or more	23.48	100

Source: Authors' elaboration on data on 120 countries, obtained from EU open data portal. Information on implementation of travel restrictions is self-collected from multiple sources.

Figure 2: Number of countries implementing travel restrictions, by day

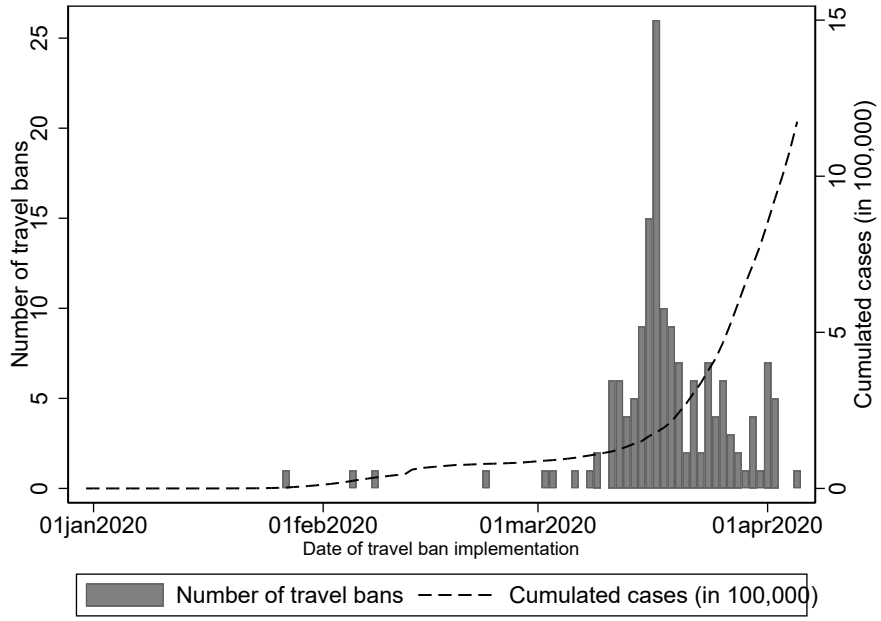

Source: Authors' elaboration on data on 120 countries, obtained from EU open data portal. Information on implementation of travel restrictions is self-collected from multiple sources.

Nevertheless, I observe that some travel restrictions were implemented very early. I ask whether

number of COVID-19 infection cases will drop in some countries while they are still rising in other countries.

these restrictions were effective in slowing down the spread of COVID-19 infections. The results in Table 6 suggest that early travel restrictions were associated with a delayed importation of the pandemic.[2] Column (1) indicates that, for a pair of countries, a bilateral travel restriction in place before the detection of the first case in either country is associated with a 6-day gap in the onset dates of the pandemic between the two countries.

Table 6: Travel restriction implementation and differences in the timing of country-specific pandemic outbreak

Dependent variable: bilateral difference in	Days until first case	Days until start	Days until 100 cases
	(1)	(2)	(3)
Travel restrictions in place before event date	6.147***	10.127	8.596***
	(0.760)	(9.589)	(1.140)
Common border	-2.317***	-2.685*	-0.197
	(0.634)	(1.557)	(0.646)
Distance between most populated cities	1.032***	0.647	1.116***
	(0.231)	(0.604)	(0.324)
Observations	1,389	1,389	1,389
R^2	0.725	0.329	0.732
Reporter FE	✓	✓	✓
Partner FE	✓	✓	✓

Source: Authors' elaboration on data from Eurostat, CEPII, and EU open data portal. The number of flight passengers is expressed in multiples of 100,000. The number of flights is expressed in multiples of 10,000. Geographic distance is expressed in multiples of 1,000 km. Bilateral air traffic intensity is measured as the total number of flights passengers between two countries from July 2018 to June 2019. The country-specific pandemic starting day is defined as the first day with a positive reported case followed by positive reported cases in the next three days. Travel restrictions are defined as a dummy variable equal to one when a travel restriction was implemented between the countries at the date of the event defined in outcome variables (first case, growth start, or 100 cases), and zero otherwise. Robust standard errors are between parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

I also explore the heterogeneity in the effectiveness of early implementation of travel restrictions. The results in Table 7 show that quickly implemented bilateral travel restrictions were more effective in delaying the onset of the COVID-19 pandemic between the two countries when bilateral travel

[2] For information, Table A2 in the appendix shows the results when travel restrictionns and air traffic intensity are included simultaneously in the regression.

Covid Economics 9, 24 April 2020: 80-99

intensity was higher.

The case of Slovakia

Slovakia offers an interesting case study when considering the role of early air travel shut downs in the spread of the COVID-19 pandemic. The country announced on March 12, 2020, that Bratislava and two other Slovakian international airports – Kosice and Poprad – were closed to all operations for at least two weeks, in response to the COVID-19 pandemic. On the same day, the country had registered 10 COVID-19 cases. Figure 3 shows the evolution in the number of cases in Slovakia along with five other countries that had comparable number of cases (between 5 and 15) on March 12, 2020. South Africa, Hungary, Serbia, Colombia, and Luxembourg all implemented travel restrictions later that Slovakia.[3] The evolution in the growth rate of number of cases suggests that the early shut down of Slovakia might have made a difference in the total number of cases. However, this interpretation has to be considered with caution given that Slovakia has a relatively small population (5.4 million people in 2018) and the country had particularly low levels of international air travel compared to the other countries before the crisis (see Table 1). Further research is needed to better understand the effectiveness of air traffic restrictions in slowing down the spread of pandemics, with a focus on the timing of restrictions' implementation.

[3]On March 11, 2020, Hungary banned inbound travel from Italy, South Korea, Iran and China for non-Hungarians in response to the COVID-19 outbreak. However, for this Figure I consider only full closures, i.e. travel restrictions to all destinations.

Covid Economics 9, 24 April 2020: 80-99

Covid Economics 9, 24 April 2020: 80-99

Table 7: Bilateral air traffic intensity, travel restrictions, and differences in the timing of country-specific pandemic outbreak

Dependent variable: bilateral difference in	Days until first case	Days until start	Days until 100 cases	Days until first case	Days until start	Days until 100 cases
	(1)	(2)	(3)	(4)	(5)	(6)
Bilateral air traffic intensity (Flights)	-2.076***	-0.824***	-1.073***			
	(0.299)	(0.160)	(0.181)			
Bilateral air traffic intensity (Passengers)				-1.445***	-0.558***	-0.660***
				(0.220)	(0.140)	(0.136)
Travel restrictions in place before first case	19.517*			23.021**		
	(11.780)			(10.907)		
Travel restrictions in place before first case	-30.345					
× Bilateral air traffic intensity (Flights)	(19.340)					
Travel restrictions in place before start		7.366***			7.258***	
		(1.180)			(1.170)	
Travel restrictions in place before start		1.349***				
× Bilateral air traffic intensity (Flights)		(0.323)				
Travel restrictions in place before 100 cases			4.769***			4.878***
			(0.792)			(0.785)
Travel restrictions in place before 100 cases			1.620***			
× Bilateral air traffic intensity (Flights)			(0.322)			
Travel restrictions in place before first case				-20.313*		
× Bilateral air traffic intensity (Passengers)				(11.900)		
Travel restrictions in place before start					1.009***	
× Bilateral air traffic intensity (Passengers)					(0.241)	
Travel restrictions in place before 100 cases						1.070***
× Bilateral air traffic intensity (Passengers)						(0.228)
Common border	-1.495	0.243	-1.740***	-2.193	-0.047	-2.091***
	(1.506)	(0.637)	(0.604)	(1.492)	(0.629)	(0.599)
Distance between most populated cities	-0.653	0.609*	0.418*	-0.545	0.651**	0.517**
	(0.594)	(0.330)	(0.233)	(0.588)	(0.326)	(0.232)
Observations	1,389	1,389	1,389	1,389	1,389	1,389
R^2	0.362	0.740	0.755	0.367	0.741	0.751
Reporter FE	✓	✓	✓	✓	✓	✓
Partner FE	✓	✓	✓	✓	✓	✓

Source: Authors' elaboration on data from Eurostat, CEPII, and EU open data portal. The number of flight passengers is expressed in multiples of 100,000. The number of flights is expressed in multiples of 10,000. Geographic distance is expressed in multiples of 1,000 km. Bilateral air traffic intensity is measured as the total number of flights passengers between two countries from July 2018 to June 2019. The country-specific pandemic starting day is defined as the first day with a positive reported case followed by positive reported cases in the next three days. The country-specific pandemic starting day is defined as the first day with a positive reported case followed by positive reported cases in the next three days. Travel restrictions are defined as a dummy variable equal to one when a travel restriction was implemented between the countries at the date of the event defined in outcome variables (first case, growth start, or 100 cases), and zero otherwise. Robust standard errors are between parentheses. Robust standard errors are between parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Covid Economics 9, 24 April 2020: 80-99

Figure 3: Number of cases by day, selected countries

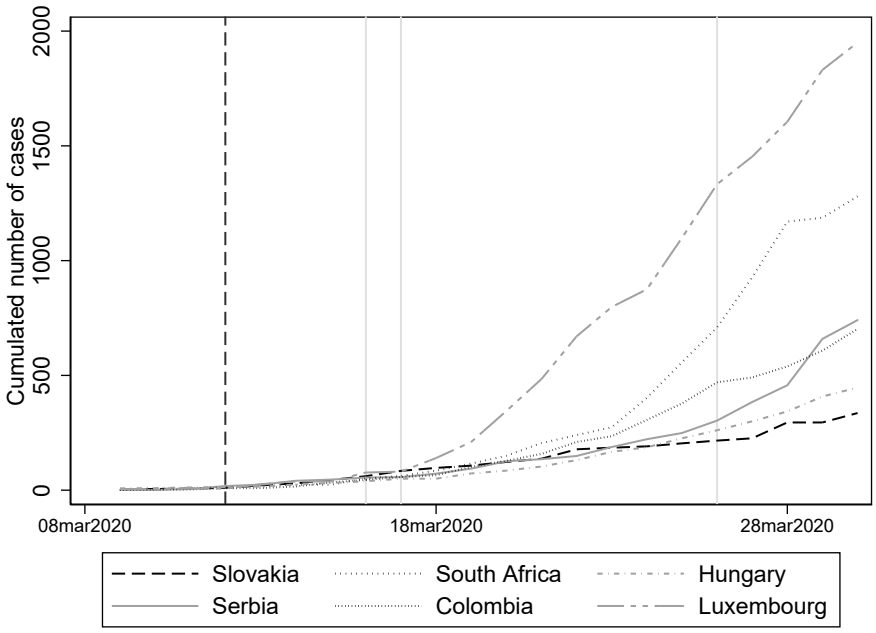

Source: Authors' elaboration on data from EU open data portal and self-collected information on traffic restriction implementations in countries.
Notes: Vertical lines correspond to the implementation of travel restrictions to all destinations, not bilateral restrictions. The dashed vertical line corresponds to march 13, 2020, when Slovakia implemented severe air traffic restrictions. The solid grey vertical lines correspond to the implementation dates of severe air traffic restrictions in Hungary, Colombia, Luxembourg, and South Africa.

