CENTRE FOR ECONOMIC POLICY RESEARCH

CEPR PRESS

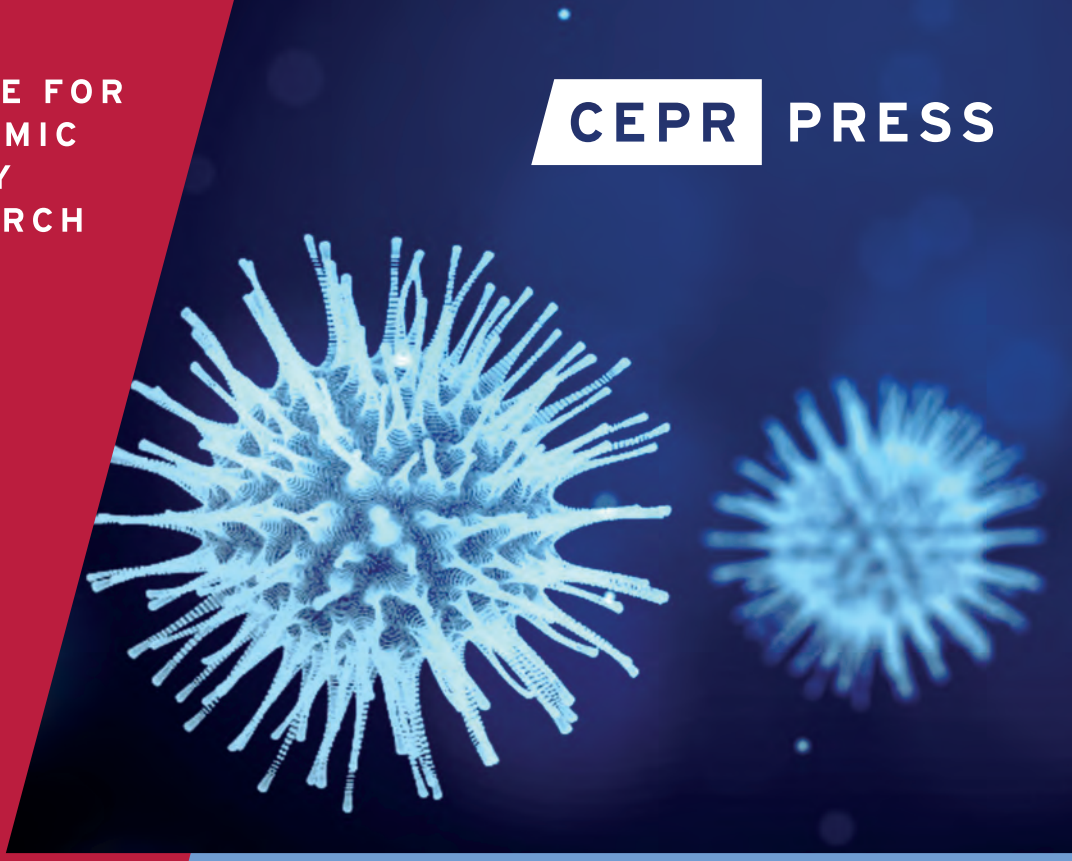

COVID ECONOMICS

VETTED AND REAL-TIME PAPERS

ISSUE 35
7 JULY 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. Other queries should be sent to covidecon@cepr.org.

The Centre for Economic Policy Research (CEPR)

The Centre for Economic Policy Research (CEPR) is a network of over 1,500 research economists based mostly in European universities. The Centre's goal is twofold: to promote world-class research, and to get the policy-relevant results into the hands of key decision-makers. CEPR's guiding principle is 'Research excellence with policy relevance'. A registered charity since it was founded in 1983, CEPR is independent of all public and private interest groups. It takes no institutional stand on economic policy matters and its core funding comes from its Institutional Members and sales of publications. Because it draws on such a large network of researchers, its output reflects a broad spectrum of individual viewpoints as well as perspectives drawn from civil society. CEPR research may include views on policy, but the Trustees of the Centre do not give prior review to its publications. The opinions expressed in this report are those of the authors and not those of CEPR.

Editorial Board

Ethics

Covid Economics will feature 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 articles.

Submission to professional journals

The following journals have indicated that they will accept submissions of papers featured in *Covid Economics* because they are working papers. Most expect revised versions. This list will be updated regularly.

American Economic Review

American Economic Review, Applied Economics

American Economic Review, Insights

American Economic Review, Economic Policy

American Economic Review, Macroeconomics

American Economic Review, Microeconomics

American Journal of Health Economics

Canadian Journal of Economics

Economic Journal

Economics of Disasters and Climate Change

International Economic Review

Journal of Development Economics

*Journal of Econometrics**

Journal of Economic Growth

Journal of Economic Theory

*Journal of the European Economic Association**

Journal of Finance

Journal of Financial Economics

Journal of International Economics

*Journal of Labor Economics**

Journal of Monetary Economics

Journal of Public Economics

Journal of Political Economy

Journal of Population Economics

*Quarterly Journal of Economics**

Review of Economics and Statistics

*Review of Economic Studies**

Review of Financial Studies

(*) Must be a significantly revised and extended version of the paper featured in *Covid Economics*.

Covid Economics
Vetted and Real-Time Papers

Issue 35, 7 July 2020

Contents

Covid Economics Issue 35, 7 July 2020

Hesitancy towards a COVID-19 vaccine and prospects for herd immunity[1]

Linda Thunström,[2] Madison Ashworth,[3] David Finnoff[4] and Stephen C. Newbold[5]

Date submitted: 2 July 2020; Date accepted: 3 July 2020

The scientific community has come together in an unprecedented effort to find a COVID-19 vaccine. However, the success of any vaccine depends on the share of the population that gets vaccinated. We design a survey experiment in which a nationally representative sample of 3,133 adults in the U.S. state their intentions to vaccinate themselves and their children for COVID-19. In the experiment, we account for uncertainty about the severity and infectiousness of COVID-19, as well as inconsistencies in risk communication from government authorities, by varying these factors across treatments. We find that 20% of people in the U.S. would decline the vaccine. General vaccine hesitancy (including not having had a flu shot in the last two years), distrust of vaccine safety, and vaccine novelty are among the most important deterrents to vaccination. Further, our results suggest that inconsistent risk messages from public health experts and elected officials reduce vaccine uptake. We use our survey results in an epidemiological model to explore conditions under which a vaccine could achieve herd immunity. We find that in a middle-of-the-road scenario with central estimates of model parameters, a vaccine will benefit public health by saving many lives but nevertheless may fail to achieve herd immunity.

1 We thank the College of Business Excellence Fund at the University of Wyoming. Authors declare no competing interests. All data and code are available upon request.
2 Associate professor, Department of Economics, University of Wyoming.
3 Ph.D. student, Department of Economics, University of Wyoming.
4 Professor, Department of Economics, University of Wyoming.
5 Assistant professor, Department of Economics, University of Wyoming.

Covid Economics 35, 7 July 2020: 1-50

Covid Economics 35, 7 July 2020: 1-50

1. Introduction

In an unprecedented effort, scientists from all over the world have come together to rapidly develop a vaccine for COVID-19 (Callaway, 2020). Vaccines have historically proven to be highly successful and cost-effective public health tools for disease prevention (Rémy et al., 2015), and already by April 2020 more than 100 COVID-19 vaccine candidates had been developed, several of which have advanced to being tested on humans (Le et al., 2020).

However, the effectiveness of a vaccine in controlling the spread of COVID-19 depends on the uptake level of the vaccine in the population. A sufficiently high uptake of an effective vaccine ensures protection for those vaccinated and may end the pandemic by generating herd immunity, thereby protecting everyone, including those still susceptible to the virus (Fine et al., 2011). For instance, Kwok et al. (2020) suggest herd immunity is reached when 69.6% of the United States' population has immunity either from a vaccine or previous infections. Recent estimates based on data from outbreaks in China suggest the basic reproductive number for COVID-19, R_0, might be higher than previously thought (as high as 5.7, see Sanche et al., 2020), which implies herd immunity may be reached first when 82.5% of the population is immune (Keeling and Rohani, 2008).

A barrier to reaching herd immunity is the prevalence of people who refuse or are hesitant to take vaccines (MacDonald, 2015). In the U.S., this share of the population has grown in recent years (Dube et al., 2013; Olive et al., 2018). For instance, the uptake level of seasonal influenza vaccines has declined, in part due to vaccine hesitancy (Larson, 2018). During the season 2017-2018 only 37% of adults got the flu vaccine, even though that flu season was particularly severe (CDC, 2020a). Recent measles outbreaks in the U.S. and elsewhere illustrate the importance of

Covid Economics 35, 7 July 2020: 1-50

vaccine hesitancy to public health, as the vaccine had succeeded in extinguishing measles in the U.S. but under and non-vaccinated communities contributed to its reappearance (De Serres et al., 2013; Sarkar et al., 2019). The World Health Organization (WHO) named vaccine hesitancy one of the top 10 threats to global health in 2019 (WHO, 2020a).

We examine the prevalence and determinants of avoidance of a COVID-19 vaccine in the U.S. and how the anticipated level of vaccine avoidance affects the ability of a vaccine to generate herd immunity. We also examine how vaccine avoidance is affected by information about health risks associated with SARS-CoV-2 (the novel coronavirus), including conflicting risk messages from public authorities. Risk perceptions have been shown to be key to vaccine decisions (Brewer et al., 2007) and effective risk communication is acknowledged as the pillar of a coordinated response to infectious disease outbreaks (Sell, 2017). Observational data suggests that risk information provided in the beginning of the COVID-19 outbreak in the U.S. affected people's health behavior in response to the pandemic (Bursztyn et al., 2020; Simonov et al., 2020). We examine whether inconsistent risk messages may similarly affect health behavior—in our case willingness to vaccinate.

To measure vaccine avoidance and how it depends on risk levels and risk communication, we design a survey experiment that elicits vaccination intentions of adults and their children. Participants consist of a nationally representative sample (N=3,133) of U.S. adults. They are randomized into eight treatments across which we vary the probability of infection, the conditional mortality rate from COVID-19, and whether the different health authorities in the U.S. provide consistent risk information. Specifically, we examine how vaccine avoidance is affected by elected or appointed White House officials communicating lower risks from COVID-19 than public health experts at the Centers for Disease Control (CDC). The lower level of risk

communicated by the White House, compared to that of the CDC, is in line with reporting in popular media at the onset of the current pandemic (e.g., CNN, 2020; MSN, 2020).[1]

Next, we use an epidemiological model with susceptible, infectious, and recovered (SIR) compartments of the population to examine how vaccine avoidance affects the ability of the vaccine to save lives and generate herd immunity in an upcoming COVID-19 season. Given uncertainty about the value of other key parameters that determine the performance of a vaccine program, we examine vaccine success across ranges of vaccine hesitancy, vaccine effectiveness in preventing infection for those vaccinated, immunity in the population, and the coronavirus basic reproduction number at the onset of the season.

While vaccine hesitancy is growing, hesitancy is not equivalent to refusal—many people who are vaccine hesitant do not entirely refuse vaccines. Instead, they either delay vaccines or are willing to take some vaccines but not others (Dube et al., 2013). Also relevant for a COVID-19 vaccine is the observation that people are more likely to reject new vaccines than familiar ones (Dube et al., 2013). A U.S.-wide study found that around 10% of the population refuse all vaccines, including seasonal influenza vaccines and those that comprise the recommended vaccine schedule for children, while around 5% refuse only one vaccine. A substantial share (40%) of those who agreed to at least one vaccine still expressed concerns about vaccines (ASTHO, 2010). The fact that many who are vaccine hesitant are likely to take some vaccines, while perhaps delayed, means it is possible that a portion of those currently reluctant to vaccinate

[1] Our study relates to a rich body of literature on consumer responses to conflicting information, spanning multiple scientific disciplines. While not an exhaustive list, examples of important work in this area are Viscusi and Magat (1992), Magat and Viscusi (1992), Viscusi (1997), Viscusi et al. (1998), Rodgers (1999), Fox et al. (2002), Hoehn and Randall (2002), Cameron (2005), Rousu and Shogren (2006), Kelly et al. (2012), Carpenter et al. (2014) Hämeen-Anttila et al. (2014), Pushkarskaya et al. (2015), Binder et al. (2016).

Covid Economics 35, 7 July 2020: 1-50

can be swayed (Leask, 2011). Here, risk communication may play an important role, given the correlation between perceived risks and vaccine acceptance (Brewer et al., 2007).

Our study provides important knowledge at a critical point in time. Measuring the share of the population that is reluctant to be vaccinated for COVID-19 can help policy makers, health care workers and other authorities to plan ahead towards minimizing the impact on public health from vaccine hesitancy. This might involve tailored public communication programs designed to persuade vaccine hesitant individuals to accept a COVID-19 vaccine, or increased efforts to ensure a high vaccine uptake level among the remainder of the population, or both. Knowing why people are hesitant to accept a COVID-19 vaccine may enable design of more effective efforts to increase the overall level of vaccine uptake in the general population.

The remainder of this paper is structured as follows. Section 2 describes the survey experiment and the SIR model. Section 3 presents the results of the experiment and the projected impact of vaccine avoidance on herd immunity and lives saved. Section 4 discusses the results, limitations of the current study, and avenues for future research on avoidance of a COVID-19 vaccine.

Covid Economics 35, 7 July 2020: 1-50

2. Methods

2.1. Survey to examine vaccine intentions and their determinants

To examine people's willingness to vaccinate for COVID-19, we design a survey experiment in which participants were asked whether they would choose to vaccinate themselves or their children. The survey experiment was approved by the IRB at University of Wyoming and was pre-registered in the AEA RCT registry (RCT ID: AEARCTR-0005576).

Covid Economics 35, 7 July 2020: 1-50

We use a between-subjects experimental design with eight information treatments (2×2×2). The experiment varies information on (1) the probability of the average American catching the coronavirus, (2) the conditional mortality rate, i.e., the probability of the average American dying if infected, and (3) the source of information for the probability of catching COVID-19 (CDC only/CDC jointly with the White House).

While there are clear benefits to measuring the prevalence and implications of vaccine avoidance at this point in time, before a vaccine is available, doing so comes with methodological challenges—currently, the true probability of infection in the U.S. and mortality rate from COVID-19 are unknown, due to limited testing. Further, these parameters might have changed by the time a vaccine is available. Any study that measures the prevalence of vaccine avoidance before health risks from COVID-19 are fully known needs to consider how this uncertainty might affect the measured prevalence of vaccine avoidance, given the important role of health risks to vaccine decisions (see e.g., Brewer et al., 2007). We deal with the uncertainty by communicating different levels of risk across treatments. Further, previous studies suggest beliefs about health risk are affected by *who* communicates the risk message (e.g., Frewer et al., 1996; Breakwell, 2000; Calman and Curtis, 2010). We focus on risk messages delivered by public health experts (i.e., the CDC) and White House officials, given they regularly address the pandemic in the U.S. media. The public health risks from COVID-19 are still highly uncertain and evolve with changing policies and individual behavioral responses to the outbreak. For this reason, our information treatments entail presentations of hypothetical scenarios to participants, as communicated to participants prior to the treatments (Step 2 below). While the contrasting risk information treatments presented to our survey subjects were within the plausible range of outcomes as understood at the time the survey was administered, no single information treatment

can be said to represent the "true" risks. Further, the information treatments attribute statements of risk made by the CDC or White House officials that are not direct quotes but rather are paraphrased summaries designed to provide a clean contrast between the severity of the risk communicated by the two sources, capturing the observation that the White House generally communicates lower risk (CNN, 2020; MSN, 2020).

If participants are treated with information about the probability of infection from the CDC only (the first four treatments to the left in Figure 1), then there is no discrepancy in the risk information presented to participants, and high probability of infection=85%, while low probability of infection=25%. If they are treated with information about the probability from both the CDC and the White House (the four treatments to the right in Figure 1), the messages from these sources are inconsistent. In these treatments, the probabilities communicated by the CDC are complemented by lower probabilities from the White House (the White House's high probability=70%; the White House's low probability=10%).[2] Hence, while the probabilities communicated by the White House also vary between high and low, they are consistently lower (more optimistic) than the probabilities communicated by the CDC).

The source of the treatment information on the conditional mortality rate of the disease is stated to be "medical experts" and high mortality=10%, while low mortality=1.5%. Shereen et al. (2020) estimate a conditional mortality rate, across 109 countries, for those infected by the virus at around 3%, and Cascella et al. (2020) report a mortality rate of 1-2%, across multiple studies.

[2] Note that we keep the disparity in risk communicated by the CDC and the White House constant (at 15 percentage points) across both high and low infection risk treatments. Viscusi (1997) shows that the disparity in the risks communicated by different information sources may affect trust in all information sources, such that a change in the disparity in percentage point probabilities across high and low infection risk treatments could have affected trust in both the CDC and the White House.

Covid Economics 35, 7 July 2020: 1-50

Covid Economics 35, 7 July 2020: 1-50

While the true conditional mortality rate for the U.S. is unknown due to limited and imperfect testing for COVID-19, some of the uncertainty about the conditional mortality rate under current health care conditions has been resolved since data collection for this study -- on June 26 2020, Johns Hopkins University and Medicine (2020) and Roser et al. (2020) reported a U.S. case fatality rate of 5.1%. The case fatality rate is higher than the true conditional mortality rate, given it is the ratio of confirmed COVID-19 caused deaths and confirmed COVID-19 cases, and actual cases are higher than confirmed cases.

Participants in our survey experiment were assembled by the survey company Qualtrics, who was instructed to recruit a sample of 3,000 survey respondents who are representative of the U.S. general population in gender, age, income, education, race, and residential region. Due to oversampling by Qualtrics, our total sample size is $N = 3{,}133$. The advantage of using Qualtrics over less costly alternatives, such as Amazon Mechanical Turk or Turk Prime, is that Qualtrics continuously performs quality checks of their participants, including with regards to background characteristics and screens for professional survey takers, which are otherwise known to contaminate online panels (e.g., see Chandler and Paolacci, 2017, and Sharpe Wessling et al., 2017). Participants received standard compensation for completing a Qualtrics survey. An additional benefit was that Qualtrics could offer rapid data collection, which was important given the information flow on COVID-19 that participants were exposed to outside of our study. Approximately 80% of the data was collected between March 24 and March 31, 2020.

The sequence of the study was as follows:

Step 1: Participants were asked questions about their gender, age, education, race, income and region of residence, to ensure the sample met U.S. national quotas for those characteristics.

Step 2: All participants received the following information about COVID-19:

Coronaviruses (CoV) are a large family of viruses that cause illness ranging from the common cold to more severe diseases.

Common signs of infection include respiratory symptoms, fever, cough, shortness of breath and breathing difficulties. In more severe cases, infection can cause pneumonia, severe acute respiratory syndrome, kidney failure and even death.

The new coronavirus (COVID-19) is still spreading globally, meaning that the risks to average Americans of catching the disease (currently or in a near future) are still uncertain, as are the risks of developing symptoms severe enough to cause deaths. In this study, we will present you with plausible estimates of these risks, based on recent knowledge of the virus and associated risks.

Public health risks caused by infectious diseases are often communicated by the Centers for Disease Control and Prevention (CDC). The CDC was established in 1946 and is the leading national public health institute in the U.S. It is a federal agency under the Department of Health and Human Services. CDC's goal is to protect public health and safety through the control and prevention of disease, injury, and disability.

Step 3: Participants were randomized into one of the eight information treatments.

If randomized into one the four treatments with CDC information only, participants saw the following statement:

The Centers for Disease Control and Prevention (CDC) has estimated that the probability of catching the coronavirus in the next 12 months is 25 [85] percent for the average American, meaning that 25 [85] out of 100 Americans are expected to catch the coronavirus.

Medical scientists have estimated that 1.5 [10] percent of Americans who catch the coronavirus will experience severe consequences leading to death, meaning that 15 [100] out of 1000 Americans who catch the virus are expected to die.

If instead randomized into one of the four treatments with both CDC and White House information, participants saw the following statement:

The Centers for Disease Control and Prevention (CDC) has estimated that the probability of catching the coronavirus in the next 12 months is 25 [85] percent for the average American. In other words, the CDC estimates that 25 [85] in 100 Americans will catch the virus.

The White House has indicated that the probability of catching the coronavirus in the next 12 months is lower, namely 10 [70] percent for the average American. In other words, the White house predicts that 10 [70] in 100 Americans will catch the virus.

Covid Economics 35, 7 July 2020: 1-50

Medical scientists have estimated that 1.5 [10] percent of Americans who catch the coronavirus will experience severe consequences leading to death. In other words, 15 [100] out of 1000 Americans who catch the virus are expected to die.

Step 4: All participants were asked to indicate their beliefs about the probability that they and their children (if they had children) will catch COVID-19 (higher/lower/same as the probability for the average American) and the probability that the average American will catch COVID-19. Similarly, they were asked about the conditional mortality risk for themselves and their children if they were to become infected (higher/lower/same as the probability for the average American), as well as the conditional mortality risk of the average American.

Step 5: All participants were asked whether they would vaccinate for COVID-19. Before the vaccine question, they were given additional information on the risks and benefits of the vaccine. Because many people worry about vaccine side-effects, and because those worries might be elevated when a vaccine is produced in a relatively short amount of time, we included information about the vaccine being approved by the Food and Drug Administration (FDA), following standard protocols. Specifically, participants were given the following information and question about whether they would take the vaccine:

Numerous pharmaceutical companies are working to develop a vaccine against the coronavirus. Before any vaccine can be provided to the public, the United States Food and Drug Administration (FDA) must approve its use. The FDA grants approval only if the vaccine is manufactured in compliance with all current regulations and medical scientists find that the vaccine is effective and has minimal side effects.

Suppose that the vaccine was approved for use by the FDA and was available today from your health care provider for free.

Also suppose that the vaccine is as effective as the flu vaccine in an average year, which is about 60 percent. In other words, 60 out of 100 people who are vaccinated would be protected from the coronavirus.

Covid Economics 35, 7 July 2020: 1-50

Would anyone in your family get the coronavirus vaccine under the conditions described above?

Participants were asked to indicate *WOULD get vaccinated* or *WOULD NOT get vaccinated*, for themselves. If they were parents of minors, they were also asked whether they would vaccinate their child, if they had one child, and whether they would vaccinate their youngest and oldest child, if they had multiple children. If they indicated one or more family members would not get vaccinated, they were presented with a series of follow-up questions designed to investigate the reasons for their choice.[3]

Step 6: Participants were asked about behavior undertaken to protect themselves from COVID-19 (hand washing, avoidance of crowds and public spaces, etc), if they (or their children) had received a flu shot in the last 2 years, if they were vaccinated for measles, if they generally followed the recommended immunization schedule for children, and questions underlying the psychological scale for vaccinations (for measles and flu vaccines) developed by Betsch et al. (2018).

Step 7: Participants were asked about their information sources on COVID-19 (family, friends, conservative media, liberal media, family physician, President Trump, etc), and their views about the trustworthiness of a variety of information sources.

Covid Economics 35, 7 July 2020: 1-50

[3] Note that we asked participants to suppose that a vaccine that was available *today*, although we expected participants to understand that a vaccine was in fact not yet available. An alternative would have been to ask about intentions to vaccinate at a future point in time, when a vaccine is more likely to be available. Our choice is based on control over the experiment environment. Participants may differ in their beliefs about when a vaccine might be available and how the risks of infection and death may evolve over the course of the outbreak, they might expect the pandemic to have concluded before a vaccine is available, herd immunity to be near, or that they personally will already have been infected. The recent polls that have measured COVID-19 vaccine hesitancy (see discussion in Section 4) vary in how they have dealt with the timing of the vaccine when asking about the willingness to vaccinate. Like our study, the poll by Pew Research Center (2020) asks about vaccine intentions if the vaccine was available today, while the polls by ABC news/Ipsos (2020) and LX/Morning Consult (2020) ask about willingness to vaccinate when a vaccine becomes available without specifying when that might be.

Step 8: Participants were asked about underlying health conditions that would put them at higher risk for severe consequences if they were to develop COVID-19 (e.g., respiratory disease, cardiovascular disease, obesity, diabetes, cancer, etc), see CDC (2020). They were also asked about risk factors for contracting the virus (being a health care worker, living in an urban area, etc).

Step 9: Participants were asked about their religious beliefs, questions underlying the social and fiscal conservatism scale developed in Everett (2013), and their views about the currently implemented social distancing measures in the U.S.

The full experimental survey can be found in the Supplementary Online Material.

Fifty-two percent of participants are female, and the mean age in our sample is 46 years. Fifty-seven percent of participants have a minimum of some college education. Twenty-five percent of participants fall into the low income category ($24,999 per year or less), 55 percent into the medium income category ($25,000-$99,999 per year), and 20% into the high income category ($100,000 per year or more). About 37% of participants identify as Republican, 41% as Democrat, and 22% identify with neither political party. Nearly 37% of our participants live in a rural area, and 81% believe in God. Around 55% of our sample had the flu shot in the last 2 years, and participants were nearly evenly split among the low, middle, and high levels of trust in government agencies. Of adult participants in our sample, 82% have followed the recommended vaccination schedule, and of participants with children, 86% have followed the recommended vaccination schedule for their children. The Supplementary Online Appendix includes a table with a more complete set of descriptive statistics as well as the variables included in the analysis reported in Table 2.

Covid Economics 35, 7 July 2020: 1-50

Covid Economics 35, 7 July 2020: 1-50

We examine the effect of the treatment information on perceived risk of COVID-19 as well as the vaccination decision by conducting single pairwise statistical tests across treatments (see Figure 1). Our ability to identify the treatment effects by excluding other control variables in the statistical analysis relies on the assumption that the randomization of participants into treatments was successful in eliminating any meaningful differences in relevant covariates across treatments. If the randomization is unsuccessful, such that the value of covariates that are a priori expected to impact the outcome variable (e.g., general vaccine hesitancy) differ across treatment groups, then Athey and Imbens (2017) and Mutz et al. (2017) argue that it is appropriate to control for those covariates in the statistical analysis. Therefore we also examine whether relevant covariates (those included in Table 2 that we expected not to be affected by the treatments) differ across treatments in meaningful ways. We follow Imbens and Rubin (2015) and identify "meaningful" differences by calculating normalized differences in mean values of covariates across pairs of treatments and designating an absolute value of the normalized difference as meaningful if it exceeds 0.25 in absolute value. For no covariate, in any pairwise comparison of treatments, do we find a value close to 0.25 (see Supplementary Online Material for details on the outcome of normalized differences in means across subject characteristics and attitudes across all treatments). We therefore conclude that the randomization in our experiment was successful, and refrain from including covariates in our statistical analysis of treatment effects. Note too that risk perceptions are not normally distributed, suggesting a non-parametric Wilcoxon Mann-Whitney test is more appropriate than a *t*-test, when analyzing the data presented in panel (a) of Figure 1.

Covid Economics 35, 7 July 2020: 1-50

2.2. An SIR model of the effect of vaccine avoidance on herd immunity

To examine the effects of vaccine avoidance on disease spread, we develop a standard epidemiological (Susceptible Infectious Recovered; SIR) model, following Hethcote (2000), Anderson and May (1991), and discussed in Pindyck (2020). The set-up of the model is as follows. The population is of known size N and consists of three classes of people—susceptible (those who have not yet been infected by SARS-CoV-2 or vaccinated and so are still susceptible to becoming infected), infected (i.e., those who are currently infected), and recovered (those who have either been vaccinated or infected and have since recovered). We assume that the three classes of people are large enough to be taken as continuous functions of time, t. The proportions of the total population in each class are $s(t)$, $i(t)$, and $r(t)$. We assume that those who are recovered from COVID-19 are immune to the virus, but it is important to note that this has not been established. While multiple studies find that people who have recovered from COVID-19 have developed antibodies, it is currently unknown whether that protects them from subsequent infections and if so for how long (WHO, 2020b).

We assume a uniform population that homogeneously mixes.[4] The average number of individual contacts with other people per day by an infectious person is β, so the average daily number of new cases is $\beta N i(t) s(t)$. Infected people recover at a rate of γ per time period. Some people die before they can recover, at a rate of m per time period. The combination of recovery and death results in a death-adjusted duration of infectivity, $\frac{1}{\gamma+m}$, so that the adjusted average number of contacts of an infected person then becomes $\sigma = \frac{\beta}{\gamma+m}$. The average number of

[4] For a recent contribution on how this assumption might affect the required population immunity level to achieve herd immunity, see Britton et al. (2020).

susceptible people who are infected by an infectious person is known as the replacement number and given by $\sigma s(t)$. At the beginning of an outbreak, $s(t) \to 1$ and the replacement number is the basic reproductive rate, R_0. The initial value problem can be written as (suppressing time notation):

1. $\frac{ds}{dt} = -\beta is, \qquad s(0) = s_0$
2. $\frac{di}{dt} = \beta is - (\gamma + m)i, \qquad i(0) = i_0$
3. $\frac{di}{dt} = \gamma i, \qquad r(0) = r_0$

As $r(t) = 1 - s(t) - i(t)$, the well-known solutions (Hethcote, 2000) can be found in the (s, i) plane by dividing equation (2) by (1) to find $\frac{di}{ds} = \frac{\beta is - (\gamma+m)i}{-\beta is} = -1 + \frac{1}{R_0 s}$. By separating variables and integrating we find that $i(t) = i_0 + s_0 - s(t) + \frac{1}{R_0}\ln[\frac{s(t)}{s_0}]$. Each solution is a trajectory from an initial condition of infectious and susceptible people in the population (i_0, s_0) to the terminal condition at which the infection has been extinguished $(0, s_T)$. Herd immunity is achieved at the threshold level of $s^* = \frac{1}{R_0}$.

Covid Economics 35, 7 July 2020: 1-50

3. Results

3.1. Prevalence of COVID-19 vaccine avoidance across risk, mortality and mixed information.

Table 1 shows the outcome of the randomization in the survey experiment into the eight treatment groups. For instance, 376 participants were randomized into the treatment group with information about high conditional mortality (10%) and high probability of becoming infected

(85%) communicated by the CDC only, while 413 participants were randomized into the treatment group with information about low conditional mortality (1.5%) and low probability of becoming infected (25%/10%) communicated by both the CDC and the White House. Note too that the probability communicated by the White house is consistently lower than that communicated by the CDC (see the four treatment groups to the right in Table 1).

Table 1. Treatment groups and their number of participants

Conditional mortality [%]	10	1.5	10	1.5	10	1.5	10	1.5
CDC infection probability [%]	85	85	25	25	85	85	25	25
White House infection probability [%]	--	--	--	--	70	70	10	10
n	376	400	384	391	413	393	363	413

First, we examine if the treatment information on the probability of developing COVID-19 affected participants' beliefs about the risk of COVID-19. Panel (a) of Figure 1 reports participants' beliefs about the risk of the average American becoming infected in the next 12 months across the eight treatment groups. It shows that the high/low probability of infection treatment succeeded in inducing variation in the perceived probability: the perceived risk of becoming infected to the average American is higher in the high probability treatments than in the low probability treatments. These differences in risk perceptions across high and low probability treatments are large in magnitude (around 20 percentage points across all four comparisons of perceived risk in high versus low probability treatments) and highly statistically significant (Wilcoxon Mann-Whitney tests generate $p<0.001$ for all four comparisons). As should be expected, the information on the conditional mortality rate from COVID-19 does not

Covid Economics 35, 7 July 2020: 1-50

affect the perceived probability of infection, given the perceived infection risk only pertains to the probability of catching the disease, not its mortality rate conditional on infection.

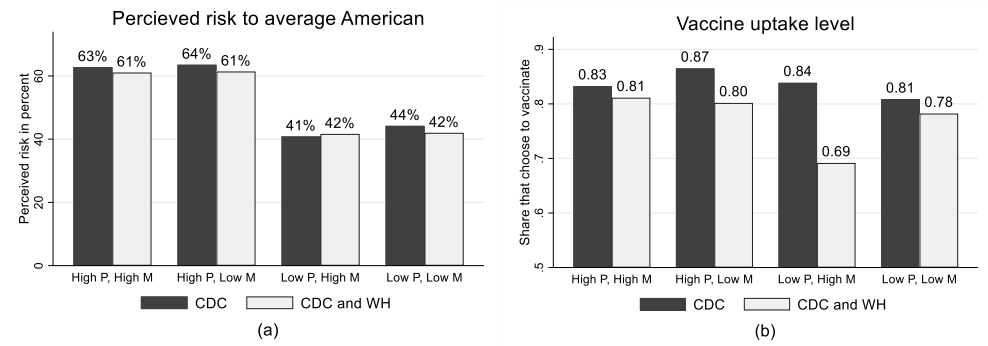

High M=High conditional mortality=10%; Low M=Low conditional mortality=1.5%, High P=High probability of infection=CDC: 85%; White House (WH): 70%, Low P=Low probability of infection= CDC: 25%; WH: 10%. Dark grey bars received information from CDC only about the probability of infection; light grey bars received information from CDC and the White House about the probability of infection.

Figure 1. Mean perceived risk of infection and vaccine uptake across treatments

Panel (a) of Figure 1 also shows that while the treatments succeeded in affecting risk beliefs, participants were not basing their risk beliefs entirely on the information provided to them in the experiment. Specifically, the mean perceived risk to the average American in the four high probability treatment groups in Figure 1 range between 61 and 63.5%, which is substantially lower than the probability of catching COVID-19 communicated by the high probability treatment—recall that participants in the high probability treatment were told either that the probability for the average American of catching COVID-19 in the next 12 months was 85% (if getting information from the CDC only) or 85% according to the CDC and 70% according to the White House. Similarly, the mean perceived risk in the low probability treatment groups range from 40.8 and 44.2%, which is substantially higher than the level of risk communicated by the

Covid Economics 35, 7 July 2020: 1-50

low probability treatments. Participants in these treatments were told either that the probability for the average American of catching COVID-19 in the next 12 months was 25% (if getting information from the CDC only) or 25% according to the CDC and 10% according to the White House.

Finally, panel (a) of Figure 1 shows that participants' mean risk perceptions might have been affected by the lower probabilities communicated by the White House, although differences are small. The mean perceived risk is lower in three out of four treatment comparisons. Specifically, if treated with information entailing high probability of catching COVID-19 and low mortality, risk perceptions are lower if the CDC information is complemented with lower risk information from the White House (Wilcoxon Mann-Whitney test: z=2.367; p=0.018). Similarly, if treated with information about high probability of catching COVID-19 and high severity, risk perceptions are lower if the CDC information is complemented with lower risk information from the White House, although this effect is only weakly statistically significant (Wilcoxon Mann-Whitney test: z=1.656; p=0.098). The exception is when participants are treated with information on low probability and high severity from COVID-19. Here, the mean perceived risk is unexpectedly higher (41.6%, compared to 40.8%) when the CDC risk information is combined with the lower risk information from the White house, although this effect is not statistically significant, as implied by a Wilcoxon Mann-Whitney test (z=-0.154; p=0.877).

Next, we proceed to the focal point of our study and examine COVID-19 vaccine uptake across treatments. When we pool participants in all treatments, we find that 19.5% (n=612/N=3,133) of adults would not vaccinate themselves and 19.7% (n=228/n=1,156) of those with children would not vaccinate their children against COVID-19.

Covid Economics 35, 7 July 2020: 1-50

Panel (b) of Figure 1 shows that the share of people who choose to vaccinate ranges from 69% to 87% across treatments. There is little consistency in the response to the information on conditional mortality from COVID-19; most pairwise differences in shares of people who choose to vaccinate across high and low levels of mortality are too small to detect statistically. The exceptions are the two treatments entailing information on low probability to catch the coronavirus, communicated by both the CDC and White House. Contrary to our prior expectations, for those treatment groups we find that people are more likely to get vaccinated when the conditional mortality is low compared to when it is high (Pearson χ^2 (8.240); p= 0.004).

In contrast, panel (b) of Figure 1 shows that the vaccination decision generally responds to the probabilities of catching the virus in the expected way—the share of people who choose to vaccinate increases as the probability of infection increases. We find that when the conditional mortality rate is high and information on the probability of infection is received from both the CDC and White House, the share of people who vaccinate is 69% if the probability of catching COVID-19 is communicated as low and 81% if the probability of catching COVID-19 is communicated as high (Pearson χ^2 (14.966); $p < 0.001$). When the conditional mortality rate is low and information on the probability of infection is only received from the CDC, the share of people who vaccinate is 81% if the probability of infection is low and 87% if the probability of infection is high (Pearson χ^2 (4.676); $p = 0.031$). While panel (b) of Figure 1 shows that the share of people vaccinating generally increases with higher probability of infection across other pairwise comparisons of low and high probability treatments as well, we cannot statistically detect those differences. Further, the exception is if the conditional mortality rate is high and the probabilities of infection are communicated by the CDC only. Here, the share who choose to

Covid Economics 35, 7 July 2020: 1-50

vaccinate is slightly higher if the probability is low, but this effect is small and far from statistically significant (Pearson χ^2 (0.051); p=0.821).

Panel (b) of Figure 1 also shows that people's vaccination decisions are affected by the White House communicating lower probabilities than the CDC. The share of people who vaccinate is consistently lower in the treatments where the White House communicates lower probabilities of catching COVID-19, compared to the CDC. The effect on vaccine uptake from adding lower probability information from the White House is particularly large in the treatments with low probability of infection and high conditional mortality rate. Here, the share that vaccinates drops from 84 to 69% (Pearson χ^2 (22.593); p<0.001), if the White House communicates lower probability of infection than does the CDC. Similarly, there is a decline in the share that vaccinates in the treatments with high infection probability and low mortality rates, if lower risk information from the White House is included. Then, the share that vaccinates drops from 81 to 78% (Pearson χ^2 (5.756); p=0.016). Other differences in shares who vaccinate for other treatments are too small for us to detect statistically. Overall, however, these results suggest that inconsistent information about risks not only may affect risk perceptions (as shown in panel (a) of Figure 1), it affects behavioral intentions.