5 Conclusions

In this study, I have reported suggestive evidence that air traffic played an important role in spreading the new corona virus COVID-19 around the world. Stronger connexions with China were associated with earlier dates of onset of the crisis. The results also show that, while very early implementations of travel restrictions were associated with delayed outbreaks of COVID-19 infections, most countries implemented such measures at a time when a relatively high number of cases had already been confirmed.

The high level of interconnection in the international travel system entails that even countries

with no direct connection to China are only one transit away. This characteristic could greatly reduce the effectiveness of bilateral travel restrictions in delaying the spread of the pandemic to new countries. Rather, timely and coordinated travel restrictions across many countries are more likely to be successful.

As countries will progressively start managing the way out of the crisis, it will be essential to take a network approach regarding the scope and timing of lifting travel restrictions. The central role of air travel for the global economy suggests that the post-crisis period will confront countries with an important trade-off between allowing a sufficient level in the international movement of people and health risks. As a takeaway for future outbreaks of infectious diseases, the results suggest that the design of a global emergency stop in international travel could be a promising avenue to limit economic damage and health risks. Such a system would require a high level of coordination at a multilateral level in order to preserve supply chains as much as possible.

Covid Economics 9, 24 April 2020: 80-99

References

Brockmann, D. and Helbing, D. (2013). The hidden geometry of complex, network-driven contagion phenomena. *science*, **342** (6164), 1337–1342.

Craig, A., Heywood, A. and Hall, J. (2020). Risk of covid-19 importation to the pacific islands through global air travel. *Epidemiology & Infection*, pp. 1–16.

da Silva Candido, D., Watts, A., Abade, L., Kraemer, M. U., Pybus, O. G., Croda, J., de Oliveira, W., Khan, K., Sabino, E. C. and Faria, N. R. (2020). Routes for covid-19 importation in brazil. *medRxiv*.

Foxwell, A. R., Roberts, L., Lokuge, K. and Kelly, P. M. (2011). Transmission of influenza on international flights, may 2009. *Emerging infectious diseases*, **17** (7), 1188.

Gilbert, M., Pullano, G., Pinotti, F., Valdano, E., Poletto, C., Boëlle, P.-Y., d'Ortenzio, E., Yazdanpanah, Y., Eholie, S. P., Altmann, M. *et al.* (2020). Preparedness and vulnerability of african countries against importations of covid-19: a modelling study. *The Lancet*, **395** (10227), 871–877.

Hsu, C.-I. and Shih, H.-H. (2010). Transmission and control of an emerging influenza pandemic in a small-world airline network. *Accident Analysis & Prevention*, **42** (1), 93–100.

Hufnagel, L., Brockmann, D. and Geisel, T. (2004). Forecast and control of epidemics in a globalized world. *Proceedings of the National Academy of Sciences*, **101** (42), 15124–15129.

Lau, H., Khosrawipour, V., Kocbach, P., Mikolajczyk, A., Ichii, H., Zacharksi, M., Bania, J. and Khosrawipour, T. (2020). The association between international and domestic air traffic and the coronavirus (covid-19) outbreak. *Journal of Microbiology, Immunology and Infection*.

Mayer, T. and Zignago, S. (2011). Notes on cepii's distances measures: The geodist database. *CEPII working paper 25*.

Zhuang, Z., Zhao, S., Lin, Q., Cao, P., Lou, Y., Yang, L. and He, D. (2020). Preliminary estimation of the novel coronavirus disease (covid-19) cases in iran: a modelling analysis based on overseas cases and air travel data. *International Journal of Infectious Diseases*.

CEPR PRESS

Appendix

Table A1: Bilateral air traffic intensity and differences in the timing of country-specific pandemic outbreak, reduced sample

Dependent variable: bilateral difference in	Days until first case	Days until start	Days until 100 cases	Days until first case	Days until start	Days until 100 cases
	(1)	(2)	(3)	(4)	(5)	(6)
Bilateral air traffic intensity (Passengers)	-0.600***	-1.548***	-0.422***			
	(0.138)	(0.240)	(0.115)			
Bilateral air traffic intensity (Flights)				-2.466***	-0.742***	-1.025***
				(0.342)	(0.163)	(0.191)
Common border	-1.798**	-0.934	0.636	-1.459**	-1.758	0.391
	(0.715)	(1.740)	(0.650)	(0.713)	(1.738)	(0.649)
Distance between most populated cities	1.709***	-0.033	1.634***	1.454***	0.443	1.833***
	(0.400)	(0.909)	(0.486)	(0.402)	(0.910)	(0.482)
Observations	784	784	784	784	784	784
R^2	0.549	0.449	0.755	0.563	0.441	0.751
Reporter FE	✓	✓	✓	✓	✓	✓
Partner FE	✓	✓	✓	✓	✓	✓

Source: Authors' elaboration on data from Eurostat, CEPII, and EU open data portal. The number of flight passengers is expressed in multiples of 100,000. The number of flights is expressed in multiples of 10,000. Geographic distance is expressed in multiples of 1,000 km. Bilateral air traffic intensity is measured as the total number of flights passengers between two countries from July 2018 to June 2019. The country-specific pandemic starting day is defined as the first day with a positive reported case followed by positive reported cases in the next three days. Robust standard errors are between parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Covid Economics 9, 24 April 2020: 80-99

Covid Economics 9, 24 April 2020: 80-99

Table A2: Bilateral air traffic intensity, travel restrictions, and differences in the timing of country-specific pandemic outbreak

Dependent variable: bilateral difference in	Days until first case	Days until start	Days until 100 cases	Days until first case	Days until start	Days until 100 cases
	(1)	(2)	(3)	(4)	(5)	(6)
Bilateral air traffic intensity (Flights)	-2.074***	-0.783***	-1.037***			
Bilateral air traffic intensity (Passengers)				-1.445***	-0.517***	-0.628***
				(0.220)	(0.127)	(0.125)
	(0.299)	(0.151)	(0.172)			
Travel restrictions in place before first case	10.292			10.420		
	(9.755)			(9.670)		
Travel restrictions in place before start		8.374***			8.456***	
		(1.172)			(1.176)	
Travel restrictions in place before 100 cases			5.663***			5.767***
			(0.742)			(0.745)
Common border	-1.632	0.199	-1.805***	-2.337	-0.073	-2.176***
	(1.499)	(0.633)	(0.603)	(1.486)	(0.627)	(0.603)
Distance between most populated cities	-0.644	0.634*	0.402*	-0.537	0.696**	0.530**
	(0.594)	(0.330)	(0.235)	(0.588)	(0.325)	(0.233)
Observations	1,389	1,389	1,389	1,389	1,389	1,389
R^2	0.362	0.739	0.751	0.366	0.739	0.747
Reporter FE	✓	✓	✓	✓	✓	✓
Partner FE	✓	✓	✓	✓	✓	✓

Source: Authors' elaboration on data from Eurostat, CEPII, and EU open data portal. The number of flight passengers is expressed in multiples of 100,000. The number of flights is expressed in multiples of 10,000. Geographic distance is expressed in multiples of 1,000 km. Bilateral air traffic intensity is measured as the total number of flights passengers between two countries from July 2018 to June 2019. The country-specific pandemic starting day is defined as the first day with a positive reported case followed by positive reported cases in the next three days. The country-specific pandemic starting day is defined as the first day with a positive reported case followed by positive reported cases in the next three days. Travel restrictions are defined as a dummy variable equal to one when a travel restriction was implemented between the countries at the date of the event defined in outcome variables (first case, growth start, or 100 cases), and zero otherwise. Robust standard errors are between parentheses. Robust standard errors are between parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A cost-benefit analysis of the Covid-19 disease

Robert Rowthorn[1]

Date submitted: 16 April 2020; **Date accepted:** 17 April 2020

The British government has been debating how and when to escape from the lockdown without provoking a resurgence of the Covid-19 disease. There is a growing recognition of the damage the lockdown is causing to economic and social life, including deaths and illness amongst the non-infected population. This paper presents a simple cost-benefit analysis based on optimal control theory and incorporating the SIR model of disease propagation. It concludes by presenting some simulations informed by the theoretical discussion. The main conclusions are: (1) the lockdown should be continued for some weeks, and (2) if there is an inexpensive way of reducing the net reproductive rate of the disease to r = 1 , this policy should be adopted within a few weeks of exiting lockdown. It is not cost-effective to linger in intermediate stages with more expensive policies designed to keep r well below unity with the hope eradicating the disease.

Covid Economics 9, 24 April 2020: 100-109

1 Emeritus Professor, Faculty of Economic, University of Cambridge. I am grateful to Paul Ormerod, Flavio Toxvaerd, Wendy Carlin and especially Nicholas Rau for their comments.

Covid Economics 9, 24 April 2020: 100-109

A Cost-Benefit Analysis of the Covid-19 Disease

There has been a debate in Britain about the best policy for dealing with the Covid-19 virus. The official policy was originally to proceed step by step and intensify, as required, the measures that encourage hygiene and social distancing. Such measures range from careful hand-washing through to the banning of large public gatherings, the closing of pubs, restaurants and many shops, quarantine or near quarantine of vulnerable people, and restrictions on national and international travel. The gradualist approach of the government was attacked by critics who called for immediate draconian action of the type observed in Italy and Spain. The government has responded by implementing an unprecedented lockdown on economic and social life. A factor behind its change heart was concern about the potential shortage of intensive care beds if the disease was not brought quickly under control. The government is now seeking ways to relax the current lockdown without provoking a surge in the disease.

The measures required to inhibit disease transmission can be very costly in economic and social terms, including the resulting illness and death in the non-infected population. These costs must be weighed against the medical benefits of intervention.The decision of when to intervene and on what scale is a classic optimum control problem. There is a well-established economic theory for dealing with such problems. This note explores the choice facing the government using a simple mathematical model based on optimal control theory. It complements the theoretical analysis with some illustrative simulations.

The economic literature on the optimal control of epidemics is sparse and its models mostly deal with individual behaviour or the externalities of individual decision-making with regard to treatment, vaccination or social distancing[1]. These are not my concern here. My interest is in the cost-benefit analysis of large-scale interventions such as lockdowns. This involves an approach that is unusual in the existing economic literature on disease. Costs and benefits in existing disease models are typically functions of the health status of individuals, computed by assigning values or weights to individuals according to their health status. This is a procedure followed here. However, unlike these models I also make an explicit allowance for the more general costs of comprehensive interventions such as lockdowns. Such costs depend on the scale and type of intervention but they are not linked in any direct way to the health status of individuals. These costs are given a central role in this paper.