Previous studies show that people find health risk information more believable if it is received from sources that share their values (Siegrist et al., 2000; Siegrist et al., 2001). For example, in a study of cancer cluster communication, Siegrist and colleagues (2001) found that people were more likely to believe clusters could occur randomly when they believed risk managers shared their values. Based on these findings, we hypothesized that individuals who identify as Republican or conservatives would respond the most to the White House information. However, we cannot detect any differences in responses to the White House information across

Covid Economics 35, 7 July 2020: 1-50

Republicans and Democrats; both groups appear to be just as responsive to the more optimistic White House information.

3.2. Determinants of vaccine avoidance

Next, we pool participants from all treatments to examine a broad range of determinants of the vaccine decision. Table 2 shows the average marginal effects generated by a probit regression on the decision to vaccinate for COVID-19 from a wider set of variables.

Covid Economics 35, 7 July 2020: 1-50

Table 2. Determinants of decision to vaccinate for COVID-19 – average marginal effects from a probit regression.

Variable	$\partial \Pr(vacc) / \partial x$	Standard errors
Treatment – Low Severity	0.023*	(0.012)
Treatment – Low Probability	-0.050***	(0.012)
Treatment – White House	-0.053***	(0.012)
Female	-0.058***	(0.014)
Low Income	-0.036*	(0.015)
High Income	-0.026	(0.018)
Rural area	0.009	(0.013)
Prevention Measures	0.016***	(0.003)
Flu Shot	0.133***	(0.012)
Vaccination Schedule	0.025*	(0.015)
Trust in Government Agencies – High	0.032*	(0.016)
Trust in Government Agencies – Low	-0.051**	(0.015)
Vaccine Confidence	0.168***	(0.015)
Vaccine Complacency	-0.020	(0.020)
Vaccine Calculation	-0.067***	(0.014)
Vaccine Collective Responsibility	-0.080***	(0.018)
Vaccine Constraint	-0.012	(0.020)
Democrat	0.000	(0.015)
Other Political Party	-0.038**	(0.016)
Belief in God	-0.045**	(0.017)

Perceived Risk of Catching COVID (High)	0.006	(0.017)
Perceived Risk of Catching COVID (Low)	-0.061***	(0.015)
Perceived Risk of Severe Consequences (High)	0.014	(0.016)
Perceived Risk of Severe Consequences (Low)	-0.042***	(0.015)
N	3,133	

*** p<0.01, ** p<0.05, * p<0.1.

Above we estimated the treatment effects with no controls for potential confounding factors. Here we include variables measuring the treatments in the regression analysis underlying the results in Table 2, to reduce the risk of omitted variable bias. However, it should be noted that the results in Table 2 are not sensitive to their inclusion, and the estimated effects of these variables in the Probit model are consistent with the effect found in Figure 1. Specifically, the variable *Treatment – Low Severity* takes the value 1 if the treatment communicated a low conditional mortality rate, and 0 if it communicated high mortality. The results in Table 2 therefore reinforces the finding in Figure 1 of the unclear role of COVID-19 severity to the vaccination decision. The average marginal effect from low severity is positive, suggesting rather counterintuitively that a higher risk of suffering severe consequences from COVID-19 might deter people from vaccinating. This result, however, is only weakly statistically significant. The variable *Treatment – Low Probability* takes a value 1 if the treatment communicated low probability of infection, and 0 if it communicated high probability. The marginal effect shows that if the probability of catching COVID-19 is low, people are 5% less likely to vaccinate. The variable *Treatment – White House* takes the value 1 if the CDC and the White House were jointly communicating the probability of catching COVID-19, with the White House communicating a lower probability than the CDC; 0 if the CDC alone communicated the probability. Table 2 shows that the vaccination decision is affected by the White House communication—people are more than 5% less likely to vaccinate if the probability of catching

Covid Economics 35, 7 July 2020: 1-50

the virus is communicated as 15 percentage points lower by the White House than the risk communicated by the CDC.

Table 2 shows that women are 6% less likely to vaccinate than men, and low-income earners are close to 4% less likely to vaccinate than medium income earners. We cannot statistically detect an effect on vaccine intentions from being a high-income earner, compared to a medium-income earner. Further, we cannot identify a difference in vaccine intentions between rural and urban households. The variable *Prevention Measures* represents the number (ranging from 0 to 12) of preventive measures a participant takes to avoid getting infected by COVID-19. These measures include washing their hands more, becoming better informed, praying to stay resilient, and eating better. We find that people who are taking more of preventive measures are more likely to vaccinate for COVID-19. Hence, while it would be entirely plausible that people who reduce their risk of infection in other ways would choose to vaccinate less, we find the opposite – those who self-protect in other ways are also more likely to vaccinate.

Table 2 shows the importance of general vaccination behavior and attitudes for the vaccination decision. The variable *Flu shot* is a dummy variable that takes the value 1 if a person took the flu shot in the last two years; 0 otherwise. The marginal effect for *Flu shot* shows that people who took the flu shot are more than 13% more likely to also take the COVID-19 vaccine. The marginal effects for *Vaccination schedule* shows that a person who has followed the recommended vaccine schedule is almost 3% more likely to take the COVID-19 vaccine, compared to someone who did not, although this effect is only weakly statistically significant.

We find that trust in government agencies matters to the decision to vaccinate for COVID-19. We split the sample in three equal size groups for low (64% or lower trust), medium (65 to 88% trust) and high trust (89% or higher) of government agencies. The marginal effect for *Trust*

Covid Economics 35, 7 July 2020: 1-50

in Government Agencies – High is a dummy variable that takes the value 1 if a person states high trust in government. Table 2 shows that people with a high trust in government agencies are 3% more likely than those of medium trust (the reference case) to get vaccinated, as implied by the positive marginal effect for *Trust in Government Agencies – High*, although this effect is only weakly statistically significant. People with a low trust in government agencies are 5% less likely to get vaccinated than those with a medium trust in government agencies. This is consistent with previous findings on the positive correlation between government trust and vaccine uptake (Lee et al., 2016).

The dummy variables *Vaccine confidence*, *Vaccine complacency*, *Vaccine calculation*, *Vaccine collective responsibility*, and *Vaccine constraint* represent the five key components of Betsch et al. (2018) vaccine hesitancy scale. *Vaccine confidence* takes a value 1 if a participant agrees with the statement "I am completely confident that vaccines are safe"; *Vaccine complacency* takes a value 1 if a participant agrees with the statement "Vaccination is unnecessary because vaccine-preventable diseases are not common anymore"; *Vaccine calculation* takes a value 1 if a participant agrees with the statement "When I think about getting vaccinated, I weigh benefits and risks to make the best decision possible"; *Vaccine collective responsibility* takes a value 1 if a participant agrees with the statement "When everyone is vaccinated, I don't have to get vaccinated too." *Vaccine constraint* takes a value 1 if a participant agrees with the statement "Everyday stress prevents me from getting vaccinated."

The average marginal effect for *Vaccine confidence* shows that people who are confident that vaccines in general are safe are 17% more likely to take the COVID-19 vaccine. We do not find an effect on the decision to vaccinate from people believing vaccines are unnecessary because vaccine preventable diseases are uncommon, as shown by the marginal effect for *Vaccine*

Covid Economics 35, 7 July 2020: 1-50

complacency. People who weigh benefits against costs for vaccines are almost 7% less likely to get vaccinated, as implied by the average marginal effect for *Vaccine calculation.* People who agree that they do not need to get vaccinated if everyone else is vaccinated are 8% less likely to get a COVID-19 vaccine, as implied by the average marginal effect for *Vaccine collective responsibility*. We do not find an effect on COVID-19 uptake from people agreeing that everyday stress prevents them from getting vaccinated, as implied by the estimated marginal effect for *Vaccine constraint*.

We do not find a difference in the probability of vaccinating for COVID-19 across Democrats and Republicans (the reference case). People identifying with neither political party are less likely than Republicans to vaccinate. The effects of the political party dummy variables remain robust if we exclude *Belief in God*. People who state they believe in God are almost 5% less likely to vaccinate, compared to non-believers.

Finally, individual risk matters to the vaccine decision. The dummy variable *Perceived Risk of Catching COVID (High)* takes a value of 1 if participants believe their risk of catching COVID-19 is higher than that of the average American, and the dummy variable *Perceived Risk of Catching COVID (Low)* takes a value of 1 if participants believe their risk of catching COVID-19 is lower than that of the average American. Compared to participants who believe their risk is the same as that of the average American (the reference case), Table 2 shows that participants who believe their risk of infection is lower than that of the average American are 6% less likely to get vaccinated, while we cannot detect a difference in vaccine intentions between participants who believe their risk is higher than that of the average American, compared to those who think their risk is the same as that of the average American. Similarly, the dummy variable *Perceived Risk of Severe Consequences (High)* takes a value of 1 if participants believe they face

Covid Economics 35, 7 July 2020: 1-50

a higher conditional mortality rate from COVID-19 than the average American. While we cannot detect a difference in vaccine intentions between those who think they are at higher risk of severed consequences than the average American, compared to those who think they face the same risk of severe consequences as the average American (the reference case), we do find that those who believe they are at lower risk of severe consequences are 4% less likely to get vaccinated, compared to those who believe they face the same risk of severe consequences as the average American. Hence, while we do not find a consistent effect of the conditional mortality rate of the *average* American on vaccine intentions, our regression results suggest beliefs about the conditional mortality rate might still matter –people who believe their conditional mortality rate is lower than that of the average American (28% of our sample) are less likely to vaccinate. Taken together, this could imply that the average conditional mortality rate is relatively unimportant to vaccine decisions, but beliefs about individual mortality risk matter.[5]

The effects reported in Table 2 are robust to a range of inclusions of other (non-significant) explanatory variables, such as college education, race, underlying health conditions that increases the risk of severe consequences if infected, or working in a high risk profession (health care, teaching). They are also robust to the inclusion of either the compressed treatment variables (the top three variables in Table 2), all eight treatment dummy variables, or no variables representing the treatment effects. Further, despite many variables measuring different aspects of vaccine hesitancy, the multicollinearity amongst the variables in the model is low, as implied by a variance inflation factor (VIF) of 1.27.

Next, we examine the main reasons people state they choose not to vaccinate against

[5] Many people might not relate to the mortality risk of the average American -- a majority (58%) of our participants believe their conditional mortality rate is higher or lower than that of the average American.

Covid Economics 35, 7 July 2020: 1-50

Covid Economics 35, 7 July 2020: 1-50

COVID-19. All participants who chose not to vaccinate themselves or their children where asked: "You indicated that one or more of your family members would not get vaccinated. Please mark the extent to which any of the below reasons mattered to your decision not to take the vaccine." The reasons are shown in Figure 2, and participants could mark if these reasons did not matter at all, mattered some, or mattered a lot. Figure 2 shows the share of participants who stated the reason mattered some or a lot, of participants who declined the vaccine.

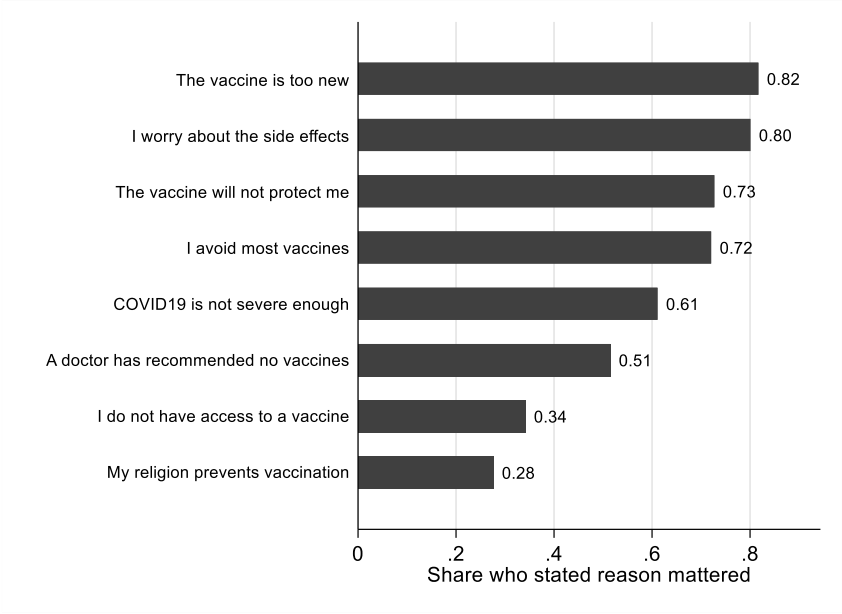

Figure 2. Reasons for declining the COVID-19 vaccine

Figure 2 shows that the most important reason to decline the COVID-19 vaccine it its novelty and worry about negative side-effects, with 82% of those who declined the vaccine agreeing with the statement "I worry the vaccine is so new we do not understand the side effects" and 80% stating that "I worry about the negative side effects of the vaccine" mattered to their decision to decline the vaccine. This is consistent with previous findings that people are particularly

skeptical to new vaccines (Dube et al., 2013). Figure 2 also shows that general hesitancy towards all vaccines spills over on the COVID-19 vaccine. Of those declining the vaccine, 72% state that general avoidance of vaccines is an important reason for also avoiding the COVID-19. Other important reasons for declining the vaccine are doubts that the vaccine will in fact provide protection from catching the virus and the belief that COVID-19 is not severe enough to warrant vaccination.

3.3. The effect of COVID-19 vaccine hesitancy on disease spread

We next use the SIR model to examine the implications of vaccine avoidance for the potential effectiveness of a mass COVID-19 vaccine program, implemented at the beginning of an upcoming COVID-19 season.

Panel (a) of Figure 3 shows the general solution to the SIR model in (s, i) space by the red line for an initial condition (i_0, s_0). We calibrate the initial conditions to roughly represent current (end April 2020) conditions in the U.S. – the population *N* is set to 327 million, i_0 represents 1 million infected (1/327, or 0.3%; note that this is roughly the number of confirmed cases, which may differ from actual cases, due to limited testing), r_0 represents 100,000 recovered (0.1/327, or 0.03%) and s_0 (99.67%) represents the rest of the population (the number of deaths is small enough to not affect the results perceptibly). In panel (a) of Figure 3, we assume a moderate level of R_0 =2.4 for the novel coronavirus, in line with estimates by Liu et al. (2020) and Fergusson et al. (2020), and adopted in Thunström et al. (2020b). Panel (b) shows the trajectory with a higher R_0 = 5.7, as estimated by Sanche et al. (2020). In the absence of a vaccine or other means of controlling the virus, the infection would spread from the initial condition in the lower right corner, following the upper red line to a peak in the proportion

Covid Economics 35, 7 July 2020: 1-50

infected when herd immunity is reached at s^*, after which the infection burns out, leaving $s_v^T = 0$ susceptible.

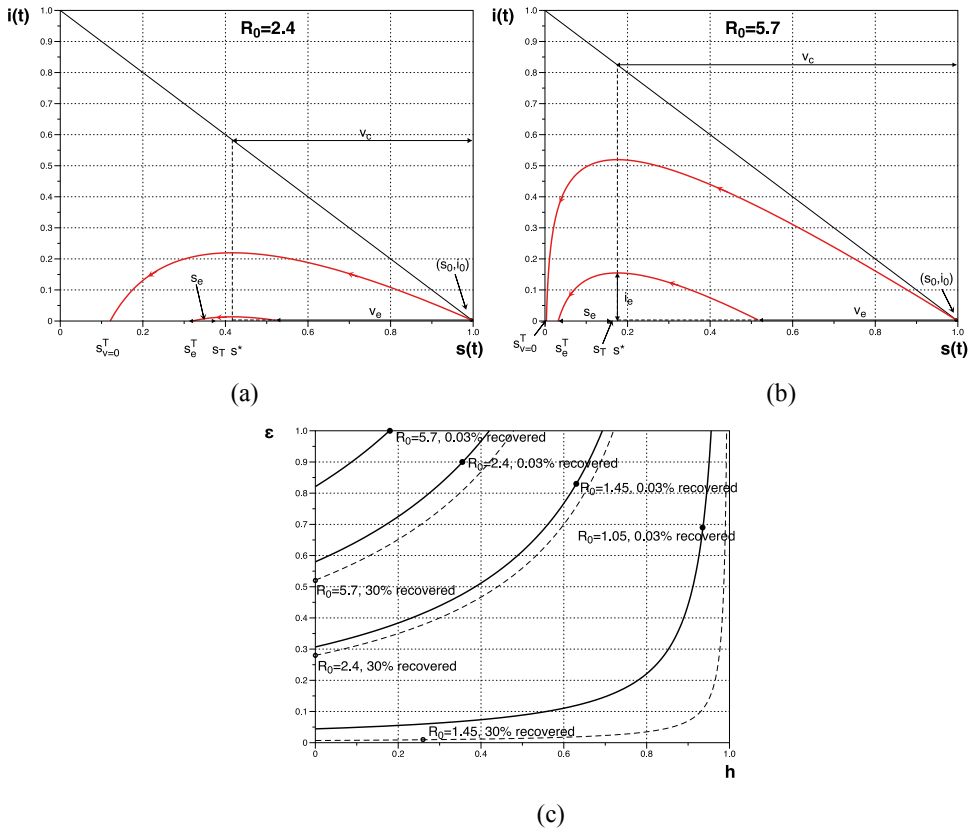

Figure 3. Panel (a): epidemic solution at $R_0 = 2.4$ and N =327 million; $i_0 = 0.3\%$; $r_0 = 0.03\%$; $s_0 = 99.67\%$; Panel (b): epidemic solution at $R_0 = 5.7$, all else same as (a). Upper red lines are the unvaccinated paths of the epidemics, lower red line the vaccinated path. Panel (c): threshold parameter combinations that allow $\nu_e = \nu_c$. Solid lines reflect initial conditions employed in panel (a), dashed lines assume 30% recovered and immune.

The lower red lines of panels (a) and (b) in Figure 3 show the epidemic solution with a vaccine. A vaccine program reduces the peak of infections and has the potential to accelerate the burnout of the epidemic, leaving more people in the susceptible (never infected) class. A mass

Covid Economics 35, 7 July 2020: 1-50

vaccine program can therefore be illustrated as an immediate reduction in the number of susceptibles at the very beginning of the season, i.e., a reduction in s_0. The critical proportion of vaccinated that ensures herd immunity and burns out the epidemic (Anderson and May, 1991; Keeling and Rhohani, 2019) is the reduction in the proportion of susceptibles from 1 to s^*, or $v_c = 1 - \frac{1}{R_0}$, as shown in panels (a) and (b). The vaccine is deployed across $v_c N$ individuals, and if mass vaccination is implemented after the initiation of the epidemic, then the critical proportion of the population that needs to be vaccinated is $v_c = s_0 - s^*$, ending with s_T susceptibles. If the vaccine is 100% effective (i.e., it entirely eliminates the risk of becoming infected for those who are vaccinated), and the entire population is willing to take the vaccine, then for an $R_0 = 2.4$, v_c requires roughly 58% of the susceptible population to be vaccinated to extinguish the epidemic and leaves 120.4 million people who never got infected (or $s_T = 0.368$). For an $R_0 = 5.7$, then v_c requires roughly 82% of the susceptible population to be vaccinated to extinguish the epidemic and leaves 47 million people who never got infected (or $s_T = 0.145$).