The Model

[1] See Toxvaerd (2020) for a brief survey of this literature. Among the articles worthy of note are Chen (2012), Chen et al. (2011), Fenichel (2013), Gersovitz (2010), Reluga (2010), Rowthorn and Toxvaerd (2015), Sethi (1978), Toxvaerd (2019), Toxvaerd and Rowthorn (2020).

Covid Economics 9, 24 April 2020: 100-109

The analysis in this paper uses a modified version of the standard SIR model of disease propagation. Ignoring births and deaths from non-Covid-19 causes, the initial population P_0 will divide in the future into three groups of people: susceptible, infected, and removed, denoted, respectively, by $S(t), I(t)$ and $R(t)$. The removed group includes people have died from the disease since time 0. They are denoted by $D(t)$ The initial population is normalised to 1, so these various quantities can be interpreted as shares. Individuals who are infected remain infectious until they recover or die. Infected individuals who recover acquire immunity, so the journey from $S(t)$ via $I(t)$ to $R(t)$ is in one direction only.

The dynamics of the disease are determined by the following equations:

$$\frac{dS(t)}{dt} = -\beta(t)S(t)I(t) \quad (1)$$

$$\frac{dI(t)}{dt} = \beta(t)S(t)I(t) - \ I(t) \quad (2)$$

$$\frac{dR(t)}{dt} = \ I(t) \quad (3)$$

$$\frac{dD(t)}{dt} = \delta I(t) \quad (4)$$

$$I(0) = I_0 \geq 0 \quad (5)$$

$$S(0) = S_0 \geq 0 \quad (6)$$

$$R(0) = R_0 \geq 0 \quad (7)$$

$$D(0) = 0 \quad (8)$$

$$I(t) + S(t) + R(t) = 1 \quad (9)$$

where snd δ are constant. Equation (1) assumes that, for a susceptible individual, the probability of becoming infected is proportional to the share of infected people in the population. The coefficient $\beta(t)$ is a variable which depends on the current intensity of social interaction. The intensity of social interaction depends, in turn, on the measures that the government puts in place to inhibit the spread of the disease. Specifically, it is assumed that

$$\beta(t) = [1 - q(t)]\beta_0 \quad (10)$$

where $q(t)$ is an index of policy severity. Such intervention comes at a cost $C(q(t))$ in the form of damage to the economy and social life. This cost is independent of the number of people currently infected and is the result of society-wide measures to control the disease. It is in addition to the costs arising directly from infection. The function $C(.)$ is assumed to be twice differentiable and such that

$$C(0) = 0, C(q_{\max}) = C_{\max} < \infty, \ C'(0) = 0 \quad (11)$$

$$C \geq 0, C' \geq 0, \ C'' > 0 \text{ for } q(t) \in [0, q_{\max}] \quad (12)$$

where $q_{\max} \leq 1$ is an upper limit beyond which it is not politically feasible to increase q. Thus, $C(q)$ is strictly convex over the relevant range. An example

is $C(q) = Aq^{1+\phi}$ with $\phi > 0$. When q is close to zero, the marginal cost of intervention is low but rises steeply at higher values of q. These are realistic assumptions. Think of hand-washing at one end of the scale and the closure of international frontiers at the other.

The government is assumed to have perfect foresight and knows that an effective vaccine will become available at time T at a cost of c_v per patient[2]. It chooses the trajectory $q(t)$ so as to minimise the followng quantity

$$J = \int_0^T e^{-\rho t}[\pi_I I(t) + C(q(t)]dt + e^{-\rho T}F(I(T), S(T)) \tag{13}$$

subject to the equations (1) to (12) where $\pi_I > 0$ is the loss from being infected.[3] It includes an allowance for the cost of treatment and for death. The terminal function $F(I(T), S(T))$ is the cost of vaccination plus the ever shrinking stream of infected individuals left over from the pre-vaccination era, all discounted back to the beginning of the programme. This can be expressed as follows

$$e^{-\rho T}F(S(T), I(T)) = e^{-\rho T}c_v S(T) + \int_T^\infty e^{-(\rho+\gamma)t}\pi_I I(T)dt \tag{14}$$

which implies that

$$F(S(T), I(T)) = c_v - c_v R(T) + \left(-c_v + \frac{\pi_I}{\rho +}\right) I(T) \tag{15}$$

A necessary condition for optimality is for $q(t)$ to maximise the following current value Hamiltonian:

$$H = -\pi_I I(t) - C(q(t)) + \lambda(t)\frac{dI(t)}{dt} + \mu(t)\frac{dR(t)}{dt} \tag{16}$$

where $\lambda(t)$ and $\mu(t)$ are the shadow prices or costate variables. The terminal values of these shadow prices are

$$\lambda(T) = c_v - \frac{\pi_I}{\rho +} \tag{17}$$

$$\mu(T) = c_v > 0 \tag{18}$$

[2] In a game theoretic paper on social distancing Reluga (2010) also assumess that vaccination will occur on a fixed date in the future.

[3] Let π_A be the monetary value that planners assign to each person who becomes infected and survives, and π_D the value they assign to those who die. The flow of costs associated with infection is as follows:

$$\begin{aligned} \text{Flow of costs} &= \pi_A \frac{d(R-D)}{dt} + \pi_D \frac{dD}{dt} \\ &= \pi_A(\gamma - \delta)I + \pi_D \delta I \\ &= \pi_I I \end{aligned}$$

where $\pi_I = (\gamma - \delta)\pi_A + \delta\pi_D$.

Covid Economics 9, 24 April 2020: 100-109

Covid Economics 9, 24 April 2020: 100-109

Differentiating equation (16)

$$\frac{\partial H}{\partial q} = -C'(q(t)) - \lambda(t)\beta_0 S(t)I(t) \tag{19}$$

If $\frac{\partial H}{\partial q} = 0$ for some $q(t)$ in the open set $(0, q_{\max})$, there may be an interior solution to the optimisation problem (assuming it has a solution). Failing this, the solution must perforce be on the boundary and hence equal to 0 or $q_{\max}$.

The evolution of the shadow prices is determined by the following equations

$$\frac{d\lambda(t)}{dt} = \rho\lambda(t) - \frac{\partial H}{\partial I} \tag{20}$$

$$\frac{d\mu(t)}{dt} = \rho\mu(t) - \frac{\partial H}{\partial R} \tag{21}$$

Suppose the equation $\frac{\partial H}{\partial q} = 0$ has a solution $q^*(t)$ in the open set $(0, q_{\max})$, If a solution to the optimisation problem exists this solution must be $q^*(t)$. The value of q that maximises H is given by:

$$q(t) = 0 \text{ if } \lambda(t)\beta_0 S(t)I(t) \geq -C'(0) \tag{22}$$

$$0 < q(t) = q^*(t) < q_{\max} \text{ if } -C'(q_{\max}) < \lambda(t)\beta_0 S(t)I(t) < -C'(0) \tag{23}$$

$$q(t) = q_{\max} \text{ if } \lambda(t)\beta_0 S(t)I(t) \leq -C'(q_{\max}) \tag{24}$$

These conditions may result in finite intervals over which the optimal $q(t)$ takes intermediate values.

There is no explicit solution to the above system of equations.

Simulations

In the absence of an explicit solution, the obvious procedure is to explore the properties of the system by means of numerical simulation. A solution can in theory be found by backward induction using a model with transversality conditions given by equations (17) and (18). In practice this may be difficult or even impossible using existing techniques. A common method for computing the solution to optimisation problems is the iterative forward-backward sweep method. However, this did not work in the present case. The iterative process normally converged but it generated solutions that were demonstrably wrong.

Being unable to compute the fully optimal path, I decided to investigate a more limited, although related, problem. The following simulations come with a health warning. I have done my best to find realistic values for the various parameters, but this is an area of great uncertainty and considerable guesswork is involved. The problem is as follows.

Suppose the government is currently operating a lockdown policy and is formulating a policy for the future. It has at its disposal four types of intervention specified by $q = q_{\max} > q_1 > q_2 > q_3$, the first of which is lockdown.

Covid Economics 9, 24 April 2020: 100-109

These interventions must be implemented sequentially, although some may be left out altogether. The optimisation problem is to choose the timing of each intervention. This will be influenced by their cost which is given by the function:

$$C(q) = C_{\max} \times \left(\frac{q}{q \max} \right)^{1+\phi} \tag{25}$$

Thus, $C_{\max}$ is the cost of lockdown. The larger is the value of ϕ the lower is the cost of the other interventions relative to lockdown.

Parameter Values

The simulations use a week as their unit of time and their horizon is $T = 52$. The monetary unit of account is thousands of UK pounds per capita per week. The simulatons use the highly accurate Runge-Kutta fourth-order algorithm to solve the differential equations.

There is uncertainty about the extent of the Covid-19 disease in the UK and therefore uncertainty about the true death rate. The extent is almost certainly greatly understated by official figures for confirmed cases (see the appendix). In my main simulations the initial conditions are $I_0 = 0.009, R_0 = 0.03, S_0 = 0.961$. I assume a death rate of 0.8 percent.

Other parameters have the following values: $\beta_0 = 3.4$, $= 1.4$, $C_{\max} = 0.20, \pi_I = 34$, $c_v = 0.05$, $\rho = 0$. Infected individuals cease to be infectious at an exponential rate of $= 1.4$ per week, which implies that they are on average infectious for 4.7 days. After two weeks 94 percent are no longer infectious. They have either recovered or died. I assume that in the absence of intervention the net reproduction ratio $r_0 = \beta_0 / \ = 2.5$. This implies that $\beta_0 = 3.4$. The cost of the vaccine c_v is £50. The discount rate ρ is zero because the time horizon is short, although one can make a case for a positive discount rate to allow for uncertainty. The per capita weekly cost of the lockdown is £200 which is approximately 35 percent of per capita GDP at factor cost, in line with the OBR estimate of what the lockdown may do to the UK to the economy (OBR, 2020).

The large value of $\pi_I = 34$ is to allow for the risk of death from infection. It is based on the assumption that planners assign a monetary value of £10,000 to each non-fatal infection, including the cost of treatment for those admitted to hospital, and £2.5 million for each fatality from the disease. Assuming that the risk of death for people who catch the disease is 0.8 percent, the expected monetary value of being infected is £34,000. These figures are little more than guesswork. The true death rate is unknown and there is no consensus in the literature about the valuation of human life. A meta-analysis for the OECD (2012) surveyed 250 studies and found that estimates of the VSL (value of statistical life) in the health sector varied from US$4,450 to US$22,100,000 at 2005 prices. It also pointed out, but did not investigate in depth, that estimates of the VSL are typically lower for people who are old or have underlying health issues, which is the situation with Covid-19.