However, the outcome of a vaccine program in panels (a) and (b) of Figure 3 relies on the highly optimistic assumptions that no-one declines the vaccine and the vaccine is 100% effective. Vaccines are typically unable to completely eliminate the risk of becoming infected, in part because many viruses can rapidly mutate. For instance, in a particularly successful season when flu vaccines are well-matched to the circulating flu viruses, the flu vaccine may reduce the risk of becoming infected with the flu virus by 40-60% (CDC, 2020c). Like the flu and measles viruses, the coronavirus is an RNA virus, known to mutate relatively quickly. However, for being part of that group, it so far has mutated slowly—data collected so far suggests the seasonal flu mutates four times faster than the novel corona virus (Moshiri, 2020). To account for vaccine avoidance and limited vaccine effectiveness, let h be the share of the population avoiding the

Covid Economics 35, 7 July 2020: 1-50

vaccine and let the vaccine have effectiveness ε, where a perfectly effective vaccine implies $\varepsilon = 1$. Further, let v be the proportion of the population that is given the opportunity to vaccinate. The resulting *effective* proportion vaccinated, v_e, is then $v_e = \varepsilon(1-h)v$.

Panel (c) of Figure 3 illustrates how the success of a vaccine program in preventing another outbreak depends on the value of key determinants of s_0 , i.e., ε, h, R_0 and the proportion of recovered and immune, r_0. We show ranges of possible values for these parameters, given their precise values are still uncertain. Contours for R_0 and r_0 trace out the combinations of ε and h when the effective proportion vaccinated just matches the critical proportion, $v_e = v_c$, documenting the (in)ability of the vaccine program to extinguish COVID-19. For combinations of ε and h below and to the right of each contour, the proportion effectively vaccinated is less than the critical proportion, $v_e < v_c$, and the vaccine program fails to prevent a COVID-19 outbreak of infections. In contrast, parameter combinations above and to the left of each contour result in the epidemic being extinguished with no new wave of infections.

Let h = 0.20 be our baseline value of vaccine avoidance, as generated by our experimental study, and let ε = 0.60 be our baseline vaccine efficiency, since participants were told in the experimental survey that the COVID-19 vaccine would be 60% effective, based on this being the upper bound of effectiveness of the flu vaccine (another, but more rapidly mutating, RNA virus). Finally, let r_0 = 0.03% be our baseline share of recovered and immune people.

At these baseline values, Figure 3 panel (c) shows that if R_0 = 2.4, then we are at a point below and to the right of the relevant contour, or $v_e < v_c$. This means the vaccine program fails to achieve herd immunity, and a next wave of infections will occur. Specifically, in this case, at most 52% of the population could be effectively vaccinated, which is too low a share of the

Covid Economics 35, 7 July 2020: 1-50

Covid Economics 35, 7 July 2020: 1-50

population to prevent an outbreak. Note that this result occurs even though Figure 3 assumes rather optimistically that the vaccine program can target 100% of the population, i.e., $v = 1$, such that all Americans are given the opportunity to take a COVID-19 vaccine. Panel (c) further shows that if R_0= 5.7, while h and r_0 remain at their baseline values, there is no level of vaccine effectiveness that could extinguish the epidemic. With lower R_0 chances of preventing a new outbreak increase. If R_0= 2.4 and h and r_0 are maintained at their baseline values, a higher vaccine effectiveness at ε = 0.75 would be needed to prevent a new outbreak. At R_0= 1.45, then $v_e > v_c$, and the vaccine program could prevent another outbreak. Further, for very low R_0, vaccine avoidance may have no impact at all on the ability of a vaccination program to extinguish the epidemic. For example, if ε = 0.10 and R_0= 1.05 (lower than any known estimates of R_0 with 0.03% recovered and immune, the epidemic burns out at any level of vaccine avoidance below 55%. The higher is R_0 the steeper the slope of the contour at any given level of h, reflecting a larger impact of vaccine avoidance.

Finally, our baseline assumption of the proportion of the population recovered when the vaccine program is initiated, i.e., r_0 = 0.03, is based on current *confirmed* recovered cases. Given limited and unreliable testing, the number of actual recovered cases could be substantially higher. Panel (c) of Figure 3 shows the important role of the recovered proportion of the population—if we assume a substantially higher proportion, e.g., 30% of the population such that r_0 = 30%, the contours are pushed down and to the right (i.e., the solid lines are pushed to the dashed). For R_0 = 1.05, the higher proportion recovered means the vaccination program extinguishes the epidemic for all combinations of vaccine avoidance and vaccine effectiveness—hence, this contour does not appear in the graph. Now consider the contours for R_0 = 5.7 and R_0 = 2.4. If r_0 = 0.03, $h = 0.20$, and $\varepsilon = 0.60$, the vaccine program will experience an additional

wave of infections as $v_e < v_c$ at both $R_0 = 5.7$ and $R_0 = 2.4$. However, if the initial proportion recovered were to be 30%, then the epidemic would be extinguished at $h = 0.20$, and $\varepsilon = 0.60$ if $R_0 = 2.4$, while not if $R_0 = 5.7$.

Next, we return to Figure 3 panels (a) and (b) to explore the implications from the inability of the vaccine program to extinguish COVID-19, see the lower red lines in both panels. Employing the baseline parameter values $r_0 = 0.03$, $\varepsilon = 0.60$, $h = 0.20$ and $R_0 = 2.4$ in panel (a), the proportion of peak infections rises by $i_e = 0.010$ above the level that would have been if the vaccine program had extinguished the epidemic. This is equivalent to almost 3.4 million excess peak infections (not shown on the figure to avoid clutter). Further, the proportion of susceptibles is $s_e = 0.048$ lower at the end of the epidemic, which is equivalent to almost 15.87 million fewer remaining susceptible people than if perfect vaccination were possible. If we consider $R_0 = 5.7$ as in panel (b) the resultant $i_e = 0.152$ and $s_e = 0.113$ provide 49.6 million excess peak infections and 36.8 million fewer susceptibles at the end of the epidemic.

Importantly, panels (a) and (b) of Figure 3 are based on the assumption of few initial recoveries, i.e., $r_0 = 0.03\%$. If $r_0 = 30\%$, then the proportion effectively vaccinated approaches the critical proportion, $v_e \rightarrow v_c$, as shown by contrasting $r_0 = 0.03\%$ and $r_0 = 0.30\%$ in panel (c). For the level of vaccine avoidance, vaccine effectiveness and R_0 underlying panel (a) in Figure 3, more than 10% of the current population would need to be recovered for $v_e = v_c$, while the same proportion needs to be 34% for $v_e = v_c$ at the higher R_0 in panel (b).

While the above illustrates how combinations of levels of vaccine hesitancy, vaccine ineffectiveness, the share of the population that is immune and R_0 pose some real challenges to the success of a COVID-19 vaccine program, it also illustrates how vaccine avoidance reduces

Covid Economics 35, 7 July 2020: 1-50

the effectiveness of a vaccine program and how, under certain conditions, it may even be the deciding factor for whether a program succeeds in extinguishing the epidemic.

4. Discussion

A vaccine for COVID-19 might be the best hope for ending the pandemic. Scientists are therefore racing to develop a vaccine, in unprecedented joint efforts within the scientific community. However, the challenge to extinguish the novel coronavirus does not end with finding an effective vaccine. The implementation of the vaccine program will be important. In this study, we focus on the challenge posed to a vaccine program of vaccine hesitancy. What if large parts of the population decline the vaccine, once it is available? Vaccine hesitancy is well-known for other types of vaccines, and has increased in recent years (Dube et al., 2013; Olive et al., 2018), for instance causing recent outbreaks of measles (De Serres et al., 2013; Sarkar et al., 2019), a disease that was extinguished in the U.S. until recently.

In a survey experiment (N=3,133) that accounts for uncertainty in probabilities of infection and conditional mortality rates, we find that around 20% of Americans would decline a COVID-19 vaccine, and this proportion is consistent across adults and their children. This result is similar to finding from recent polls. A survey by Pew Research Center (2020) finds that 27% of U.S. adults would not get a coronavirus vaccine if it was available today, a poll by ABC news/Ipsos (2020) suggests that 25% of U.S. adults were unlikely to get vaccinated if an effective coronavirus vaccine was developed, and a poll by LX/Morning Consult (2020) finds that 9% of U.S. adults would not get vaccinated if a vaccine became available while another 15% do not know if they would get vaccinated.

Covid Economics 35, 7 July 2020: 1-50

Covid Economics 35, 7 July 2020: 1-50

Our exploration of the determinants of vaccine avoidance suggests that the probability of infection matters—the higher the probability of the average American to catch the virus, the more likely people are to choose to vaccinate. Further, having had a flu shot in the last two years is one of the most important determinants of vaccine avoidance observed in our study. Given family physicians have records (although perhaps incomplete) of who had flu shots, this presents an opportunity to identify people at particularly high risk of being skeptical towards a future COVID-19 vaccine. Related, we find that general vaccine hesitancy also increases avoidance of a COVID-19 vaccine, and the novelty of the vaccine and concerns about its safety are likely to decrease the uptake in the population even further.

The outcome of a standard epidemiological model suggests COVID-19 vaccine avoidance has important implications for public health. If we assume a vaccine will be available at the beginning of the next COVID-19 season (or a season thereafter), and employ the currently best available information about key parameters of the model (vaccine avoidance, vaccine effectiveness, COVID-19 infectiousness, amount of recovered and immune individuals), we find that a COVID-19 vaccine program can reduce the amount of infections but will likely fail to generate herd immunity on its own. We extend the analysis to examine combinations of key parameter values, including vaccine avoidance, that could ensure the vaccine program's ability to prevent a new COVID-19 outbreak. All else equal, vaccine avoidance poses a greater challenge to the vaccine program if the infectiousness of COVID-19 is high, the effectiveness of the vaccine in preventing infections is low, or the acquired immunity level in the population when entering the next COVID-19 season is low.

Knowing about COVID-19 vaccine avoidance before a vaccine is available can help government agencies, health care workers, and other authorities mitigate the impact of vaccine

Covid Economics 35, 7 July 2020: 1-50

avoidance. Such efforts may involve developing policies and a preparedness for the vaccine avoidance. It might also involve public information campaigns designed to increase confidence in the effectiveness and safety of the vaccine. Here, our results offer insights that may be helpful. First, our results underscore the importance of a uniform risk message from government authorities. We find that inconsistent information from government authorities about COVID-19 risks may affect not only risk perceptions, but also health related behaviors—vaccine avoidance increases if the White House communicates lower risks to COVID-19 than does the CDC. This result relates to findings that risk information in the news had direct effects on people's health behavior during the pandemic (Bursztyn et al., 2020). Second, we find that distrust in the government is higher amongst those who decline the vaccine. To address COVID-19 vaccine hesitancy, broader public health campaigns may therefore be less effective. Instead, efforts might focus on reaching out to health care providers (the most trusted source of vaccine safety information, see e.g., Freed et al., 2011) and local authorities, including religious leaders. To identify effective strategies to reduce COVID-19 vaccine hesitancy, government agencies will likely benefit from the knowledge gathered during recent measles outbreaks—the COVID-19 vaccine decliners in our study share many attitudes with those generally avoiding vaccines (MacDonald, 2015). Policy makers may also consider regulations that require people to have COVID-19 vaccinations in order to attend schools and workplaces, similar in spirit to the bans of philosophical exemptions from vaccinations in the wake of the recent measles outbreaks (Kuehn, 2019).

Our study has several short comings. First, we only consider one level of vaccine effectiveness (60%) and it is possible that the level of vaccine avoidance observed in our study is unique to that level. Second, our SIR model relies on the common assumptions that the

Covid Economics 35, 7 July 2020: 1-50

population homogenously mixes, and all individuals are equally susceptible and infectious, if infected. Britton et al. (2020) show that instead employing the more realistic assumption of heterogeneous mixing in the population likely reduces the level of population immunity required to achieve herd immunity, although extensive, and country specific, explorations of mixing patterns would be required to narrow down the extent to which this might happen. The underlying immunity level in the population required to achieve herd immunity in our model might therefore be too high. Third, the uncertainty about both probabilities of infection and conditional mortality rates is yet to be resolved, through more information generated by increased testing, such that the upper and lower bounds of those variables employed in our survey might turn out to be too low/high. Some of the uncertainty has resolved since the time of data collection. The current case fatality rate (generally an overstatement of the conditional mortality rate) in the U.S. is 5.1% (Johns Hopkins University and Medicine (2020); Roser et al. (2020), which is lower than our upper bound of the conditional mortality rate (10%) and higher than our lower bound (1.5%). This implies that the true conditional mortality rate is likely closer to 1.5% than to 10%, which might be known to the general population when a vaccine becomes available. That said, our results suggest the conditional mortality rate for the average American, within the range explored in our study, might not matter much to vaccine decisions.[6] Future research might examine more in detail the role of beliefs about conditional mortality rate to vaccine decisions.

We also encourage future research to examine the relationship between vaccine effectiveness and vaccine avoidance, as well as the relationship between vaccine costs and avoidance. In our

[6] Also, if we calculate average vaccine avoidance for the treatments with low conditional mortality rate only (n=1,597), we find that the prevalence of vaccine avoidance amounts to 81%, i.e., about the same as the average vaccine avoidance for the study as a whole.

Covid Economics 35, 7 July 2020: 1-50

study, participants were asked to consider a costless COVID-19 vaccine. Getting a vaccine will be affected by some cost, whether it be a time, financial or inconvenience cost, such that the vaccine avoidance in our study is likely on the lower end. A study that measures the cost people are willing to bear for a COVID-19 vaccine may be helped by the literature examining willingness to pay for health risk reductions, see e.g., Sloan et al. (1987), Smith and Desvousges (1987), Viscusi and Evans (1990), Viscusi and Aldy (2003), Edwards (2008), Hammitt and Haninger (2010), Alberini and Ščasný (2013), Finkelstein et al. (2013), Gerking et al. (2017). Related, Serra-Garcia and Szech (2020) find that costs to COVID-19 antibody tests have substantial negative effects on people's willingness to get tested.

In the U.S., conflicting risk messages about COVID-19 are regularly communicated to the public, by the media and public authorities, and have been shown to affect other types of health behavior related to COVID-19 (see e.g., Bursztyn et al., 2020; Simonov et al. 2020). We encourage future research to further explore the mechanisms by which conflicting risk information affect vaccination decisions. For instance, Viscusi (1997) finds that disparity in the risks communicated by different sources might affect how people process information—the greater the disparity, the lower the trust in all information sources, Fox et al. (2002) show that negative information dominates positive information when consumers form private values about risky food, while Viscusi (1997) and Viscusi et al. (1999) show that people may place different weights on information sources, depending on whether the sources communicate high or low risk, even independent of the sources' trustworthiness. Results from these studies could be applied to COVID-19 risks, and generate important novel insights into risk communication in response to this and future pandemics. Finally, other studies show that some preventive behaviors during the current pandemic are motivated by prosocial attitudes (Campos-Mercade et

al., 2020; Jordan et al., 2020; Thunström et al., 2020a). Future studies might similarly examine if prosocial messaging might increase a COVID-19 vaccine uptake.

Covid Economics 35, 7 July 2020: 1-50

Covid Economics 35, 7 July 2020: 1-50

References

ABC news/Ipsos poll (2020), Topline and Methodology, https://www.scribd.com/document/460461528/ABC-News-Ipsos-Poll-May-7-2020, Retrieved June 22 2020.

Alberini, A., & Ščasný, M. (2013). Exploring heterogeneity in the value of a statistical life: Cause of death v. risk perceptions. *Ecological Economics*, *94*, 143-155.

Anderson, R.M. and May, R.M. (1991) *Infectious diseases of humans*. New York: Oxford Science Publication.

ASTHO. Communicating Effectively About Vaccines: Summary of a Survey of U.S. Parents and Guardians. Arlington, VA: Association of State and Territorial Health Officials, 2010.

Athey, S. and Imbens, G.W. (2017). The econometrics of randomized experiments. In Banerjee, A.V. & Duflo, E. (eds.). *Handbook of Economic Field Experiments,* 73-140. Amsterdam, Netherlands: North-Holland.

Betsch, C., Schmid, P., Heinemeier, D., Korn, L., Holtmann, C., and Böhm, R. (2018). Beyond confidence: Development of a measure assessing the 5C psychological antecedents of vaccination. *PLoS One*, *13*(12).

Binder, A. R., Hillback, E. D., and Brossard, D. (2016). Conflict or caveats? Effects of media portrayals of scientific uncertainty on audience perceptions of new technologies. *Risk Analysis*, *36*(4), 831-846.

Breakwell, G. M. (2000). Risk communication: factors affecting impact. *British Medical Bulletin*, *56*(1), 110-120.

Brewer, N. T., Chapman, G. B., Gibbons, F. X., Gerrard, M., McCaul, K. D., and Weinstein, N. D. (2007). Meta-analysis of the relationship between risk perception and health behavior: the example of vaccination. *Health psychology*, *26*(2), 136.

Britton, T., Ball, F., and Trapman, P. (2020). A mathematical model reveals the influence of population heterogeneity on herd immunity to SARS-CoV-2. *Science*, doi: 10.1126/science.abc6810

Bursztyn, L. Rao, A., Roth, C. and Yanagizawa-Drott, D. (2020) Misinformation during a pandemic, Becker Friedman Institute Working paper 2020-44, https://bfi.uchicago.edu/wp-content/uploads/BFI_WP_202044.pdf. Retrieved May 1.

Callaway, E. (2020) The race for coronavirus vaccines: a graphical guide -- Eight ways in which scientists hope to provide immunity to SARS-CoV-2. *Nature*, 580, 576-577. doi: 10.1038/d41586-020-01221-y

Calman, K., and Curtis, S. (2010). *Risk communication and public health*. Oxford University Press.

Campos-Mercade, P., Meier, A., Schneider, F., & Wengström, E. (2020). Prosociality predicts health behaviors during the COVID-19 pandemic. *University of Zurich, Department of Economics, Working Paper*, (346).

Carpenter, D. M., Elstad, E. A., Blalock, S. J., and DeVellis, R. F. (2014). Conflicting medication information: prevalence, sources, and relationship to medication adherence. *Journal of Health Communication*, *19*(1), 67-81.

Covid Economics 35, 7 July 2020: 1-50

Cascella, M., Rajnik, M., Cuomo, A., Dulebohn, S. C., and Di Napoli, R. (2020). Features, evaluation and treatment coronavirus (COVID-19). In *Statpearls [internet]*. StatPearls Publishing.

Centers for Disease Control and Prevention (CDCa) https://www.cdc.gov/flu/about/season/flu-season-2017-2018.htm, Retrieved April 20, 2020.

Centers for Disease Control and Prevention (CDCb), https://www.cdc.gov/coronavirus/2019-ncov/downloads/COVID19-What-You-Can-Do-High-Risk.pdf, Retrieved April 4, 2020.

Centers for Disease Control and Prevention (CDCc), https://www.cdc.gov/flu/vaccines-work/vaccineeffect.htm, Retrieved April 26, 2020.

Chandler, J., and Paolacci, G. (2017). Lie for a dime: When most prescreening responses are honest but most study participants are impostors. *Social Psychological and Personality Science*, 8(5), 500-508.

CNN, https://www.cnn.com/2020/04/07/politics/white-house-press-secretary-coronavirus/index.html. Retrieved April 20 2020.

De Serres, G., Markowski, F., Toth, E., Landry, M., Auger, D., Mercier, M., Bélanger P., Turmel B., Arruda, H., Boulianne, N., and Ward B.J. (2013). Largest measles epidemic in North America in a decade—Quebec, Canada, 2011: contribution of susceptibility, serendipity, and superspreading events. *The Journal of Infectious Diseases*, *207*(6), 990-998.

Dubé, E., Laberge, C., Guay, M., Bramadat, P., Roy, R., and Bettinger, J. A. (2013). Vaccine hesitancy: an overview. *Human Vaccines & Immunotherapeutics*, *9*(8), 1763-1773.

Covid Economics 35, 7 July 2020: 1-50

Edwards, R. D. (2008). Health risk and portfolio choice. *Journal of Business & Economic Statistics*, *26*(4), 472-485.

Everett, J. A. (2013). The 12 item social and economic conservatism scale (SECS). *PLOS ONE*, *8*(12).

Ferguson, N. M., Laydon, D., Nedjati-Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunubá, Z., Cuomo-Dannenburg, G., Dighe, A., Dorigatti, I., Fu, H., Gaythorpe, K., Green, W., Hamlet, A., Hinsley, W., Okell, L. C., van Elsland, S., Thompson, H., Verity, R., Volz, E., Wang, H., Wang, Y., Walker, P. G. T., Walters, C., Winskill, P. Whittaker, C., Donnelly, A., Riley, S., Ghani, A. C. (2020) Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand. *Imperial College Response Team*, http://hdl.handle.net/10044/1/77482.

Fine, P., Eames, K., and Heymann, D. L. (2011). "Herd immunity": a rough guide. *Clinical Infectious Diseases*, *52*(7), 911-916.

Finkelstein, A., Luttmer, E. F., and Notowidigdo, M. J. (2013). What good is wealth without health? The effect of health on the marginal utility of consumption. *Journal of the European Economic Association*, *11*(suppl_1), 221-258.

Fox, J. A., Hayes, D. J., and Shogren, J. F. (2002). Consumer preferences for food irradiation: how favorable and unfavorable descriptions affect preferences for irradiated pork in experimental auctions. *Journal of Risk and Uncertainty*, *24*(1), 75-95.

Freed, G. L., Clark, S. J., Butchart, A. T., Singer, D. C., and Davis, M. M. (2011). Sources and perceived credibility of vaccine-safety information for parents. *Pediatrics*, *127*(Supplement 1), S107-S112.

Covid Economics 35, 7 July 2020: 1-50

Frewer, L. J., Howard, C., Hedderley, D., and Shepherd, R. (1996). What determines trust in information about food-related risks? Underlying psychological constructs. *Risk Analysis*, *16*(4), 473-486.

Gerking, S., Adamowicz, W., Dickie, M., and Veronesi, M. (2017). Baseline risk and marginal willingness to pay for health risk reduction. *Journal of Risk and Uncertainty*, *55*(2-3), 177-202.

Hämeen-Anttila, K., Nordeng, H., Kokki, E., Jyrkkä, J., Lupattelli, A., Vainio, K., and Enlund, H. (2014). Multiple information sources and consequences of conflicting information about medicine use during pregnancy: a multinational Internet-based survey. *Journal of Medical Internet Research*, *16*(2), e60, doi: 10.2196/jmir.2939

Hammitt, J. K., and Haninger, K. (2010). Valuing fatal risks to children and adults: Effects of disease, latency, and risk aversion. *Journal of Risk and Uncertainty*, *40*(1), 57-83.

Hethcote, H. W. (2000). The mathematics of infectious diseases. *SIAM review*, *42*(4), 599-653.

Hoehn, J. P., and Randall, A. (2002). The effect of resource quality information on resource injury perceptions and contingent values. *Resource and Energy Economics*, *24*(1-2), 13-31.

Imbens, G. W., and Rubin, D. B. (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.

Johns Hopkins University and Medicine (2020), Mortality Analyses, https://coronavirus.jhu.edu/data/mortality, Retrieved June 26 2020.

Covid Economics 35, 7 July 2020: 1-50

Jordan, J., Yoeli, E., and Rand, D. G. (2020, April 3). Don't get it or don't spread it? Comparing self-interested versus prosocially framed COVID-19 prevention messaging. Doi: 10.31234/osf.io/yuq7x

Keeling, M.J. and Rohani, P. (2008). *Modeling infectious diseases in humans and animals*. Princeton, NJ: Princeton University Press.

Kelly, D. L., Letson, D., Nelson, F., Nolan, D. S., and Solís, D. (2012). Evolution of subjective hurricane risk perceptions: A Bayesian approach. *Journal of Economic Behavior & Organization*, *81*(2), 644-663.

Kuehn B. (2019) Measles Vaccine Exemptions. *JAMA*, 321(14):134. doi:10.1001/jama.2019.3063

Kwok, K. O., Lai, F., Wei, W. I., Wong, S. Y. S., and Tang, J. W. (2020). Herd immunity–estimating the level required to halt the COVID-19 epidemics in affected countries. *Journal of Infection*, Forthcoming.

Larson, H. J. (2018). The state of vaccine confidence. *The Lancet*, *392*(10161), 2244-2246.

Le, T.T., Andreadakis, Z., Kumar, A., Gómez Román, R., Tollefsen, S., Saville, M., and Mayhew, S. (2020) The COVID-19 vaccine development landscape. *Nature Reviews Drug Discovery*. Retrieved April 11. doi: 10.1038/d41573-020-00073-5.

Leask, J. (2011). Target the fence-sitters. *Nature*, *473*(7348), 443-445.

Lee, C., Whetten, K., Omer, S., Pan, W., and Salmon, D. (2016). Hurdles to herd immunity: Distrust of government and vaccine refusal in the US, 2002–2003. *Vaccine*, *34*(34), 3972-3978.

Covid Economics 35, 7 July 2020: 1-50

Liu, Y., Gayle, A. A., Wilder-Smith, A., and Rocklöv, J. (2020) The reproductive number of COVID-19 is higher compared to SARS coronavirus. *Journal of Travel Medicine,* 27(2). Doi: 10.1093/jtm/taaa021

LX/Morning Consult (2020), National Tracking Poll #200395: March 24-25, 2020. https://www.nbcwashington.com/wp-content/uploads/2019/09/LX_COVID_Vaccine_Poll_crosstabs.pdf. Retrieved June 22 2020.

MacDonald, N. E. (2015). Vaccine hesitancy: Definition, scope and determinants. *Vaccine*, *33*(34), 4161-4164.

Magat, W. A. and Viscusi, W. K. (1992). *Informational approaches to regulation* (Vol. 19). MIT press.

Mochiri, N. (2020) Coronavirus seems to mutate much slower than seasonal flu. https://www.livescience.com/coronavirus-mutation-rate.html. Retrieved April 28 2020.

MSN News, https://www.msn.com/en-us/news/politics/analysis-the-white-house-continues-to-downplay-the-coronavirus-threat-to-much-of-the-country/ar-BB12evZY. Retrieved April 20.

Mutz, D.C., Pemantle, R, and Pham, P. (2017). The perils of balance testing in experimental design: messy analyses of clean data. *The American Statistician*, 73(1), 32-42.

Olive, J. K., Hotez, P. J., Damania, A., and Nolan, M. S. (2018). The state of the antivaccine movement in the United States: A focused examination of nonmedical exemptions in states and counties. *PLoS Medicine*, *15*(6).

Covid Economics 35, 7 July 2020: 1-50

Pushkarskaya, H., Smithson, M., Joseph, J. E., Corbly, C., and Levy, I. (2015). Neural correlates of decision-making under ambiguity and conflict. *Frontiers in Behavioral Neuroscience*, *9*, 325, doi: 10.3389/fnbeh.2015.00325

Rémy, V., Largeron, N., Quilici, S., and Carroll, S. (2015). The economic value of vaccination: why prevention is wealth. *Journal of Market Access & Health Policy*, *3*(1), doi: 10.3402/jmahp.v3.29284.

Rodgers, W. (1999). The influences of conflicting information on novices and loan officers' actions. *Journal of Economic Psychology*, *20*(2), 123-145.

Roser, M., Ritchie, H., Ortiz-Ospina, E. and Hasell, J. (2020) Coronavirus Pandemic (COVID-19). *Published online at OurWorldInData.org.* https://ourworldindata.org/coronavirus. Retrieved April 28 2020.

Rousu, M. C., and Shogren, J. F. (2006). Valuing conflicting public information about a new technology: the case of irradiated foods. *Journal of Agricultural and Resource Economics*, 642-652.

Pew Research Center (2020) Most Americans expect a COVID-19 vaccine within a year; 72% say they would get vaccinated, https://www.pewresearch.org/fact-tank/2020/05/21/most-americans-expect-a-covid-19-vaccine-within-a-year-72-say-they-would-get-vaccinated/. Retrieved June 22 2020.

Pindyck, R.S. (2020). COVID-19 and the welfare effects of reducing contagion. https://mitsloan.mit.edu/sites/default/files/2020-04/Contagion2020.pdf, Retrieved May 4 2020.

Covid Economics 35, 7 July 2020: 1-50

Sanche S., Lin, Y.T., Xu, C., Romero-Severson, E., Hengartner, N., and Ke, R. (2020). High contagiousness and rapid spread of severe acute respiratory syndrome coronavirus 2. *Emerging Infectious Disease*. 2020 Jul [*11 April 2020*]. doi: 10.3201/eid2607.200282

Sarkar, S., Zlojutro, A., Khan, K., and Gardner, L. (2019). Measles resurgence in the USA: how international travel compounds vaccine resistance. *The Lancet Infectious Diseases*, *19*(7), 684-686.

Sell, T. K. (2017). When the next disease strikes: How to communicate (and how not to). *Health Security*, *15*(1), 28-30.

Serra-Garcia, M., and Szech, N. (2020). Demand for COVID-19 Antibody Testing, and Why It Should Be Free, CESifo Working Paper No. 8340, https://ssrn.com/abstract=3623675

Sharpe Wessling, K., Huber, J., and Netzer, O. (2017). MTurk character misrepresentation: Assessment and solutions. *Journal of Consumer Research*, 44(1), 211-230.

Shereen, M. A., Khan, S., Kazmi, A., Bashir, N., and Siddique, R. (2020). COVID-19 infection: origin, transmission, and characteristics of human coronaviruses. *Journal of Advanced Research*.

Siegrist, M., Cvetkovich, G., and Roth, C. (2000). Salient value similarity, social trust, and risk=benefit perception. *Risk Analysis*, 20(3), 353–362.

Siegrist, M., Cvetkovich, G. T., and Gutscher, H. (2001). Shared values, social trust, and the perception of geographic cancer clusters. *Risk Analysis*, 21(6), 1047–1053.

Covid Economics 35, 7 July 2020: 1-50

Simonov, A., Sacher, S. K., Dubé, J. P. H., and Biswas, S. (2020). The persuasive effect of fox news: non-compliance with social distancing during the covid-19 pandemic (No. w27237). National Bureau of Economic Research.

Sloan, F. A., Viscusi, W. K., Chesson, H. W., Conover, C. J., and Whetten-Goldstein, K. (1998). Alternative approaches to valuing intangible health losses: the evidence for multiple sclerosis. *Journal of Health Economics*, *17*(4), 475-497.

Smith, V. K., and Desvousges, W. H. (1987). An empirical analysis of the economic value of risk changes. *Journal of Political Economy*, *95*(1), 89-114.

Thunström, L., Ashworth, M., Shogren, J. F., Newbold, S., and Finnoff, D. (2020a). Testing for COVID-19: Willful ignorance or selfless behavior?. *Behavioural Public Policy*, 1-18, doi:10.1017/bpp.2020.15.

Thunström, L., Newbold, S., Finnoff, D., Ashworth, M., Shogren, J.F. (2020b). The benefits and costs of social distancing to flatten the curve for Covid-19, *Journal of Benefit-Cost Analysis,* 1-17, doi:10.1017/bca.2020.12.

Viscusi, W. K. (1997). Alarmist decisions with divergent risk information. *The Economic Journal*, *107*(445), 1657-1670.

Viscusi, W. K., and Aldy, J. E. (2003). The value of a statistical life: a critical review of market estimates throughout the world. *Journal of Risk and Uncertainty*, *27*(1), 5-76.

Viscusi, W. K., and Evans, W. N. (1990). Utility functions that depend on health status: estimates and economic implications. *The American Economic Review*, 353-374.

Covid Economics 35, 7 July 2020: 1-50

Viscusi, W. K., and Magat, W. A. (1992). Bayesian decisions with ambiguous belief aversion. *Journal of Risk and Uncertainty*, *5*(4), 371-387.

Viscusi, W. K., Magat, W. A., and Huber, J. (1999). Smoking status and public responses to ambiguous scientific risk evidence. *Southern Economic Journal*, *66*(2), 250-270.

World Health Organization (WHOa), https://www.who.int/news-room/feature-stories/ten-threats-to-global-health-in-2019, Retrieved April 21 2020.

World Health Organization (WHOb), https://www.who.int/news-room/commentaries/detail/immunity-passports-in-the-context-of-covid-19, Retrieved April 29 2020.

Covid Economics 35, 7 July 2020: 1-50

A simple method to estimate firms' liquidity needs during the Covid-19 crisis with an application to Italy[1]

Fabiano Schivardi[2] and Guido Romano[3]

Date submitted: 1 July 2020; Date accepted: 2 July 2020

We propose a simple method based on firms' balance sheets and sectoral predictions of sales growth to determine the firms that will become illiquid month by month as the Covid-19 crisis unfolds. We apply the method to the population of Italian incorporated businesses to the end of 2020. We find that at the peak, around 200,000 companies employing 3.3 million workers would become illiquid. The progression is fast, with 180,000 firms turning illiquid already by April. The liquidity shortage, defined as the "negative" liquidity stock of illiquid firms, amounts to 72 billion. We evaluate the Italian government liquidity decree, which provides guarantees for bank loans under four different facilities of increasing complexity. Assuming that firms have access to all the facilities, almost all firms are able to cover their liquidity shortfalls. The issue is the speed of implementation: the facilities supplying more liquidity are more complex to administrate, and many firms require these facilities to cover their liquidity shortfalls. Overall, we conclude that even in the case of a second wave after the summer, which would increase the liquidity shortfall substantially, firms' liquidity needs are manageable under the current schemes of liquidity provision.

1 We are grateful to webinar participants at the OECD and the ECB for useful comments. We are also grateful to Letizia Sampoli and her team of Cerved sector analysts for providing us sales forecasts for more than 500 sectors. David Kwon provided superb research assistance. The views expressed in the paper are those of the authors and do not necessarily reect those of Cerved Group.

2 Luiss University, EIEF and CEPR.

3 Chief Economist, Cerved Group Research Department.

Covid Economics 35, 7 July 2020: 51-69

1 Introduction

A fundamental question to predict the economic effects of the Covid-19 pandemic is to understand their persistence, that is, if the economy will return quickly to its pre-crisis level or if it will take a long time to reabsorb the fall in output. This depends on how many companies will go bankrupt from the liquidity crisis due to the fall in sales. Bankruptcies have long lasting effects, prolonging the negative consequences of the shock. They amplify the real (through input/output relationships) and financial (through trade payables and receivables) contagion to other companies, which can have a chain effect on the entire economy. When this happens, bad loans grow, and the infection is extended to the financial sector. Therefore, it is important to provide firms with liquidity to avoid bankruptcies at this scale. This is a shared goal, and the response by policymakers has generally been to provide *whatever it takes*. In fact, most governments have set up some form of credit guarantee, particularly for small and medium enterpreises (SME) OECD (2020). However, these policies need to be credible; therefore, it is important to determine *how much it takes*: are the schemes that governments provide sufficient to avoid massive liquidity-induced bankruptcies?

In this paper, we develop a simple accounting framework to determine which firms will have liquidity constraints and to what extent. The general logic is very straightforward and is based on three ingredients: the initial stock of liquidity, an estimate of the evolution of cash flow month by month and the budget equation determining the evolution of liquidity. The framework uses firms balance sheets to obtain pre-pandemic output and costs as well as the initial stock of liquidity. It requires as an input the month by month estimate of sales growth at the sectoral level which, given each firm's previous year sales, make it possible to forecast sales evolution at the level of the firm. Costs are predicted using inputs' elasticities, which allow us to use sales growth forecasts–mediated by the elasticity of each input–to determine monthly outflows. Given an initial stock of liquidity, the budget equation determines the stock of liquidity month by month. When this value turns negative, we classify a firm as illiquid. The absolute value of the negative liquidity is the amount of the liquidity shortfall. Summing across all illiquid firms produces the aggregate liquidity needs to avoid firms going bankrupt. The method is transparent and straightforward to implement.[1] It has been used by various institutions (Bank of Italy 2020, European Commision 2020, OECD 2020).

We apply the method to the population of Italian incorporated businesses, around 650,000 companies producing three quarters of the Italian private sector output. We consider the period from March 2020 until the end of the year. Sales growth from more than 500 sectors is forecasted by Cerved, a data provider and credit rating agency that also supplies firm balance sheets. The forecasts are carried out by Cerved sectoral experts, who take into account both the legislation (the lockdown) and other economic factors (drop in demand, effects of social distancing, disruption of supply chains etc.). We set all financial outflows and tax payments to zero, based on legislative decrees that allow firms to postpone them. Using time series data and taking into account a job retention scheme enacted by the government, we assume that the elasticity of intermediate goods and services expenditure to sales is 0.5 and that of labor is 0.75. In the baseline scenario, the lockdown is active from mid-March to the beginning of May for non essential sectors. Then, activity gradually reverts back to normal at different speeds according to sectoral characteristics. We experiment with both input costs elasticities and the evolution of the pandemic, allowing for a second wave in the Fall.