It is assumed that $q_{\max} = 0.75$ $q_1 = 0.70, q2 = 0.65, q3 = 0.6$. The net reproduction rate $r_0 = 2.5$ is reduced by lockdown to $(1 - 0.75) \times 2.5 = 0.625$

CEPR PRESS

whereas the alternative interventions reduce r_0 to $r = 0.75, r = 0.875$ and $r = 1.0$ respectively. The cost of each intervention depends on ϕ and is shown in Table 1.

Results

Table 2 shows the optimal intervention times and values of the programme (J in equation (13)) for various values of ϕ. The first thing to note is that in none of the scenarios is the lockdown abandoned immediately, although there is significant variation about how long it is kept in place, ranging from 2 to 5 weeks. The second thing to note is that under all of the scenarios the time spent in the intermediate states is short. In every case, within less than 4 weeks after exiting lockdown, policy has transitioned to the minimal intervention option with $q_3 = 0.6$. This is interesting, because the net reproduction ratio in this case is $r = 1$ and the disease is not on the way to eradication.

To the extent that one can talk about realism in the present context, the most realistic scenario is shown in line (3) of Table 2. Under this scenario, the total cost arising from the disease, including economic impacts and the burden of disease, is £6,500 per head of population. The full lockdown is maintained for 3 weeks, which is about average for the scenarios shown in the table. The table also shows what happens when the time horizon in this scenario is extended from 52 to 104 weeks (line 5). The results are hardly altered. This suggests that uncertainty about the terminal date is not a serious problem. The same is true if there is a general reduction in the monetary valuation of infection (π_I). If this valuation is cut by a half (line 6) the lockdown is abandoned a week earlier and the policy arrives a bit earlier at the least expensive option. The table also shows (line 9) a scenario in which the numbers of people who have been infected and those who are still infectious are initially much lower than in the other scenarios. This implies a higher death rate and hence a higher value of π_i. This modification has surprisingly little effect on the optimal timetable, although by the end of the year there are, naturally, a lot fewer people who have been infected.

Concluding Remarks

People are starting to ask the obvious question with regard to the Covid-19 virus. "Is the cure worse than the disease?" Governments are urgently seeking cost-effective policies that will enable them to exit the lockdown without setting off a renewed surge of infection. Although they are highly speculative, the estimates presented here may throw some light on the subject. The most robust conclusion is that, if a relatively inexpensive way can be found to reduce the net reproduction ratio to $r = 1$, that is the policy to aim for in the medium term (a few months). It would be a mistake to stick with more expensive policies that aim to eradicate the disease by keeping r well below 1.

Covid Economics 9, 24 April 2020: 100-109

Covid Economics 9, 24 April 2020: 100-109

Table 1: The Cost of Intervention

	Weekly Cost of Intervention *£ thousands per capita*			
	$q_{\max} = 0.75$ $r = 0.0625$	$q_1 = 0.70$ $r = 0.75$	$q_2 = 0.65$ $r = 0.875$	$q_3 = 0.60$ $r = 1.00$
$\phi = 0.5$	0.200	0.193	0.186	0.179
$\phi = 1$	0.200	0.187	0.173	0.160
$\phi = 2$	0.200	0.174	0.150	0.128
$\phi = 4$	0.200	0.152	0.113	0.082

Table 2: Optimal Timing of Interventions

		Value of Programme *£ per capita*	Optimal Timing				Total Infected *Millions*
			$q_{\max} = 0.75$	$q_1 = 0.7$	$q_2 = 0.65$	$q_3 = 0.6$	
(1)	$\phi = 0.5$	8, 400	0 to 4.8	4.8 to 5.2	5.2 to 5.6	5.6 to 52	3.7
(2)	$\phi = 1.0$	7,700	0 to 4.2	4.2 to 4.8	4.8 to 5.4	5.4 to 52	3.8
(3)	$\phi = 2.0$	6,500	0 to 3.2	3.2 to 4.2	4.2 to 5.4	5.4 to 52	3.9
(4)	$\phi = 4.0$	4,700	0 to 2.0	2.0 to 3.6	3.6 to 5.0	5.0 to 52	4.4
(5)	$\phi = 2$*	11,800	0 to 2.8	2.8 to3.8	3.8 to 5.8	5.8 to 104	4.0
(6)	$\phi = 2$**	6,100	0 to 1.8	1.8 to 2.8	2.8 to 4.2	4.2 to 52	4.5
(7)	$\phi = 2$***	6,000	0 to 1.8	1.8 to 3.4	3.4 to 5.4	5.4 to 52	1.1

*T = 104 instead of 52; ** $\pi_I = 17$ instead of 34; *** $I_0 = 0.0018, R_0 = 0.006, \pi_I = 48$.

Covid Economics 9, 24 April 2020: 100-109

Appendix

The Covid Symptom Tracker (King's, 2020) shows a steeply increasing number of currently infected people rising to a peak of 2.0 million on 1 April followed by a decline to 580 thousand on 16 April. The lockdown began on 26 March, but took some days to take full effect. Assume that prior to 1 April the disease coefficients were $\beta = 3.4$ and $\quad = 1.4$, and thereafter $\beta = 0.25 \times 3.4 = 0.85$ and $\quad = 1.4$. The former assumption implies that by 15 April $I = 0.033$ (2.0 million) and $R = .024$ (1.6 million). The latter assumption implies that by 15 April $I = .009$ (0.6 million) and $R = 0.080$ (5.3 million). Thus, by 15 April almost 6 million people had been infected.

These figures are far removed from the official figure for confirmed cases. They may be exaggerated by time-varying selection bias and misdiagnosis, but nevertheless they suggest that Covid-19 is already quite widespread. Accordingly, I have arbitrarily assumed a figure of 2 million recovered and 600 thousand currently infected in the base year.

The death rate δ is given by the following formula[1]i

$$\delta = \quad \times \frac{\text{deaths}}{\text{recovered}} \tag{1}$$

If deaths are equal to 14 thousand, $\quad = 1.4$ and recovered equal to 2 million, $\delta = 1.0$ percent. With 5.3 million recovered the death rate is 0.4 percent. Going forward, the simulations assume 0.8 percent.

[1] See footnote 4

Covid Economics 9, 24 April 2020: 100-109

Bibliography

Chen, F. (2012); A Mathematical Analysis of Public Avoidance Behavior During Epidemics Using Game Theory, *Journal of Theoretical Biology*, 302, 18-28.

Chen, F., M. Jiang, S. Rabidoux and S. Robinson (2011): Public Avoidance and Epidemics: Insights from an Economic Model, *Journal of Theoretical Biology,* 278, 107-119.

Fenichel, E. P.(2013): Economic Considerations for Social Distancing and Behavioral Based Policies During an Epidemic,*Journal of Health Economics*, 32(2), 440-451.

Gersovitz, M. (2010): Disinhibition and Immizeration in a Model of Susceptible -Infected-Susceptible (SIS) Diseases, *mimeo.*

King's College, London (2020): https://covid.joinzoe.com/

OBR (2012): Coronavirus: reference scenario. https://obr.uk/coronavirus-reference-scenario/

OECD (2012): The Value of Statistical Life: A Meta-Analysis,ENV/EPOC/ WPNEP(2010)9/FINAL.

Reluga, T. C.(2010): Game Theory of Social Distancing in Response to an Epidemic,*PLoS Computational Biology*, 6(5):e1000793.

Rowthorn, R. and F. Toxvaerd (2015): The Optimal Control of Infectious Diseases via Prevention and Treatment, *mimeo.*

Sethi, S. P. (1978): Optimal Quarantine Programmes for Controlling an Epidemic Spread, *Journal of the Operational Research Society,* 29(3), 265-268.

Toxvaerd, F. (2019): Rational Disinhibition and Externalites in Prevention, *International Economic Review, 60(4), 1737-1755.*

Toxvaerd, F. (2020): Equilibrium Social Distancing, Cambridge-INETInet Working Paper Series No: 2020/08.

Toxvaerd, F. and R. Rowthorn (2020): On the Management of Population Immunity, *mimeo.*

Covid Economics Issue 9, 24 April 2020

The costs and benefits of home office during the Covid-19 pandemic: Evidence from infections and an input-output model for Germany[1]

Harald Fadinger[2] and Jan Schymik[3]

Date submitted: 15 April 2020; **Date accepted:** 17 April 2020

We study the impact of working from home on (i) infection risk in German regions and (ii) output using an input-output (IO) model of the German economy. We find that working from home is very effective in reducing infection risk: regions whose industry structure allows for a larger fraction of work to be done from home experienced much fewer Covid-19 cases and fatalities. Moreover, confinement is significantly more costly in terms of induced output loss in regions where the share of workers who can work from home is lower. When phasing out confinement, home office should be maintained as long as possible, to allow those workers who cannot work from home to go back to work, while keeping infection risk minimal. Finally, systemic industries (with high multipliers and/or high value added per worker) should be given priority, especially those where home office is not possible.

1 Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B06) is gratefully acknowledged.
2 Professor of Economics, University of Mannheim and CEPR Research Fellow.
3 Post-Doctoral Fellow in Economics, University of Mannheim.

Covid Economics 9, 24 April 2020: 110-137

Covid Economics 9, 24 April 2020: 110-137

1 Introduction

Reactivating the economy after the pandemic shutdown is a key challenge for policy makers because there is a clear trade-off between keeping infections rates low and minimizing output losses. We show that *(i)* home office is a very effective tool for reducing infection rates: regions with fewer workers that can work from home due to the nature of their occupation and industry composition have experienced higher Covid-19 infection rates and fatalities; *(ii)* the economic costs of confinement are significantly higher in regions where a smaller fraction of jobs can be done in home office.

To provide evidence for the first point, we regress regional infection rates on the region-specific share of workers who can work from home (*WFH* share) and find a strongly significant negative correlation. Our estimates imply that a one percentage point higher WFH share is associated with 20 fewer infections and 0.9 fewer fatalities per 100 thousand inhabitants as of April 9, 2020. To give a concrete example how this translates into absolute numbers, consider Berlin (WFH share of 0.45 and population of 3,8 Million) and Niederbayern (WFH share of 0.38). If Berlin had the same share of workers working from home as Niederbayern, it would have experienced around 5,300 additional Covid-19 infections (more than double the actual number) and 200 additional deaths (more than 4 times the actual number) by April 9, 2020. We also estimate a simple epidemiological model to show that working from home more generally translates into siginificantly reduced infection rates.