We find that the effects of the pandemic are very quick, with more than 180,000 firms, employing 3.1 million workers, already becoming illiquid in April. The number of illiquid

[1] A brief note illustrating it was posted on the economists blog lavoce.info at the end of March and a code with a mock dataset is available http://docenti.luiss.it/schivardi/policy-wor/policy-work/.

Covid Economics 35, 7 July 2020: 51-69

firms peaks at 200,000 (employing 3.3 million workers) in September, and then it decreases very slightly for the rest of the year. The amount of liquidity shortage, that is, the value of "negative" liquidity of illiquid firms, is 40 billion in April. Then, it keeps increasing until the end of the year, when it reaches 72 billion. Of these, more than 50 million are in firms with less than 500 employees.

Next, we use the scheme to evaluate the coverage provided by the Italian Liquidity Decree, which supplies public guarantees to bank loans issued as a response to the pandemic. The decree designs different facilities to provide loan guarantees, with the amount of guarantee decreasing with the size of the loan and with different conditions for small (less than 500 employees) and large firms. We find that theoretical coverage is complete: assuming that firms have access to all the facilities, basically all firms are able to cover their liquidity shortfalls. The issue then becomes implementation: in fact, the semi-automatic facilities supply limited coverage, and only the more complex ones bring coverage close to full. However, the complex facilities require both bank screening (as the guarantee is not full) and an approval from a government agency. Given that we forecast that many firms will become illiquid quickly, there could be a congestion effect that prevents the full theoretical coverage from becoming actualized. The key issue is therefore the speed of implementation. We propose a simple scheme, according to which firms with good pre-crisis rating are granted credit semi-automatically so that banks can deploy their screening capabilities for firms with less solid rating, distinguishing illiquid but solvent firms from insolvent firms.

As stated above, our methodology has been employed by several economic institutions (Bank of Italy 2020, European Commision 2020, OECD 2020) with comparable results. We are only aware of three papers independently developed and related to our work. Carletti, Oliviero, Pagano, Pelizzon & Subrahmanyam (2020) estimate the drop in profits and the equity shortfall following the Covid-19 crisis for a sample of Italian firms, finding that they are substantial. Compared to us, they focus on profits rather than liquidity, use a smaller sample of 81,000 firms (our sample of 650,000 firms is the universe of incorporated firms) and use a different methodology to estimate demand and costs. In particular, they assume zero elasticity for intermediate goods and services (intermediates in what follows) which, as we show in our exercise, has a strong influence on the estimates. De Vito & Gomez (2020) focus on listed firms in 26 countries. Compared to our sample with a large majority of private firms, listed firms on average hold more liquidity, so the share of firms becoming illiquid early on is much smaller: only 10% within six months under the most adverse scenario, against 31% in our case. This is in line with the idea that listed firms are financially more solid and lends support to the view that government schemes should target SMEs first. McGeever, McQuinn & Myers (2020) carry out a similar exercise for Ireland. Compared to us, they have no information on the firm's stock of liquidity before the crisis and only focus on a subset of highly affected sectors, for which they assume that sales go completely to zero for three months. Also their treatment of costs differ. They estimate liquidity needs of between 2.5 and 5.7 billion. In related work, Schivardi, Sette & Tabellini (2020) use our framework to analyze the possibility that government loan guarantees might induce "zombie" lending, that is, the provision of credit to firms that were already insolvent before the crisis. They conclude that due to the nature of the shock, which hits firms independently from their economic and financial conditions, the amount of "zombie" lending is likely to be limited.

The rest of the paper is organized as follows. Section 2 illustrates the method and describes the preferred parameterization. Section 3 applies it to the Italian population of incorporated businesses and Section 4 evaluates the Italian Liquidity Decree. Section 5 explores alternative parameterizations in terms of the evolution of the pandemic and the elasticities of inputs and Section 6 concludes.

2 The method

In this section, we illustrate the accounting scheme and the choice of the parameters and forecasts to implement it.

2.1 The accounting scheme

We construct an accounting framework that allows us to estimate the liquidity needs of firms during the Covid-19 crisis. The framework identifies the month in which a firm becomes illiquid (if any) and the amount of the liquidity shortage afterwards. The general logic is simple and is based on three ingredients:

1. the initial stock of liquidity in the firm's balance sheets
2. an estimate of the evolution of cash flow month by month
3. the budget equation governing the evolution of liquidity.

Specifically, for firm i in month m of 2020, given an initial stock of liquidity L_{i0} in February, sales S_{im} and outlays C_{im}, the evolution of liquidity L_{im} is:

$$L_{im} = L_{im-1} + S_{im} - C_{im} \tag{1}$$

for m =March, April,..., December 2020. To implement Equation 1, we therefore need L_{i0} as well as the monthly evolution of cash flow (sales minus costs). Based on this, we can determine the month in which L_{im} turns negative. In this case, a firm is defined as illiquid. For each month m, the total (that is, for the whole economy) liquidity shortage TLS is the sum of all the liquidity shortages of illiquid firms:

$$\text{TLS}_m = \sum_{L_{im}<0} |L_{im}|. \tag{2}$$

The method requires information on firms' balance sheets. Balance sheets are typically available for all listed firms. In many countries, they are also available for unlisted incorporated firms. The Orbis database of Beurau Van Dick contains data for many countries and is widely used in research (Kalemli-Ozcan, Sorensen, Villegas-Sanchez, Volosovych & Yesiltas 2015). In our application to Italy, we use the Cerved database, which covers all incorporated Italian businesses. By law, they are obliged to file their balance sheets every year to the Firm Registry. We focus on the 650,000 non-financial firms which account for approximately three quarters of private sector GDP. We use the most recently available balance sheets, which are from 2018. Ideally, one would use those of 2019; however, note that this is not a major limitation. In fact, while the exercise predicts liquidity needs firm by firm, we are interested in the aggregate values. As long as the distribution of firms' conditions is invariant between 2018 and 2019, the aggregate results will be unaffected by idiosyncratic firm movements that leave the distribution unchanged.

From the balance sheets, we obtain the initial value of liquidity, defined as the value of liquid assets reported in the balance sheets. For sales, we consider the sales of 2018 and assume that absent the Covid-19 crisis, monthly sales would have been equal to 1/12 of the total sales of 2018. We then apply forecasts of sales growth described in more detail in the next subsection. In terms of firm costs, we assume that following the Italian government decrees enacted during the crisis, all financial payments and taxes are suspended. Moreover, we also assume that firms freeze their investment expenditures.[2] The only outlays left are

[2]Accounting for financial payments is straightforward, and in an initial version of the procedure implemented, before the government had frozen financial payments, we had taken them into account. We assumed that each month, a firm had to pay 1/12 of the interest expenses reported in the balance sheets of 2018 and 1/12 of the mortgage payments, which were estimated as a fraction of long term debt.

Covid Economics 35, 7 July 2020: 51-69

cost of labor and intermediates. To estimate them, we use the elasticity of each input to sales, which allows us to determine the evolution of costs from that of sales. Given an elasticity of each input to sales, $\varepsilon_{WS}, \varepsilon_{MS}$, and sectoral estimates of the drop in sales for each month compared to the pre Covid-19 value d_{im}, Equation 1 becomes:

$$L_{im} = L_{im-1} + (1 - d_{im})S_i - (1 - \varepsilon_{WS} d_{im}) * W_i - (1 - \varepsilon_{MS} d_{im}) * M_i \tag{3}$$

where $L_i = L_{i2018}/12$ is monthly sales according to 2018 sales and similarly for labor costs W_i and intermediates M_i. To implement Equation 3, we need to determine the values of d_{im}, ε_{WS} and ε_{MS}.

2.2 Sales forecasts and inputs elasticities

Sales forecasts for approximately 500 sectors were produced by Cerved sector experts. Cerved is a data provider and credit rating agency that computes firms default probability (the score) used by banks to process credit applications. Sectoral forecasts are used by Cerved in their predictive exercises. In the basic scenario, the lockdown lasts until the beginning of May, as it actually did, and applies at the sectoral level according to the various Government decrees that were issued during the acute phase of the pandemic. The key distinction is between essential sectors (food, health, delivery), which were allowed to continue production, and non essential sectors, which had to shut down. The scenario then assumes a period of partial opening that also varies by sector, and after that, activity gradually recovers. In addition to the legal constraints, the sectoral estimates also take into account sectoral exposure to Covid-19 specific effects, such as the possibility to work remotely, the effects of social distancing, the reduction in mobility, etc. The appendix reports a detailed description of the procedure, as well as a 2-digit aggregation of sales growth for 2020 with respect to 2019. The most affected sector is air transportation, which records a drop of 46%. At the opposite end, online retail trade increases by 30% (see Appendix Table 3). In the base scenario, it is assumed that the pandemic gradually disappears. Cerved also computed a pessimistic scenario in which the virus returns in the Fall. We evaluate the effects of the pessimistic scenario on firms' liquidity in Section 5.1.

Table 1 reports descriptive statistics of firms, dividing them according to the predicted sectoral drop in sales for 2020 with respect to 2019. We separate sectors in groups with a drop of 20% or larger, between 20 and 10%, between 10 and 0%, and with non negative sales growth. The group with the largest drop is by far the most populated, with more than 300,000 firms. The two intermediate groups have approximately 130,000 firms, and less than 60,000 firms record sales increases. Firm characteristics are very similar across groups: while the mean values of employment, sales and liquidity differ somewhat due to the skewed distribution with some large outliers, the 25^{th}, 50^{th} and 75^{th} percentiles are remarkably similar, with a large prevalence of small firms in all clusters: the 75^{th} percentile of employment varies between 7 and 9. Average liquidity is around 400,000 euros, but the median is much slower, at around 30,000 euros in all groups, while the 25^{th} percentile is always below 10,000 euros. This indicates that many firms have small liquidity buffers.

Next, we consider two indicators of financial fragility: leverage, which is defined as debt over equity, and Cerved Group Credit Score, a riskiness indicator computed by Cerved that takes values from 1 (very safe) to 10 (very risky).[3] As it turns out, firms are also very similar across groups in terms of financial conditions. The statistics in Table 1 therefore indicate that, not surprisingly, the crisis hit sectors with an intensity uncorrelated to sectoral characteristics, at least related to size and financial health.

[3]Given that leverage is very skewed due to the large number of small firms with either no debt or very low equity, we trim leverage at the 1^{st} and 99^{th} percentiles.

Covid Economics 35, 7 July 2020: 51-69

Table 1: Firm Statistics by Change in Sales

	Mean	p25	Median	p75	S.D.	N. Obs.
> 20% Decrease						
Employees	13.5	0	2	8	294	316,529
Sales	3,797	146	452	1,404	104,170	316,855
Liquidity	419	7	28	112	30,846	316,529
Leverage	1.09	0	0	0.757	2.9	291,176
Risk Class	5.43	4	5	6	1.64	311,238
10 − 20% Decrease						
Employees	14.1	0	2	7	161	135,470
Sales	4,278	107	348	1,143	109,171	135,777
Liquidity	448	7	29	118	16,568	135,470
Leverage	0.907	0	0	0.494	2.64	127,563
Risk Class	5.21	4	5	6	1.54	133,717
0 − 10% Decrease						
Employees	21.6	0	3	9	349	127,882
Sales	3,131	105	318	1,042	57,506	127,992
Liquidity	320	7	27	105	6,725	127,882
Leverage	0.854	0	0	0.444	2.52	118,746
Risk Class	5.05	4	5	6	1.58	126,359
10% Increase						
Employees	17.4	0	2	7	233	57,917
Sales	6,489	154	585	2,209	70,270	57,953
Liquidity	583	8	34	135	16,072	57,917
Leverage	1.19	0	0	0.891	3.05	53,110
Risk Class	5.42	4	5	6	1.66	56,961

Note: The table reports descriptive statistics for firms characteristics, split according the drop in sales. Leverage is debt over equity and is trimmed at the 1st and 99th percentile. Risk class is from the Cerved Credit Score, which takes discrete values between 1 (very safe) to 10 (very risky), with unit intervals.

The last input we need is a value for the two elasticities of labor and intermediates to sales. To obtain a rough estimate of the two values, we use the balance sheets of Italian non financial incorporated companies between 2005 and 2015, which contains a total of 3.9 million firm-year observations. We regress the percentage annual change (the log difference) in intermediate expenditure and the wage bill on the percentage change in sales, controlling for year and firm fixed effects. For intermediates, we obtain $\varepsilon_{MS} = 0.70$. While purchases of goods are highly elastic, due to the strong pro-cyclical behavior of inventories (Khan & Thomas 2007), services expenditure, which includes rents and fixed contracts for the provision of telecom services, royalties etc., is more difficult to cut in the short run. Moreover, these estimates are based on all changes in sales, including both small and positive ones, but during the Covid-19 crisis the most important changes are large and negative. In the presence of non-linearities and asymmetries, estimates based on the whole sample might not adequately capture the response to a large negative shock. To check how firms respond to such a shock, we repeated the regressions using only observations for which the change in sales was below -0.1. The number of observations drops to 1.05 million and the elasticity to 0.62, indicating that large negative shocks are more difficult to accommodate. We therefore assume a conservative value for the elasticity of intermediates of 0.5.[4]

[4]We are aware of the fact that our estimates do not account for endogeneity. However, we only use them

Covid Economics 35, 7 July 2020: 51-69

Covid Economics 35, 7 July 2020: 51-69

For labor, the same exercises as above delivers $\varepsilon_{WS} = 0.46$ for the whole sample and $\varepsilon_{WS} = 0.40$ when focusing on sales drops of at least 10%. However, the labor elasticity during the Covid-19 crisis critically depends on public supplement schemes. The Italian government provides a job retention scheme that allows all firms to reduce paid work by any amount and have the government pay workers for the income loss (the *Cassa integrazione guadagni* or the Fund to integrate income). This greatly increases the elasticity of the wage bill to production. The scheme has been extensively used by Italian firms: in April, the number of hours paid was almost 900 million, equal to the *total* amount paid in 2009, the worst year of the financial crisis, when GDP contracted by more than 5%. To account for this, we set ε_{WS} to 0.75. This measure is renewed on a monthly basis, so it might be discontinued before the end of the year. For simplicity, we assume that both elasticities are constant throughout the exercise. Changing the values of the elasticities obviously affects the absolute vales but not the general conclusions, as we show in some robustness exercises below.

One final point is that input elasticities are likely to be asymmetric. We do not allow for full adjustment – elasticities are smaller than 1 – for sales drop, allowing for frictions in reducing inputs in the short run. However, for the few sectors that expand sales, assuming low elasticities will boost cash flow, as the firm is allowed to expand sales with proportionally smaller increases in costs. This might be sensible to the extent that some of the costs have a fixed component. At the same time, elasticities for increases in output might be higher than those for decreases.[5] To take this into account, we have experimented by assuming that ε_{WS} and ε_{MS} are smaller than one only when sales contract and equal to 1 when sales expand. Unit elasticity is consistent with a constant return to scale production function and no fixed costs.[6] We have experimented with this asymmetric parameterization, finding that it makes very little difference because, as shown in Appendix Table 3, very few sectors increase their sales during 2020.

3 Results

We now apply our scheme to the universe of Italian incorporated companies. For each month, we compute the number of firms that are illiquid (firms for which $L_{im} < 0$) and the workers employed by these firms and plot their evolution from March to December of 2020 in Figure 1. The crisis is clear very quick: already in April, 180,000 firms, employing 3.1 million workers, become illiquid. The peak is reached in September, with 201,946 firms, and then the number decreases very slowly for the rest of year. In terms of workers, in September, 3.5 millions are employed by illiquid firms, which amounts to 12% of total Italian employment.

Figure 2, Panel a) plots the total liquidity shortage (TLS) defined in Equation 2 as a total and separately for firms above and below the 500 employee threshold.[7] The TLS is 12 billion in March, jumps to 40 billion in April and increases steadily until September by

to get a rough idea of the magnitude and interpret them conservatively to account for the exceptional nature of the situation. In Section 5.2, we experiment with alternative values.

[5] Regression estimates only using positive sales changes deliver a slightly larger elasticity for intermediates and a lower elasticity for labor.

[6] Note that we are not accounting for the cost of capital. That is, we do not impute any capital costs in case of sales growth. Using survey data for Italian manufacturing firms with information on capacity utilization, Pozzi & Schivardi (2016) show that the average degree of capacity utilization is 81%, with a standard deviation of 13%; the 5^{th} and the 95^{th} percentile are 60 and 98%. This suggests that most firms will be able to deliver the small increases in sales we predict for a few sectors without resorting to increases in the capital stock and therefore without cash outflows.

[7] We use the 500 employee threshold to define SMEs because this is the definition used by the Italian government in the liquidity decree that we analyze below.

Figure 1: Illiquid firms and workers

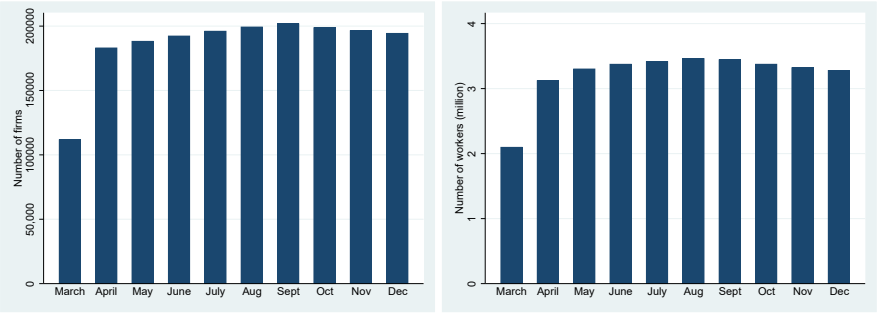

(a) Number of illiquid firms (b) Workers in illiquid firms

Note: The figure reports the total number of illiquid firms (Panel a) and the total number of workers in such firms (Panel b) using Equation 3 to detect the firms for which liquidity has hit the zero constraint.

Covid Economics 35, 7 July 2020: 51-69

approximately 5 billion per month. After that, the growth slows to 2 billion per month, reaching a peak of 72 billion in December. SMEs account for 2/3 of the total at the beginning of the crisis and for 3/4 of the total in December.

Figure 2: Total liquidity shortage for all firms and by firm size

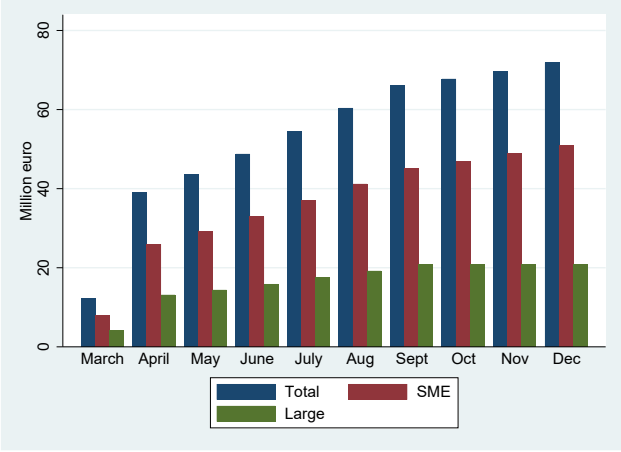

Note: The figure reports the value of the total liquidity shortage (TLS) defined in Equation 4 for all firms and distinguishes between firms above and below the 500 employees threshold.

TLS measures the value of the "negative" liquidity accumulated by firms that hit the zero liquidity constraint. As such, it does not include the value of liquidity that firms had before the crisis that is lost before hitting the zero liquidity constraint. To account for this, we also compute the total liquidity loss (TLL), defined as the total liquidity lost by illiquid firms:

$$\text{TLL}_m = TLS + \sum_{L_{im}<0} L_{i0}. \tag{4}$$

where L_{i0} is the initial liquidity stock. Figure 3 reports the evolution of TLL and TLS to

facilitate comparisons. In March, TLS is 12 billion and TLL is 17 billion. The difference grows over time and reaches a maximum of 22 billion in December. Therefore, accounting for the liquidity lost makes a difference, as illiquid firms do deplete a substantial liquidity stock. However, TLS represents the largest component of TLL.

Figure 3: Total liquidity loss

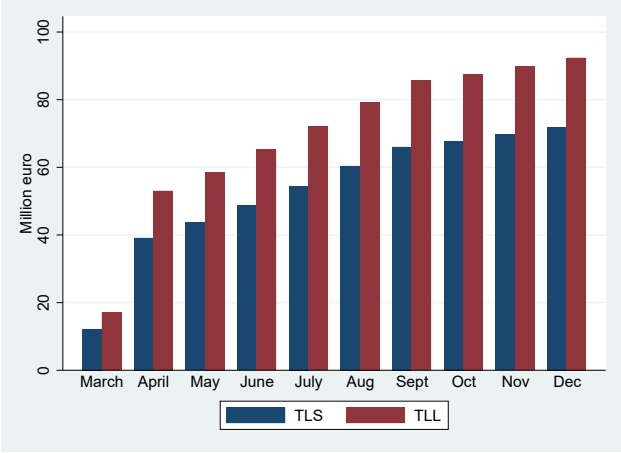

Note: The figure reports the values of the total liquidity shortage (TLS) and of total liquidity loss (TLL).

Table 2 reports firms' characteristics separately for liquid and illiquid firms as of December 2020. The two groups are very similar in terms of size, with illiquid firms slightly larger both in terms of sales and employment. Illiquid firms also have more trade credit and less debt at the mean. However, their equity is lower: the mean is 1.2 million (median 47,000) against 1.7 million (median 86,000) for liquid ones. Not surprisingly, the biggest difference emerges in terms of the stock of liquidity, which is less than one fourth of that of liquid firms both at the mean and at the median. Lower equity and liquidity implies that illiquid firms were ex ante slightly more risky, with an average leverage of 1.56 against 0.8 (but the median firm has zero leverage in both groups) and an average risk class of 5.9 against 5.0. All in all, these values indicate that firms that turn illiquid were more financially fragile before the crisis. However, the difference is mostly in the liquidity holdings, while other characteristics are relatively similar.

4 Evaluating the Italian government liquidity guarantee decree

The response of many governments to firms' liquidity needs following the Covid-19 crisis has been to set up schemes of credit guarantees for bank loans, particularly to SMEs. According to an OECD survey of policies enacted by governments in 54 countries to contrast the crisis, as of April 20th, 2020, 52 of them had set up some form of government-provided financial support for SMEs (see OECD 2020, Table 3). The analytical framework described in the previous Section allows us to evaluate the extent of coverage provided by such schemes. We therefore use it to analyze the coverage provided by the Italian scheme, set up in Decree n. 23 of April 8, 2020 (*Decreto liquidità*). The decree offers public guarantees that decrease with the amount of the loan. In particular, for firms with less than 500 employees (SMEs

Covid Economics 35, 7 July 2020: 51-69

Table 2: Liquid and Illiquid Firm Characteristics

	Mean	p25	Median	p75	S.D.	N.Obs.
Liquid						
Employees	15	0	2	7	297	443,258
Sales	3,598	119	370	1218	96,927	444,037
Commercial Credit	835	0	5	231	17,911	443,258
Equity	1,725	24	86	350	73,183	444,029
Debt	1,386	0	0	71	126,481	443,911
Liquidity	559	12	44	164	28,448	443,258
Leverage	0.798	0	0	0.392	2.45	423,587
Risk Class	5.04	4	5	6	1.54	437,481
Illiquid						
Employees	16.9	1	3	9	233	194,540
Sales	4,951	155	503	1,598	90,533	194,540
Commercial Credit	1,048	0	10	314	13,064	194540
Equity	1,216	11	47	214	80,356	194,540
Debt	937	0	0	195	14,636	194,540
Liquidity	105	3	11	39	1,336	194,540
Leverage	1.56	0	0	1.5	3.45	167,008
Risk Class	5.9	5	6	7	1.61	190,794

Note: The table reports descriptive statistics for firms' characteristics before the crisis separately for liquid and illiquid firms. Leverage is debt over equity and is trimmed at the 1st and 99th percentile. Risk class are from the Cerved Credit Score, which takes discrete values between 1 (very safe) to 10 (very risky), with unit intervals.

in what follows), it offers:

- **Measure 1**: Full guarantee up to the minimum between 30.000 and 25% of 2019 sales.
- **Measure 2**: For firms with less than 3,2 million turnover, 25% of 2019 sales, with 90% government guarantee and 10% Confidi (an association for mutual guarantees) guarantee.
- **Measure 3**: Up to 5 million with 90% government guarantee.
- **Measure 4**: Up to the maximum between 25% of sales and twice the labor costs of 2019, with a guarantee from 90% to 70% according to firm and loan size.

Measures 3 and 4 require the approval of a Government agency. Firms with more than 500 employees have access to Measure 4 only.

The government claims that this scheme mobilizes 400 billion euros which, according to the numbers seen in the previous section, should be more than enough to cover liquidity shortages. To check for this, we let firms borrow the maximum amount according to the measures above and check which firms cannot cover their liquidity shortage with such borrowing. Figure 4 shows that coverage is indeed complete: at peak, just 153 firms, employing less than 13,000 workers, cannot cover their liquidity shortages.

Of course, this is the maximum theoretical coverage assuming that firms have access to the maximum loan supply the decree allows for. However, as discussed above, the procedural complexity of the measure increases with the amount it supplies. For example, Measure 1 offers SMEs up to 30,000 euro fully guaranteed. This measure is being implemented

Covid Economics 35, 7 July 2020: 51-69

Covid Economics 35, 7 July 2020: 51-69

Figure 4: Illiquid firms and workers without and with the Decree

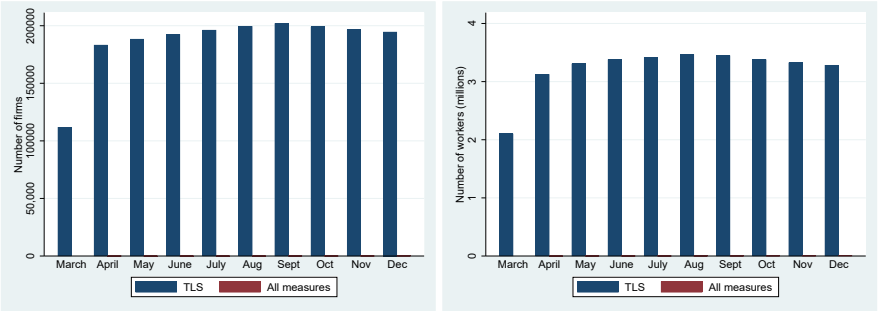

(a) Number of illiquid firms (b) Number illiquid workers

Note: The figure reports the total number of illiquid firms (Panel a) and of workers in such firms (Panel b) without and with the liquidity decree.

rather quickly, as it entails no risk for banks. As the loan amount increases, the government guarantee stops being complete. This means that banks might need some time to process applications, as well as to obtain the approval from the government agency. However, many firms become illiquid very quickly, so it is essential that credit flows to firms quickly. To check the amount of coverage from the different measures, Figure 5 reports the liquidity shortages after borrowing the maximum amount on measure 1 and after adding measures 2, 3 and 4 sequentially. We perform this exercise only for SMEs, as large firms only have access to measure 4.

Figure 5: Firms with liquidity shortfalls according to the liquidity measure

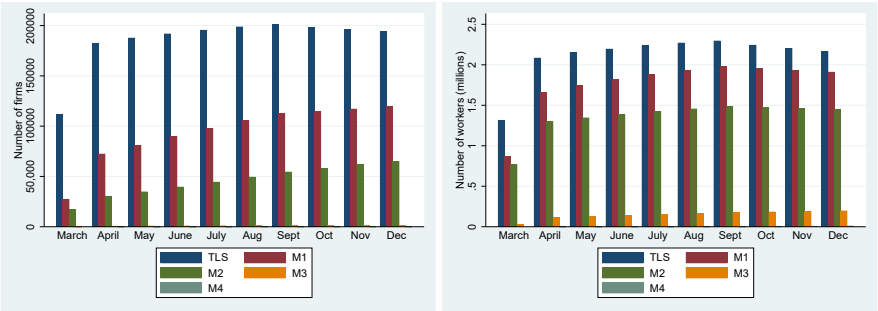

(a) Number of illiquid firms (b) Number of illiquid workers

Note: The figure reports the liquidity shortfalls with no borrowing, with borrowing from measure 1, then when adding measures 2, 3 and 4. We only consider SMEs.

At peak, more than 100,000 firms are not fully covered by measure 1, and around 65,000 by measure 2. It is only with measure 3 that we obtain almost full coverage (Panel a). Things are even more dramatic in terms of workers, as measures 1 and 2 are sufficient only for small firms. In fact, almost 1.5 million workers are in firms that cannot cover their liquidity needs with measure 2. With measure 3, which supplies up to 5 million euros, the number of uncovered firms and workers drastically drops to 1,430 firms and 197,473 workers. However, the measure only entails a 90% guarantee, so banks will have to screen borrowers. Additionally, the measure requires approval from the government. According to

our calculations, over 60,000 firms will need it, implying that the banking system and the government agency will have to process a large number of applications in a short period of time.

One way to speed up the process is to provide a two-stage procedure, using algorithms that measure credit risk in a timely manner based on scoring models in the first stage. If a company has a positive score, the credit should be given with a lean and fast investigation. Banks' specific skills in credit assessment should be committed to companies with weak scores to distinguish between companies that still have development prospects, despite their negative quantitative indicators, from those that do not. By their nature, the scores do not incorporate soft information, that is, the information that banks develop through direct relationships with their customers or that can be collected through direct investigation of the applicants. Soft information is important in cases in which hard information raises a red flag.

To assess how much this approach would reduce the preliminary investigation, we used the Cerved Credit Score. Of the 110,000 companies that will need liquidity in April, around 90,000 fall into the first seven classes, considered solvent. These firms should get credit quickly with simplified procedures. The reduction in the number of detailed investigations would allow banks to devote more time and resources to carefully but quickly screen the 20,000 companies in the risk area.

Covid Economics 35, 7 July 2020: 51-69

5 Alternative parameterizations

We now experiment with two changes in the parameterization. First, we re-run the exercise under a more pessimistic scenario on the evolution of the pandemic and the size and persistence of the drop in sales. Second, we vary the elasticities of inputs. To facilitate comparisons, we report graphs in which the alternative parametrization are compared with the basic one.

5.1 Pessimistic scenario

As we write, the Covid-19 pandemic has greatly receded since its peak in March and April. However, there is a concrete possibility of a second wave of the epidemic after the summer. Cerved sectoral experts have also produced sales growth predictions in this pessimistic case. The predictions are based on the assumption that the contagion comes back starting in September and picks up in October with a less strict lockdown due to the experience of the first phase, which has made the population more aware of how to contain the contagion. Also in this case, the impact and length of the new lockdown are differentiated by sector to take social distancing needs into account. The last column of Table 3 in the Appendix reports the sectoral growth rates at the yearly level under this pessimistic scenario. GDP would fall by 12% in 2020 and sales would decrease on average by 18%, with a more than 20% drop for 33 out of 79 sectors.

Figure 6, Panel a) reports the results for the number of firms. By construction, the number of illiquid firms is the same until August and jumps discontinuously in October to 236,000 against less than 200,000 in the basic scenario. It remains stable around that value until the end of the year. Workers in illiquid firms (unreported for brevity) jump to 4 million, with an increase of 600,000 units with respect to the basic scenario.

Panel b) of Figure 6 plots the TLS. The increase is more substantial than it is for the number of firms because in this case, not only does the extensive margin contribute to the increase (more firms become illiquid), but so does the intensive margin (illiquid firms accumulate further negative cash flows). At the peak in December, TLS reaches 106 billion, 34 billion above the basic scenario. This indicates that despite the fact that the learning

Figure 6: Illiquid firms and TLS in the base and pessimistic scenarios

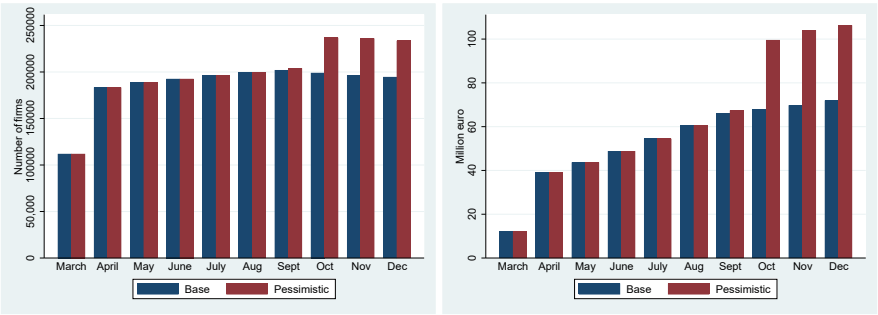

(a) Number of illiquid firms (b) TLS

Note: The figure reports the number of illiquid firms and the TLS for the basic and the pessimistic scenarios.

process of the first lockdown might mitigate the adverse consequences during the second wave, the impact on the economic system will be substantial.

5.2 Elasticities of inputs

As discussed above, the preferred values of $\varepsilon_{WS} = 0.75, \varepsilon_{MS} = 0.5$ were based on modifying the regression elasticities to take into account a policy intervention (the temporary layoff scheme that increases the labor elasticity with respect to the estimated value) and the speed of the crisis, for which the elasticity of intermediates expenditure was lowered with respect to the estimated value. We now assess how results change if we stick to the estimated values. We consider 3 scenarios: one in which the labor cost elasticity is equal to its estimated value of 0.46 while keeping the intermediate cost elasticity at 0.5, one in which the labor cost elasticity is at its basic value of 0.75 while the intermediate elasticity equals its estimated value of 0.7, and one in which both elasticities are at their estimated values of 0.46 (labor) and 0.7 (intermediates).

Figure 7, Panel a), plots the number of firms for the basic and three alternative scenarios. When we decrease the elasticity of the cost of labor from 0.75 to 0.46, more firms become illiquid, as the cost reduction following a drop in sales gets smaller. The effect is however not dramatic: at peak (September), 15,000 more firms, employing 445,000 workers, become illiquid. The increase in TLS (Panel b) at the end of the year is around 7 billion, a 10% increase with respect to the basic parametrization.

When we increase the elasticity of intermediates from 0.5 to 0.7, the changes are more substantial. The total number of illiquid firms in September drops by 40.000, which represents a 20% decrease. TLS also drops substantially from 72 to 43 billion, a 40% drop. The stronger influence of intermediates is a consequence of the fact that they represent a much larger share of costs than labor: for the median firm, the ratio of labor to intermediate costs is in fact 20%.

In the final experiment, we change both elasticities simultaneously, setting them to their estimated values of 0.46 for labor and 0.7 for intermediates. In line with the previous results, the increase in the elasticity of intermediates more that counteracts the decrease in the labor elasticity so that the number of illiquid firms and the TLS is lower than in the basic case.

Summing up, a second lockdown would have strong consequences on firms' liquidity needs. In terms of input elasticity, the results are more sensitive to the intermediates; elasticity than they are to the labor elasticity, as the former represents a larger share of

Covid Economics 35, 7 July 2020: 51-69

Figure 7: Illiquid firms and TLS under different elasticities of inputs

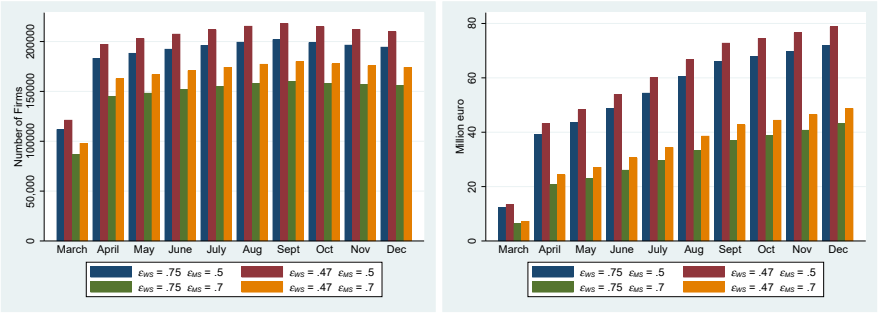

(a) Number of illiquid firms (b) TLS

Note: The figure reports the number of illiquid firms and the TLS for four different parametrizations of the elasticities of inputs.

firms' expenditures.

6 Conclusions

We have constructed an accounting scheme to predict firms' liquidity needs during the Covid-19 crisis. We have applied it to the universe of incorporated firms in Italy. We find that a substantial number of firms become illiquid early on. At the same time, the liquidity shortage is large but not unbearable: it amounts to less than 4% of Italian GDP. This is because the large drop in sales goes together with a drop in costs, which limits the effects on the cash flow. We also show that the measures enacted by the Italian government, in the form of credit guarantees to supply firms with liquidity, can cover pretty much all the liquidity needs. An important issue is then the speed of implementation: given that we find that many firms become illiquid very quickly, it is of paramount importance that governments' measures are enacted quickly. Finally, a second lockdown would have large effects on liquidity needs.

References

Bank of Italy (2020), Rapporto sulla stabilità finanziaria. Number 1/2020, April.

Carletti, E., Oliviero, T., Pagano, M., Pelizzon, L. & Subrahmanyam, M. G. (2020), The Covid-19 shock and equity shortfall: Firm-level evidence from Italy. CEPR Discussion Paper No. DP14831.

De Vito, A. & Gomez, J.-P. (2020), 'Estimating the COVID-19 cash crunch: Global evidence and policy', *Journal of Accounting and Public Policy* **39**.

European Commision (2020), Identifying Europe's recovery needs. Commission Staff working document.

Kalemli-Ozcan, S., Sorensen, B., Villegas-Sanchez, C., Volosovych, V. & Yesiltas, S. (2015), How to construct nationally representative firm level data from the Orbis global database: New facts and aggregate implications. National Bureau of Economic Research WP No. 21558.