Based on this observation, we then look at the economic costs of strict confinement using a static 62-sector model with input-output linkages calibrated to the German economy. Industries vary in terms of their WFH share and their value-added multipliers – the magnitude of value added changes implied by a given change in industry employment – and the industry composition varies across regions. We consider a situation where only workers who can work from home are allowed to work (confinement). This policy reduces the labor supply by 58%. We show that the cost of strict confinement is very high: the output loss for the German economy

is around 1.6% of GDP per week. However, this cost varies a lot across regions due to differences in regional industry composition. We find that having a 10 percentage point higher fraction of jobs that can be done at home is associated with a 0.4% smaller weekly GDP reduction from strict confinement. The highest economic costs of confinement are faced by those regions where the share of industries characterized by high multipliers and a low WFH share is particularly large.

Finally, we consider policies for phasing out confinement. Our first message is that – in all industries and regions – workers who can work from home should stay in home office as long as a significant infection risk is present. This allows sending workers who cannot work from home back to their work place, while keeping infection risk as low as possible. Second, confinement should be phased out in a region-specific way that balances increases in infection risk associated with sending more people to work with output losses per worker. Finally, systemic industries should be given priority, especially those in which home office is not possible. These are either industries producing key inputs used by virtually all sectors in the economy and a low WFH share, such as Petroleum and Coke Refinement, Chemicals, or Construction, or industries with very high value added per worker, like Pharmaceuticals, Telecommunications or Cars.

This study relates to the literature studying the effects of Covid-19 confinement rules on the economy. Dingel and Neiman (2020) estimate the fraction of jobs that can be done at home in the U.S. and find similar industries to be intensive in those jobs as we do for Germany. Barrot et al. (2020) study the costs of the shutdown in France - their estimate is a weekly loss of about 1% of French GDP. We estimate the weekly GDP loss for Germany if only jobs that can be done at home were to remain in the labor force to be a weekly GDP loss of 1.6%. Koren and Peto (2020) show that U.S. businesses that require face-to-face communication or close physical proximity are particularly vulnerable to confinement. We relate to their results by showing that there is a tight link between regional variation in jobs that can be done at home to Covid-19 infections and then quantify the output loss from taking jobs that require physical presence out of the labor force.

Covid Economics 9, 24 April 2020: 110-137

Figure 1: Share of Jobs that can be Performed at Home

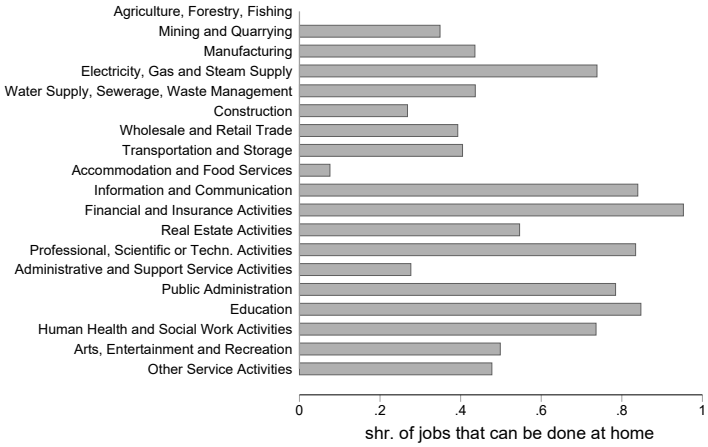

Notes: The Figure plots the estimated shares of jobs that can be done at home in the German economy across broad NACE Rev. 2 industries.

Hartl et al. (2020) identify a trend break in German Covid-19 infections growth subsequent to the implementation of social distancing policies.

The remainder of this paper is organized as follows. Section 2 presents empirical evidence on the relation between jobs that can be done at home and the spread of Covid-19 across regions. Based on that relation, section 3 discusses reactivation policies and the economic costs of confinement based on a simple structural model of the German production network. Finally, section 4 concludes.

2 Working from Home and the Spread of Covid-19

To estimate how many jobs can be performed at home and how "working from home" (WFH) can be seen as an effective social distancing measure, we classify the feasibility of WFH for all occupations and merge this classification with occupational employment counts for Germany. Our measure for potential WFH jobs is based on Eurostat data. Overall, we estimate that a

Covid Economics 9, 24 April 2020: 110-137

maximum of 42% of jobs in Germany could potentially be done from home. This number seems reasonably close to the 37% of WFH jobs that Dingel and Neiman (2020) calculate for the U.S. economy. The three sectors with the highest share of jobs that can be done from home are 'Financial and Insurance Activities' (NACE Rev. 2 code K), 'Information and Communication Services' (J), and 'Education' (P). The three sectors with the lowest share of jobs that can be done from home are 'Agriculture, Forestry and Fishing' (A), 'Accommodation and Food Services' (I) and 'Construction' (F).

What is key in our analysis is that we want an exogenous measure of the regional ability to work from home that is not driven by the endogenous response of people due to the spread of Covid-19 infections. This is because in a region with more Covid-19 infections more people will be endogenously induced to work from home, leading potentially to a spurious positive relationship between working from home and Covid-19 cases. We thus aggregate the fraction of WFH jobs in each industry to the regional level using regional employment shares of each industry. This gives us a region-specific measure of workers' ability to perform their jobs from home. If more social distancing causally reduces the spread of Covid-19 infections we expect to find a negative relationship between Covid-19 infections and the WFH share.

Plotting the variation in Covid-19 cases and WFH across regions on a map in Figure 2 suggests that there is some regional clustering for both, WFH and the spread of Covid-19 within former Eastern and Western territories. To rule this out and to control for other confounding factors potentially correlated with Covid-19 infections and the WFH share, we control for differences in population, area, economic activity, former Eastern German region status and the share of workers in the 'Accommodation and Food Services' industry within each region. We hence correlate the regional WFH share with Covid-19 infections and case fatalities to evaluate the impact of social distancing at the workplace on the spread of Covid-19. We regress the measures of Covid-19 on regional WFH shares including these control variables in Table 1.

Covid Economics 9, 24 April 2020: 110-137

Figure 2: Regional Clustering of Covid-19 and Working from Home Jobs

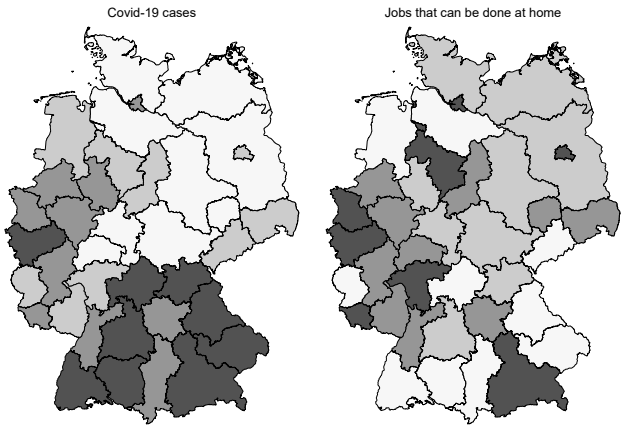

Notes: The Maps plot Covid-19 cases per 100 thd. inhabitants (left) or the share of jobs that can be done from home (right) across NUTS-2 regions in Germany. Darker colors correspond to higher values. Data are from Robert-Koch-Institut (based on April 9, 2020) and Eurostat.

This statistical association suggests that a one percentage point higher share of jobs that can be conducted from home is associated with 20 fewer infections and 0.9 fewer fatalities per 100 thousand inhabitants as of April 9, 2020. Figure 3 indeed shows a strong negative correlation between WFH jobs and disease spread, both in terms of infections and fatalities. Consider the following thought experiment to interpret this correlation. The region Lower Bavaria (Niederbayern) is strongly affected by Covid-19 infections and its regional WFH share is relatively low at 38% compared to 45% in Berlin. If Berlin had a WFH share as low as Lower Bavaria, there would be more than 4,000 additional infections and about 50 additional fatalities to be expected.

To quantitatively assess the impact of WFH on regional Coronavirus infection rates we now use a very simple epidemiological model. For simplicity, we structurally estimate a basic SIS model (see Hethcote (1989)), because it allows for an explicit solution of the infection rate as a function of param-

Covid Economics 9, 24 April 2020: 110-137

Covid Economics 9, 24 April 2020: 110-137

Figure 3: Covid-19 and Working from Home

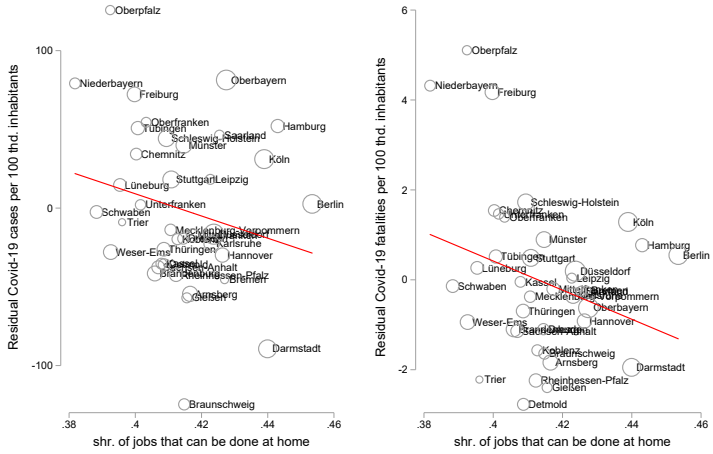

Notes: The Figure displays scatterplots of Covid-19 cases (left) or deaths (right) per 100 thd. inhabitants and the share of jobs that can be done from home across NUTS-2 regions in Germany. Individual dots are population-weighted. Data are from Robert-Koch-Institut (based on April 9, 2020) and Eurostat.

eters, which is not the case for more complicated models.[1] The infection rate $I_r(t)$, defined as the number of infected persons per population t days after an initial outbreak in region r, can be expressed as follows:

$$I_r(t) = \frac{e^{(\lambda_r - \gamma)t}}{\frac{\lambda_r}{\lambda_r - \gamma}(e^{(\lambda_r - \gamma)t} - 1) + I_0^{-1}}, \tag{1}$$

where λ_r is the contact rate (the average number of contacts per infective per day), γ is the removal rate (or recovery rate) and I_0 is the initial infection rate on day 0. We allow the contact rate to depend on the fraction of people that have jobs that can be done at home and posit the functional form $\lambda_r = \lambda_0 + \beta\, WFH_r + \delta\, GDP_r$. Thus, the region-specific contact rate depends on regional economic activity (log GDP), the WFH share with

[1] Strictly speaking, the SIS model is probably not adequate because it assumes that infected individuals do not acquire immunity from the disease. So far, the evidence suggests that recovered individuals acquire at least temporary immunity. However, this assumption does not make much of a difference in the early stages of the Covid-19 outbreak because initially, the entire population is susceptible.

Table 1: Conditional Correlations between Covid-19 and Working from Home

	Covid-19 Cases		Covid-19 Fatalities	
	(1)	(2)	(3)	(4)
Pct. WFH jobs	-20.47**	-16.37***	-0.941***	-0.750***
	(9.249)	(0.0058)	(0.334)	(0.0002)
Controls	yes	yes	yes	yes
Population weights	no	yes	no	yes
NUTS-2 regions	38	38	38	38

Notes: Dependent variables are the number of Covid-19 cases or the number of Covid-19 fatalities per 100 thousand inhabitants up to April 9, 2020 at the NUTS-2 level based on data from the Robert-Koch-Institut. Controls are region-specific population, area, GDP (all in logs), a dummy indicator for regions in former Eastern Germany and the share of workers in the 'Accommodation and Food Services' industry. Standard errors are heteroskedasticity-robust. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

slope β and an intercept λ_0. The unknown parameters are thus λ_0, β, δ and γ. We set $t = 0$ for the first large outbreak in Germany, dated with February 27, 2020 (after the Heinsberg outbreak) such that our data from April 9, 2020 are 42 days after the initial outbreak and $I_0 = 10^{-8}$.[2] We then estimate equation (1) using first a grid search over parameters to find starting values and then running a non-linear least squares estimator. The point estimate of the key parameter of interest β, is -0.33. This parameter estimate is significant at the 5% level with a standard error of 0.145. Thus, a 1 percentage point increase in WFH is associated with a 1/3 percentage point drop in the contact rate. As an example, moving from a WFH share of 0.38 (Niederbayern) to a WFH share of 0.45 (Berlin), reduces the contact rate by 2.3 percentage points (0.023=0.33(0.45-0.38)). This relatively small drop in the contact rate leads however to large quantitative effects on the infection rate. To illustrate that this drop in the contact rate matters,

[2]While the estimate of γ depends on the choices of the initial date and I_0, the estimate of β is not sensitive to these parameters. We also estimate the model using infection rates from April 14, 2020 and obtain similar results.

Covid Economics 9, 24 April 2020: 110-137

CEPR PRESS

Figure 4: Covid-19 Infection Rates and Working from Home

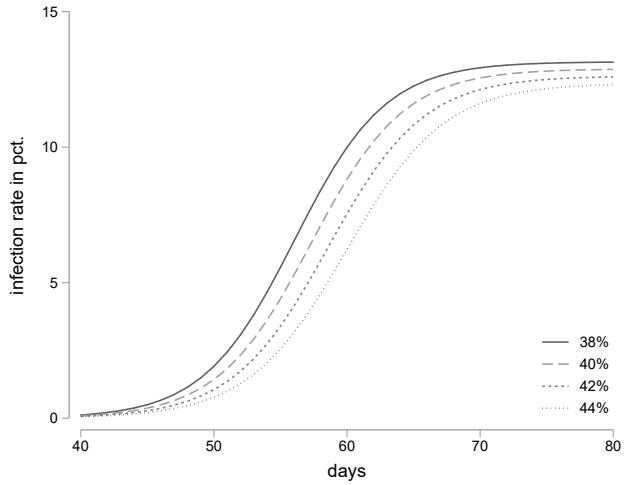

Notes: The Figure plots infection rate paths for different WFH shares based on the estimated SIS model (1).

we use the estimated epidemiological model to predict infection rates in each region using the empirical WFH shares in Figure 2. Even though we want to emphasize here that the results of the quantitative analysis should be interpreted with some caution as we are not trained epidemiologists, a general takeaway is that keeping the contact rate at the workplace as low as possible is key in reducing the spread of Coronavirus. According to our analysis, an effective way of doing this is to increase the WFH share as much as possible. Furthermore, one should note that even though the share of jobs that can potentially be done from home provides some source of exogenous variation to Covid-19 infections, this is likely a noisy measure of the true share of jobs that stay at home during the confinement phase. Consequently, we might underestimate the effect of confinement at the workplace on Covid-19 infections, here.

Covid Economics 9, 24 April 2020: 110-137

3 How can Confinement Rules be Phased-Out?

3.1 Modelling the Economic Costs of Confinement in the German Production Network

Motivated by the strong statistical association between WFH jobs and the spread of Covid-19, we want to evaluate the impact of strict confinement strategies on economic output. To analyze the sectoral effects of the Covid-19 shock we start from a standard model of production networks as in Jones (2013). The economy consists of many industries, linked with each other through an input-output network. The goods in each industry are produced by a representative firm that uses capital, labor and other industries' goods as inputs for production according to industry-specific Cobb-Douglas production functions. To take into account that the social distancing policies are implemented in the short run, we do not allow for long-run adjustments of economic factor allocations across sectors. We thus assume that producers may choose their intermediate inputs optimally but that sectoral capital and labor endowments remain fixed. Furthermore, we assume for simplicity that jobs that can be done from home and jobs that can be done only at the workplace are perfect substitutes.[3] Confinement is then modelled as an industry-specific shock to labor supply: We assume that only workers in WFH jobs remain in the labor force and compare value added under this hypothesis with value added when the sectoral labor force is fully available. In addition, to asses the benefit of loosening confinement in an industry-specific way, we also compute the percentage effect on GDP of increasing each industry's labor force by one percent and the marginal value added effects of letting an additional worker of a given industry return to the workforce.

Based on these assumptions, the relative GDP change of increasing the labor supply in an individual industry i by a share ΔL_i is given by the product of ΔL_i, the industry-specific labor share $(1-\alpha_i)$ and μ_i, where μ_i measures the input-output multiplier of industry i. It is given by $\mu_i =$

[3] If these jobs were instead complementary on average, our model will underestimate the negative economic effects of confinement.

Covid Economics 9, 24 April 2020: 110-137

$(I-\Gamma)^{-1}\beta_i$, where $(I-\Gamma)^{-1}$ denotes the Leontief inverse of the input-output matrix Γ and where β_i is the final demand share of industry i.

A typical element of the Leontief inverse can be interpreted as the percentage increase in the output of downstream sector i following a 1% increase in output of upstream sector j. Thus, a typical element μ_i of the resulting vector of IO multipliers reveals how a 1% increase in output of sector i affects value added, both directly and via the impact on the output of other sectors. Thus, sectors with high multipliers either provide inputs to many other industries or they have a high final expenditure share (value added share in GDP).

Multiplying μ_i by the importance of labor in the industry's production, $(1-\alpha_i)$ and summing the effect across sectors tells us how aggregate value added (GDP) changes in response to a 1% shock to labor supply of each industry.

We calibrate the model using an input-output table of the German economy for 2016 (the latest available year) disaggregated into 62 industries from Eurostat. Data on industry-specific labor shares in value added $(1-\alpha_i)$ and employment L_i are also sourced from Eurostat for the same year.

3.2 The Aggregate Economic Costs of Confinement

As a benchmark, we first use our stylized model to quantify the economic costs of an extreme confinement policy where only workers in jobs that require no physical presence remain in the labor force of each industry, i.e. we set the labor force in each sector to one minus the WFH share. Our input-output model for the German economy suggests that such a confinement policy that effectively reduces the labor force by 58% of workers (however, with a large variation across industries) is very costly in terms of output loss. Such a policy causes an overall annualized cost of 80.9% GDP loss translating to a 1.6% loss in GDP for every week with the policy is in place.

Covid Economics 9, 24 April 2020: 110-137

Covid Economics 9, 24 April 2020: 110-137

3.3 Industry-Based Policies

We then use our structural model to ask how strict confinement can be liberalized in a way that minimizes physical presence of workers while maximizing aggregate output. Since industries have different positions in the German production network and thus contribute different amounts of value added to aggregate GDP, we next study the industry-specific multipliers taking the German input-output structure into account. Specifically, we quantify the marginal increase in value added of sending 1% of the sectoral workforce of each industry back to work. Figure 5 shows the sectoral value added multipliers for the ten industries with the highest multipliers.[4] Amongst these are industries providing business services such as legal services, ICT services or finance, which provide key inputs for most other sectors in the economy. Similarly, we find high value added multipliers in construction, public administration, and manufacturing of motor vehicles and machinery.[5] For instance, the multiplier of 0.1 in the sector "Legal, accounting and consulting services" means that an increase of 1% of the workforce in this industry implies a 0.1% increase in GDP. Since some sectors with large value added multipliers might be rather small in the economy in terms of absolute size, while other sectors with relatively small multipliers are large, we use our value added multipliers to evaluate the economic impact of letting an additional worker return to the workplace in terms of their Euro increase in the German GDP. In Figure 6 we show the absolute effects for the industries with the largest values per worker.[6] The industries with the largest level effects on GDP are a mixture of business services (rental and leasing, telecommunications, insurance), supply industries (water, electricity and gas) and manufacturing industries (coke and petroleum products, pharmaceuticals, vehicles and chemicals).[7] In all of these indus-

[4] We provide a full list of industry multipliers in Appendix A.

[5] Barrot et al. (2020) conduct a similar exercise for the production network in France. Also their analysis suggests that business services, construction, public administration and real estate are among the industries where a marginal phasing-out of social distancing has the largest marginal effects on GDP.

[6] We omit real estate services, which has by far the highest value added per worker, as an outlier industry because of measurement problems.

[7] We provide a full list of industry values in Appendix B.

Covid Economics 9, 24 April 2020: 110-137

Figure 5: Largest Industry Value Added Multipliers

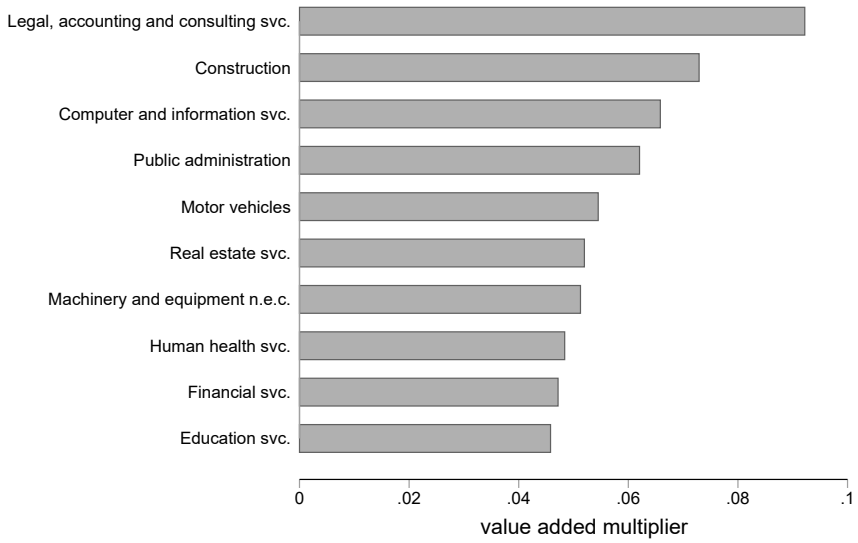

Notes: The Figure displays the 10 largest value added multipliers of individual NACE Rev. 2 industries in the German production network. Multipliers indicate the marginal increase in industry value added by sending an additional 1% of industry employees back to work.

tries, the impact of an individual worker on weekly GDP is substantially above 2 thousand Euros per worker. Hence, letting an additional worker return to the workplace would increase the annualized GDP by more than 100 thousand Euros. These industries are characterized by both, high levels of multipliers and high values of value added per worker. Hence, our model suggests that a policy where industries with a large impact on GDP are granted some priority to phase out confinement rules might help to effectively reactivate the economy while keeping infections at a sufficiently low level.

3.4 Region-Based Policies

While industry-specific policies are the preferred option, since they allow minimizing physical presence in the workplace, they may be difficult to

Covid Economics 9, 24 April 2020: 110-137

Figure 6: Absolute Industry Effects

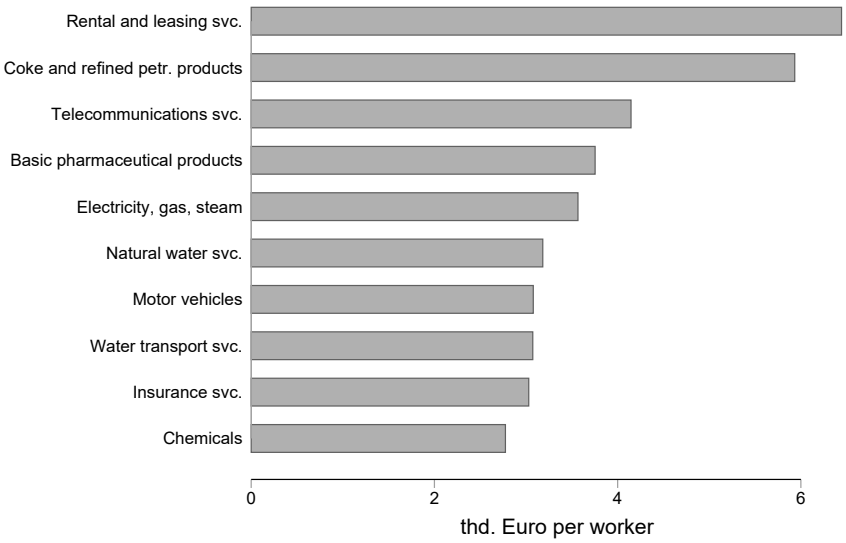

Notes: The Figure displays by how many Euros German weekly GDP would increase by sending an individual worker back to work for the 10 individual NACE Rev. 2 industries with the largest impact, taking into account size differences across industries. We exclude the 'Real estate svc.' (L) industry as an outlier.

implement. We therefore also consider the effect of phasing out confinement by region, without variation in policy by industry. To do so, we take regional differences in industry activity into account. We aggregate the industries across NUTS-2 regions in Germany based on regional employment shares from Eurostat. Assuming a constant ratio of valued added to employment within each sector across German regions, we construct industry value added within each region. As we did before for the individual industries, we now consider the regional weekly change in GDP for each additional worker sent back to work. Figure 7 lists this weekly Euro increase per worker for each individual NUTS-2 region. Metropolitan areas such as Berlin, Hamburg, Düsseldorf or Oberbayern (including Munich) would experience the largest GDP increases per additional worker measured in absolute Euro terms per week. Regional weekly GDP increases per worker are in the range between above 2 thousand and 6 thousand Euro.

Covid Economics 9, 24 April 2020: 110-137

Figure 7: Absolute Regional Effects

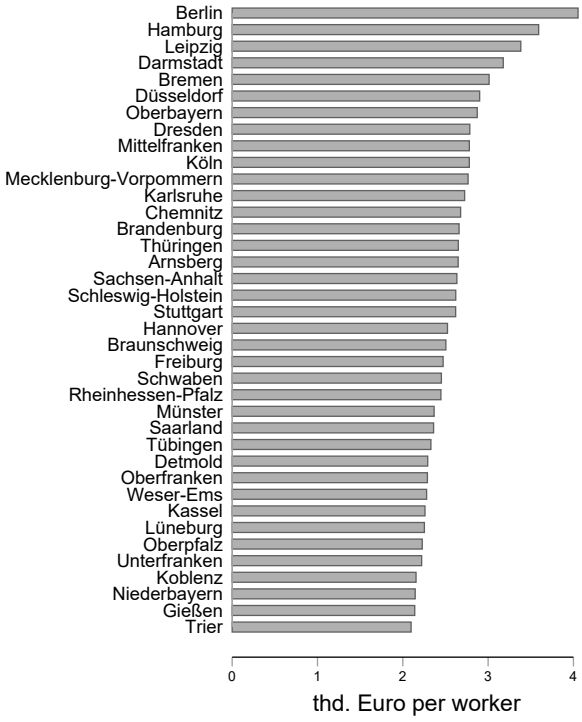

Notes: The Figure displays by how many Euros regional weekly GDP would increase by sending an individual worker back to work for each NUTS-2 region, taking into account size differences across industries.

Next, we ask how regional GDP would be affected when strict confinement is implemented. As we did for the aggregate level of GDP (in subsection 3.2), we consider a change in the labor force from the full regional labor force to WFH workers only. The map in Figure 8 illustrates that the regions that are hurt most from social distancing policies due to their industry structure are mostly concentrated in Southern Germany, in particular in Baden-Württemberg and Bavaria. On the other hand, regions that are hurt the least from confinement because they have a high share of industries where a large share of workers can work from home are regions of former

CEPR PRESS

Figure 8: Regional Changes in Weekly GDP from Phasing Out Confinement

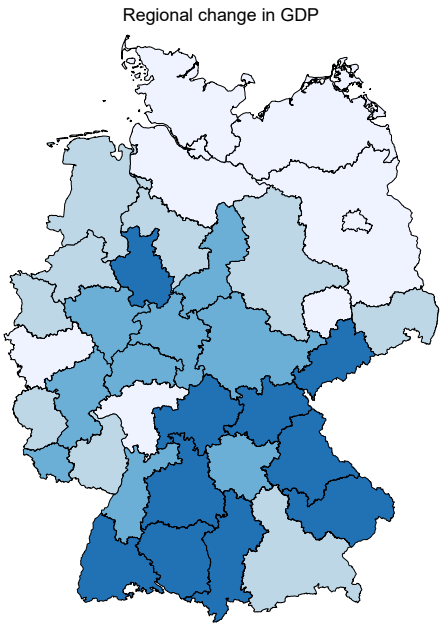

Notes: The Map displays the regional weekly GDP increase in % from giving up strict confinement where only WFH jobs are included in the workforce across NUTS-2 regions in Germany. Regions in darker blue tones have higher multipliers.

Eastern territories of Germany, Northern Germany and Cologne.[8] Overall, the heterogeneity across regions is substantial but the costs are large in all regions. The cost of strict confinement in Berlin, which is the region affected the least, is 1.36% of annual GDP. Tübingen, the region affected the most, experiences instead a weekly GDP loss of 1.77%, about 30% more than in Berlin. Correspondingly, the regions where losses from confinement are largest would also gain the most from reducing it.

Lastly, we consider the relation between WFH jobs in a region and how much a region would suffer in economic terms from confinement in Figure 9. As expected, there is a strong negative relation between both measures. Having a ten percentage point higher fraction of jobs that can be done at home is associated with having a 0.4% smaller weekly GDP reduction from

[8] We list the effects on all regions in Appendix C.

Covid Economics 9, 24 April 2020: 110-137

CEPR PRESS

Figure 9: Weekly Output Changes from Giving Up Confinement and Working from Home

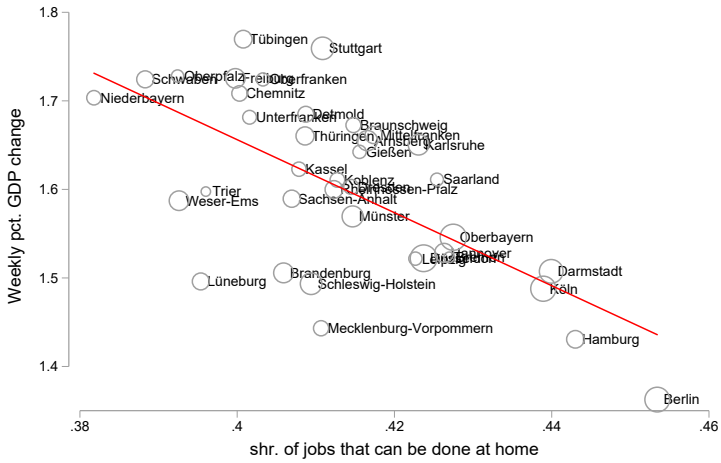

Notes: The Figure displays a scatterplot of the share of jobs that can be done from home and the regional weekly GDP increase in % from giving up strict confinement where only WFH jobs are included in the workforce across NUTS-2 regions in Germany. Individual dots are population-weighted.

strict confinement. Those regions that lose the most substantial amounts of their workforce due to the introduction of strict social distancing rules are also those regions that hurt most from confinement. The variation of this labor supply explains almost 42% of the variation in the weekly GDP change. The remaining 58% can then be explained by differences in the regional sectoral structure and the input-output effects.

4 Conclusion

We have discussed the impact of social distancing on Covid-19 infections on the one hand, and the German economy on the other hand. While social distancing is very effective in reducing infection rates, it also imposes substantial economic costs on the economy. We discuss different policies to reactivate the German economy, while keeping as many workers at home as possible. We have identified those systemic industries and regions that ben-

Covid Economics 9, 24 April 2020: 110-137

efit most from lifting confinement rules on the basis of a stylized structural model. We find that the industries where lifting confinement has the largest level effects on GDP are a mixture of business services such as telecommunications or insurances, supplier industries such as water, electricity and gas supply and manufacturing industries such as petroleum products, pharmaceuticals, vehicles and chemicals. Furthermore, the regions potentially benefitting economically the most from loosening confinement are mostly concentrated in Southern Germany, in particular in Baden-Württemberg and Bavaria, which are the same regions where working from home is difficult due to their industry structure.

Covid Economics 9, 24 April 2020: 110-137

References

Barrot, Jean-Noël, Basile Grassi, and Julien Sauvagnat. Sectoral Effects of Social Distancing. mimeo, Bocconi University, 2020.

Dingel, Jonathan and Brent Neiman. How Many Jobs Can be Done at Home? mimeo, University of Chicago, 2020.

Hartl, Tobias, Klaus Wälde, and Enzo Weber. Measuring the Impact of the German Public Shutdown on the Spread of Covid-19. *Covid Economics*, 1:25–32, 2020.

Hethcote, Herbert W. *Three Basic Epidemiological Models*, pages 119–144. Springer, 1989.

Jones, Charles I. Misallocation, Economic Growth, and Input-Output Economics. volume 2 of *Advances in Economics and Econometrics, Tenth World Congress*, chapter 10, pages 419–454. Cambridge University Press, 2013.

Koren, Miklós and Rita Peto. Business Disruptions from Social Distancing. *Covid Economics*, 2:13–31, 2020.

Covid Economics 9, 24 April 2020: 110-137

Appendix to "The Costs and Benefits of Home Office During the Covid-19 Pandemic"

A Full List of Value Added Multipliers

Industry	Description	Value added multiplier
M69 - 70	Legal and accounting svc.; svc. of head offices; management consultancy svc.	0.0922
F	Constructions and construction works	0.0729
J62 - 63	Computer programming, consultancy and related svc.;Information svc.	0.0658
O	Public administration and defence svc.; compulsory social security svc.	0.0620
C29	Motor vehicles, trailers and semi-trailers	0.0545
N80 - 82	Security and investigation svc.; svc. to buildings and landscape; office administrative, office support and other business support svc.	0.0520
C28	Machinery and equipment n.e.c.	0.0513
Q86	Human health svc.	0.0484
K64	Financial svc., except insurance and pension funding	0.0472
P	Education svc.	0.0458
G46	Wholesale trade svc., except of motor vehicles and motorcycles	0.0454
C25	Fabricated metal products, except machinery and equipment	0.0453
H52	Warehousing and support svc. for transportation	0.0421
C20	Chemicals and chemical products	0.0416
C24	Basic metals	0.0357
H49	Land transport svc. and transport svc. via pipelines	0.0345
C10 - 12	Food, beverages and tobacco products	0.0327
G47	Retail trade svc., except of motor vehicles and motorcycles	0.0312
M71	Architectural and engineering svc.; technical testing and analysis svc.	0.0295
C27	Electrical equipment	0.0293

Covid Economics 9, 24 April 2020: 110-137

Covid Economics 9, 24 April 2020: 110-137

Industry	Description	Value added multiplier
C22	Rubber and plastic products	0.0250
D	Electricity, gas, steam and air conditioning	0.0247
B	Mining and quarrying	0.0237
G45	Wholesale and retail trade and repair svc. of motor vehicles and motorcycles	0.0236
C26	Computer, electronic and optical products	0.0228
N78	Employment svc.	0.0223
I	Accommodation and food svc.	0.0198
K65	Insurance, reinsurance and pension funding svc., except compulsory social security	0.0193
Q87 - 88	Residential care svc.; social work svc. without accommodation	0.0193
H53	Postal and courier svc.	0.0189
C33	Repair and installation svc. of machinery and equipment	0.0177
J58	Publishing svc.	0.0168
C23	Other non-metallic mineral products	0.0167
C19	Coke and refined petroleum products	0.0158
M74 - 75	Other professional, scientific and technical svc. and veterinary svc.	0.0157
K66	Svc. auxiliary to financial svc. and insurance svc.	0.0149
C17	Paper and paper products	0.0145
C31 - 32	Furniture and other manufactured goods	0.0140
S96	Other personal svc.	0.0134
L	Real estate activities	0.0131
J61	Telecommunications svc.	0.0128
M73	Advertising and market research svc.	0.0127
S94	svc. furnished by membership organisations	0.0127
E37 - 39	Sewerage svc.; sewage sludge; waste collection, treatment and disposal svc.; materials recovery svc.; remediation svc. and other waste management svc.	0.0126
C30	Other transport equipment	0.0102
C16	Wood and of products of wood and cork, except furniture; articles of straw and plaiting materials	0.0092
C21	Basic pharmaceutical products and pharmaceutical preparations	0.0085
C18	Printing and recording svc.	0.0080

Industry	Description	Value added multiplier
A01	Products of agriculture, hunting and related svc.	0.0074
R90 - 92	Creative, arts, entertainment, library, archive, museum, other cultural svc.; gambling and betting svc.	0.0070
J59 - 60	Motion picture, video and television programme production svc., sound recording and music publishing; programming and broadcasting svc.	0.0069
R93	Sporting svc. and amusement and recreation svc.	0.0064
C13 - 15	Textiles, wearing apparel, leather and related products	0.0060
H51	Air transport svc.	0.0055
M72	Scientific research and development svc.	0.0047
N77	Rental and leasing svc.	0.0041
N79	Travel agency, tour operator and other reservation svc. and related svc.	0.0034
A02	Products of forestry, logging and related svc.	0.0022
E36	Natural water; water treatment and supply svc.	0.0015
S95	Repair svc. of computers and personal and household goods	0.0013
H50	Water transport svc.	0.0006
A03	Fish and other fishing products; aquaculture products; support svc. to fishing	0.0002

Covid Economics 9, 24 April 2020: 110-137

B Full List of Absolute Industry Effects

Industry	Description	Impact on GDP (thd. Euro)
I	Accommodation and food svc.	24.6
Q87 - 88	Residential care svc.; social work svc. without accommodation	24.7
S95	Repair svc. of computers and personal and household goods	28.4
N78	Employment svc.	30.4
G47	Retail trade svc., except of motor vehicles and motorcycles	31.1
A01	Products of agriculture, hunting and related svc.	31.5
N80 - 82	Security and investigation svc.; svc. to buildings and landscape; office administrative, office support and other business support svc.	32.0
H53	Postal and courier svc.	34.0
S94	svc. furnished by membership organisations	43.2
S96	Other personal svc.	45.7
C18	Printing and recording svc.	46.6
K66	Svc. auxiliary to financial svc. and insurance svc.	49.0
Q86	Human health svc.	50.1
A03	Fish and other fishing products; aquaculture products; support svc. to fishing	50.2
C16	Wood and of products of wood and cork, except furniture; articles of straw and plaiting materials	50.8
C10 - 12	Food, beverages and tobacco products	51.3
M73	Advertising and market research svc.	52.3
C13 - 15	Textiles, wearing apparel, leather and related products	52.7
R90 - 92	Creative, arts, entertainment, library, archive, museum, other cultural svc.; gambling and betting svc.	53.1
H49	Land transport svc. and transport svc. via pipelines	54.4
P	Education svc.	54.8
M71	Architectural and engineering svc.; technical testing and analysis svc.	57.5
G45	Wholesale and retail trade and repair svc. of motor vehicles and motorcycles	58.1

Covid Economics 9, 24 April 2020: 110-137

Covid Economics 9, 24 April 2020: 110-137

Industry	Description	Impact on GDP (thd. Euro)
F	Constructions and construction works	58.2
C31 - 32	Furniture and other manufactured goods	59.6
N79	Travel agency, tour operator and other reservation svc. and related svc.	60.9
R93	Sporting svc. and amusement and recreation svc.	62.5
C25	Fabricated metal products, except machinery and equipment	64.4
C33	Repair and installation svc. of machinery and equipment	66.4
M69 - 70	Legal and accounting svc.; svc. of head offices; management consultancy svc.	67.9
H52	Warehousing and support svc. for transportation	68.6
C22	Rubber and plastic products	70.2
O	Public administration and defence svc.; compulsory social security svc.	70.3
M74 - 75	Other professional, scientific and technical svc. and veterinary svc.	70.3
C23	Other non-metallic mineral products	74.7
B	Mining and quarrying	75.8
G46	Wholesale trade svc., except of motor vehicles and motorcycles	76.9
J58	Publishing svc.	77.0
C17	Paper and paper products	77.6
C24	Basic metals	80.5
A02	Products of forestry, logging and related svc.	87.4
C28	Machinery and equipment n.e.c.	89.2
C27	Electrical equipment	90.1
J62 - 63	Computer programming, consultancy and related svc.;Information svc.	97.2
E37 - 39	Sewerage svc.; sewage sludge; waste collection, treatment and disposal svc.; materials recovery svc.; remediation svc. and other waste management svc.	104.0
C30	Other transport equipment	105.4
M72	Scientific research and development svc.	109.2
C26	Computer, electronic and optical products	114.7
K64	Financial svc., except insurance and pension funding	121.3
H51	Air transport svc.	127.0

Industry	Description	Impact on GDP (thd. Euro)
J59 - 60	Motion picture, video and television programme production svc., sound recording and music publishing; programming and broadcasting svc.	128.7
C20	Chemicals and chemical products	144.3
K65	Insurance, reinsurance and pension funding svc., except compulsory social security	157.5
H50	Water transport svc.	159.8
C29	Motor vehicles, trailers and semi-trailers	160.1
E36	Natural water; water treatment and supply svc.	165.5
D	Electricity, gas, steam and air conditioning	185.5
C21	Basic pharmaceutical products and pharmaceutical preparations	195.2
J61	Telecommunications svc.	215.6
C19	Coke and refined petroleum products	308.5
N77	Rental and leasing svc.	335.1
L	Real estate activities	662.3

Covid Economics 9, 24 April 2020: 110-137

CEPR PRESS

C Full List of Regional Effects of Lifting Confinement

Region	Weekly change in GDP
Tübingen	1.77%
Stuttgart	1.76%
Oberpfalz	1.73%
Freiburg	1.73%
Schwaben	1.72%
Oberfranken	1.72%
Chemnitz	1.71%
Niederbayern	1.70%
Detmold	1.68%
Unterfranken	1.68%
Braunschweig	1.67%
Mittelfranken	1.66%
Thüringen	1.66%
Arnsberg	1.66%
Karlsruhe	1.65%
Gießen	1.64%
Kassel	1.62%
Saarland	1.61%
Koblenz	1.61%
Dresden	1.60%
Rheinhessen-Pfalz	1.60%
Trier	1.60%
Sachsen-Anhalt	1.59%
Weser-Ems	1.59%
Münster	1.57%
Oberbayern	1.55%
Hannover	1.53%
Bremen	1.52%
Düsseldorf	1.52%
Leipzig	1.52%
Darmstadt	1.51%
Brandenburg	1.51%
Lüneburg	1.50%
Schleswig-Holstein	1.49%
Köln	1.49%
Mecklenburg-Vorpommern	1.44%

Covid Economics 9, 24 April 2020: 110-137

Region	Weekly change in GDP
Hamburg	1.43%
Berlin	1.36%

Covid Economics 9, 24 April 2020: 110-137

Covid Economics 9, 24 April 2020: 110-137

D The German Input-Output Network