Khan, A. & Thomas, J. K. (2007), 'Inventories and the business cycle: An equilibrium analysis of (s, s) policies', *American Economic Review* **97**(4), 1165–1188.

Covid Economics 35, 7 July 2020: 51-69

McGeever, N., McQuinn, J. & Myers, S. (2020), SME liquidity needs during the COVID-19 shock. Bank of Ireland Financial Stability Note, No. 2.

OECD (2020), Corporate sector vulnerabilities during the Covid-19 outbreak: Assessment and policy responses. OECD policy briefs on Tackling Coronavirus.

Pozzi, A. & Schivardi, F. (2016), 'Demand or productivity: What determines firm growth?', *The RAND Journal of Economics* **47**(3), 608–630.

Schivardi, F., Sette, E. & Tabellini, G. (2020), 'Identifying the real effects of zombie lending', *Review of Corporate Fiance Studies* . Forthcoming.

Covid Economics 35, 7 July 2020: 51-69

A Sectoral growth forecasts

Sectoral forecasts are produced by Cerved, which adopted a sectoral methodology to forecast the 2020/2019 growth rate of sales for the Italian firms. This was based on both quantitative evidence (i.e. sectors in lockdown, length of lockdown, % of smart working) and qualitative assessment (i.e. impact of social distancing) by sector experts. In particular, Cerved analysts considered the following factors:

- **Lockdown provisions.** The Italian government introduced lockdown provisions relating to commercial and production activities in order to contain the spread of contagion, with less restrictive measures on certain activities specifically indicated on the basis of Ateco codes ("essential activities"). A detailed breakdown of Ateco classification has been considered to take into account different periods of lockdown for different sectors.

- **Operability of firms during lockdown.** Many firms continued to operate in spite of the lockdown through smart working or e-commerce. When available, statistics on smart working and e-commerce intensity by sector have been considered.

- **Social distancing impact on demand and supply.** Social distancing impact on demand and supply is strongly dependent on the firm sector: for example, restaurants must guarantee a minimum distance between tables (impact on supply); on the other hand, Covid-19 can change the customer preferences. For example, customers may lower their demand of services if they perceive a risk of contagion (e.g. local transport services)

- **Impact of lower mobility on demand and supply.** Lockdown provisions and the necessity of social distancing have strongly reduced mobility, with effects beyond transport services (e.g. tourism services). Forecasts considered such impacts.

- **Impact on some specific sectors in terms of extra demand.** Covid-19 has raised the demand of particular goods or services. For example, demand of medical protection, plexiglass articles, and e-commerce services boomed after Covid. Analysts considered this factor in their forecasts.

- **International trade impact.** Covid-19 has strongly affected international trade with an impact on both demand and supply. The dependence on international trade in terms of intermediates and the impact of Covid-19 on destination markets have been considered.

- **Other relevant variables for specific sectors.** Other variables or law provisions have been considered if they influence the trend of specific sectors (e.g. raw materials prices, government incentives on demand of bikes).

Based on such factors, Cerved produced forecasts for more than 500 sectors, characterised by homogenous law provisions and homogenous supply and demand conditions.

For each sector, analysts produced differentiated year on year growth rate forecasts for four different periods in 2020:

- Normal: January and February 2020

- Lockdown: since March 9th, with a length depending on the specific sector;

- Transition: a period which depends on the social distancing provisions and/or impact of Covid-19 on demand and supply;

- New normal: time left to the end of 2020, given the lockdown and the transition period.

Covid Economics 35, 7 July 2020: 51-69

Covid Economics 35, 7 July 2020: 51-69

The forecast exercise has then been made coherent with a consensus macroeconomic scenario, thanks to the correlation between firm sales, firm value added and GDP. Based on the described assumptions and the bottom-up approach, sales of Italian firms would decrease by 12.9% between 2020 and 2019 (from 2,411 billions to 2,100 billions in the sample analyzed). The sectoral variability of forecasts is very high. If we consider the detailed breakdown of more than 500 sectors, the strongest loss is in Ateco 591400 (Film projection activity, -65%), while Ateco 479110 (E-commerce) is the best performer (+35%). In Table 3, we show an Ateco 2 digit aggregation of the sectoral forecasts, ranked by loss. We consider the four clusters (based on the magnitude of the sales loss in 2020) used to split firms in descriptive Table 1.

The first group is comprised by 16 Ateco 2 digit sectors with a major loss of more than 20% between 2020 and 2019. There are different sub-groups according to the main causes of the fall in turnover. First, sectors majorly impacted by lockdown, social distancing and reduced mobility: transport, travel agency activities, accommodation services, catering, film production. Second, oil extraction activities, which are affected by the lockdown by the sharp reduction in the oil price and by reduced mobility. Finally, the manufacturing of motor vehicles is expected to be affected by the sharp decline in international trade and by more cautious consumer behavior.

In the second cluster, we identify 27 sectors with a drop in turnover between 10% and 20% and four main subgroups. First, non-essential manufacturing sectors that have suffered from the impact of the lockdown on production and can be hit by difficulties in supplying components from abroad due to the disruption of global supply chains. Second, real estate and related activities which suspended activity during the lockdown. Another sector that has suffered a drop in turnover due to the lockdown is the retail trade (other than food, i.e. especially clothing), for which the social distancing measure are expected to make recovery slower. Finally, the last group of sectors includes recreational activities for which supply will continue to be reduced in order to comply with social distancing measures.

The third cluster includes 27 sectors with a drop of sales between 10% and 0%. These sectors include activities that have been classified as essential but were nevertheless impacted by Covid-19 due to the decline in demand of downstream sectors: suppliers of electricity and gas, professional and technical services, accounting, wholesale trade. Health services are also part of this group because despite being at the forefront of addressing the health emergency, many activities not related to the emergency were suspended. In addition, another group of sectors is penalized by the collapse of non-home consumption (e.g. drinks and fisheries/aquaculture).

The last cluster includes sectors that have experienced a stable growth in turnover in spite of Covid-19, such as utilities less linked to industrial production (water supply) and food-related sectors. Finally, the best performers are the pharmaceutical sector, which has received a strong boost to directly address the health emergency, and e-commerce, which has benefited from a surge in demand from consumers.

Covid Economics 35, 7 July 2020: 51-69

Table 3: Annual sales growth by sector, base and pessimistic scenarios

Nace code	Sector Description	Base	Pess
51	Air transport	-46.0	-55.4
79	Travel agency, tour operator and related activities	-43.8	-55.0
55	Accommodation	-41.6	-51.2
56	Food and beverage service activities	-30.1	-44.9
59	Motion picture, video television, sound recording	-29.4	-43.4
6	Extraction of crude petroleum and natural gas	-27.8	-36.3
9	Mining support service activities	-27.7	-36.3
29	Manufacture of motor vehicles. trailers and semi-trailers	-25.3	-35.2
49	Land transport and transport via pipelines	-24.1	-33.1
50	Water transport	-23.2	-33.8
52	Warehousing and support activities for transportation	-22.9	-31.7
45	Trade and repair of motor vehicles and motorcycles	-21.9	-28.9
14	Manufacture of wearing apparel	-21.3	-29.6
91	Libraries, archives, museums and other cultural activities	-20.6	-25.3
93	Sports activities and amusement and recreation activities	-20.6	-25.3
24	Manufacture of basic metals	-20.4	-28.5
92	Gambling and betting activities	-19.8	-26.8
90	Creative arts and entertainment activities	-19.7	-24.5
41	Construction of buildings	-19.2	-27.0
8	Other mining and quarrying	-17.7	-25.3
82	Office administrative and other business support activities	-17.7	-23.9
28	Manufacture of machinery and equipment n.e.c.	-17.5	-24.5
31	Manufacture of furniture	-16.9	-25.7
15	Manufacture of leather and related products	-16.9	-25.7
23	Manufacture of other non-metallic mineral products	-16.4	-24.1
25	Manufacture of fabricated metal products†	-16.2	-22.9
19	Manufacture of coke and refined petroleum products	-16.0	-25.1
47	Retail trade†	-15.6	-21.5
43	Specialised construction activities	-15.4	-21.8
18	Printing and reproduction of recorded media	-15.4	-21.3
71	Architectural and engineering activities	-14.8	-19.5
30	Manufacture of other transport equipment	-14.5	-22.6
13	Manufacture of textiles	-14.3	-21.3
77	Rental and leasing activities	-14.2	-20.3
58	Publishing activities	-13.4	-17.7
16	Manufacture of wood and wood products, except furniture	-13.4	-18.6
22	Manufacture of rubber and plastic products	-13.3	-19.3
42	Civil engineering	-11.8	-17.2
26	Manufacture of computer, electronic and optical products	-11.7	-16.9
72	Scientific research and development	-11.6	-17.8
27	Manufacture of electrical equipment	-10.9	-17.4
53	Postal and courier activities	-10.5	-7.4
73	Advertising and market research	-10.5	-16.0
35	Electricity, gas, steam and air conditioning supply	-9.4	-12.1
74	Other professional, scientific and technical activities	-9.1	-12.8
86	Human health activities	-8.6	-12.3
78	Employment activities	-8.5	-10.7
69	Legal and accounting activities	-8.2	-12.4

46	Wholesale trade, except of motor vehicles and motorcycles	-8.0	-12.0
85	Education	-7.9	-15.2
62	Computer programming, consultancy and related activities	-7.9	-9.2
80	Security and investigation activities	-7.8	-11.2
96	Other personal service activities	-7.6	-9.6
12	Manufacture of tobacco products	-7.6	-12.8
3	Fishing and aquaculture	-6.7	-10.6
11	Manufacture of beverages	-6.4	-10.5
33	Repair and installation of machinery and equipment	-6.2	-8.3
32	Other manufacturing	-6.1	-10.4
20	Manufacture of chemicals and chemical products	-5.9	-8.6
63	Information service activities	-5.3	-6.2
70	Activities of head offices; management consultancy activities	-5.0	-6.1
60	Programming and broadcasting activities	-5.0	-8.0
88	Social work activities without accommodation	-4.9	-6.5
5	Mining of coal and lignite	-4.5	-7.8
81	Services to buildings and landscape activities	-4.0	-6.3
17	Manufacture of paper and paper products	-3.6	-5.4
61	Telecommunications	-2.5	-3.4
1	Crop and animal production and related service activities	-2.3	-5.0
38	Waste collection, treatment and disposal activities	-1.5	-2.5
39	Remediation activities and other waste management services	-1.5	-2.5
66	Activities auxiliary finance and insurance	0.0	0.0
68	Real estate activities	0.0	-1.7
10	Manufacture of food products	0.0	-0.2
87	Residential care activities	0.4	-1.8
36	Water collection, treatment and supply	0.6	0.6
37	Sewerage	0.6	0.6
47.1+47.2	Other in-store retail trade‡	9.9	10.7
21	Manufacture of basic pharmaceutical products	10.2	13.0
4791	Retail sale via mail order houses or via Internet	30.2	40.0

§ Except machinery and equipment
† Except for motor vehicles and motorcycles. Food, beverages, tobacco, via mail order or via internet
‡ Retail sale of food, beverages and tobacco in specialized stores and in non-specialized stores

Note: The table reports sectoral growth rates on 2020 with respect to 2019 under the basic and pessimistic scenarios.

Covid Economics 35, 7 July 2020: 51-69

Covid Economics Issue 35, 7 July 2020

Quantifying the impact of non-pharmaceutical interventions during the COVID-19 outbreak: The case of Sweden[1]

Sang-Wook (Stanley) Cho[2]

Date submitted: 29 June 2020; Date accepted: 1 July 2020

This paper estimates the effect of non-pharmaceutical intervention (NPI) policies on public health during the recent COVID-19 outbreak by considering a counterfactual case for Sweden. Using a synthetic control approach, I find that strict initial lockdown measures played an important role in limiting the spread of the COVID-19 infection and that Swedish policymakers would have eventually reduced the infection cases by more than half had they followed those policies. As people dynamically adjust their behavior in response to information and policies, the impact of NPIs becomes visible with a time lag of around 5 weeks. An alternative difference-in-differences research design that allows for changes in behavioral patterns also confirms the effectiveness of a strict lockdown policy. Finally, extending the analysis to excess mortality, I find that the lockdown measures would have lowered excess mortality in Sweden by 23 percentage points, with a steep age gradient of more than 30 percentage points for the most vulnerable elderly cohort. The outcome of this study can help policymakers lay out future policies to further protect public health, as well as facilitate an economic plan for recovery.

Covid Economics 35, 7 July 2020: 70-95

1 I thank the editor Charles Wyplosz as well as Nicola Aravecchia, Hansoo Choi and Julián P. Díaz for their constructive feedback and comments.
2 Senior Lecturer, University of New South Wales.

Covid Economics 35, 7 July 2020: 70-95

1. Introduction

Coronavirus disease 2019 (COVID-19) is a viral respiratory illness caused by a new coronavirus, first reported in Wuhan, Hubei Province, China in November 2019. Over the next few months, the illness rapidly spread to almost every country. In response, the WHO declared COVID-19 a pandemic on March 11, 2020. As vaccines or medicines for COVID-19 have yet to be available, most countries around the world resorted to non-pharmaceutical interventions (NPIs), or community mitigation strategies, to help slow the spread of the illness. Some of the NPIs involve government measures to close schools and workplaces, canceling and restricting public events and gatherings, shutting down public transport and stay-at-home requirements, as well as restrictions on domestic and international travel, not to mention general public information campaigns. By late March, nearly every country in Europe have implemented these policies basically putting themselves into a nationwide lockdown. These government policies remained in place until late May with a gradual easing of some of the harshest measures. One country, however, stood out for its decision to remain open: Sweden. In fact, Swedish officials chose not to implement a nationwide lockdown, trusting that people would voluntarily do their part to stay safe. For example, while high schools and universities have switched to distance learning, elementary and preschools have remained open. In addition, while the government recommended people to stay at home, many non-essential businesses such as restaurants, gyms and bars were still open, while gatherings up to 50 people were allowed. Given this divergence in the policy measures between Sweden and the rest of Europe, I study the public health impact of NPIs by asking how the trajectories of the COVID-19 infection and mortality would have evolved had Sweden opted for more stringent lockdown measures.

In order to study this counterfactual scenario, I first employ the synthetic control method (SCM) pioneered by Abadie and Gardeazabal (2003) and analyze how a parallel (or "synthetic") version of Sweden would have evolved had it enforced a mandatory lockdown policy. This parallel version of Sweden is first constructed through a data-driven process with weights assigned to all possible donor countries that would best approximate the pre-lockdown characteristics of Sweden (our "treatment" unit). Once the policy intervention takes place, we can trace its effect with the evolution of the untreated synthetic control unit to assess the counterfactual situation corresponding to the parallel regime where strict lockdown measures were in place. The causal effect of the lockdown is measured by the post-intervention difference in infection rates of the treatment and the synthetic control unit. It has been shown that the synthetic control method offers several advantages over traditional difference-in-differences or fixed-effect models as not only is the procedure a transparent data-driven one but also it allows the effect of unobservable country heterogeneity to vary over time as discussed by

Covid Economics 35, 7 July 2020: 70-95

Abadie, Diamond and Hainmueller (2010) and Imbens and Wooldridge (2009). I further quantify the causal effect of counter-COVID measures by using a difference-in-differences (DD) research design that allows for additional variables regarding people's behavior. This would enable us to understand how much of the observed infection rate dynamics is attributed to the effect of NPIs by itself relative to voluntary changes in people's behavior for fear of infection.

The key findings from regression analysis are as follows. I find that the lockdown measures played an important role in limiting the spread of the COVID-19 infection and that Swedish policymakers would have contained the infection cases by more than half had they followed similar policies implemented elsewhere. I also find that as people dynamically adjust their behavior in response to information and policies, the impact of NPIs does not manifest immediately but only with a time lag of approximately five weeks. Profiling excess mortality for the synthetic Sweden, I find that the excess mortality rate in Sweden would have been reduced by approximately 23 percentage points had the policymakers followed strict counter-COVID measures. The effectiveness in death prevention becomes disproportionately higher by age, with more than a 30 percentage point reduction in the excess mortality rate for the elderly cohort aged 85 and above.

This paper contributes to the ongoing discussion on the effectiveness of NPI policy response to the COVID-19 shock, see Chen and Qiu (2020); Gonzalez-Eiras and Niepelt (2020); Ullah and Ajala (2020); Goodman-Bacon and Marcus (2020); Chernozhukov, Kasahara and Schrimpf (2020) and the contributions in the volume by Baldwin and di Mauro (2020). Empirically, this paper extends cross-country experiences in the policy effectiveness. Castex, Dechter and Lorca (2020) shows that the effectiveness of NPIs differ by various socio-economic and public health systems, and the effectiveness of lockdown policies is declining with GDP per capita, population density and surface area; and increasing with health expenditure and proportion of physicians in the population. In terms of scope and methodology, the paper is closest in spirit to Born, Dietrich and Müller (2020), hereinafter BDM, that conducts a similar counterfactual lockdown scenario for Sweden using the synthetic control method. Documenting infection dynamics of one month post-lockdown, they find that the counterfactual Sweden did not differ from actual infection dynamics observed in Sweden. In their discussion, they attribute this outcome to the voluntary precautions taken by the general public that essentially had the same impact as a mandatory lockdown.

This paper extends BDM in the following aspects. First, I consider post-lockdown period extending for two months, which completely covers the time horizon during which the initial lockdown measures were fully in place outside Sweden. Consistent with BDM, I also find that during the first half, the infection dynamics in the counterfactual Sweden was not

lower than those in actual Sweden. However, over time, the counterfactual Sweden shows a significant slowdown in the infection rate, which demonstrates that the lockdown measures would eventually have a containment effect in the longer horizon. Second, using a difference-in-differences approach, I formally control for the behavioral changes using Google Mobility Tracker and show that the mandatory lockdown measures would have significantly reduced the infection rate in comparison to a voluntary social distancing scenario.

The rest of the paper is organized as follows. Section 2 describes the methodology and data for the synthetic control approach. Section 3 presents the main estimation results and robustness checks. Section 4 extends the analysis to mortality and discusses the role of voluntary social distancing, followed by a difference-in-differences estimation in Section 5. Finally, conclusion is provided in Section 6.

Covid Economics 35, 7 July 2020: 70-95

2. Data and Methodology

In this section I describe the synthetic control method (SCM) proposed by Abadie and Gardeazabal (2003), and later developed in Abadie, Diamond and Hainmueller (2010) and Abadie, Diamond and Hainmueller (2015). The SCM is a popular approach for comparative case studies, which has also been used to quantify the economic effects of shocks or policy interventions. [1]

Under the synthetic control approach, we can generate a counterfactual designed to capture how the infection rates would have evolved in Sweden had it followed a similar policy approach (or a mandatory lockdown) taken by other European countries. This counterfactual (or synthetic control) unit would track the actual path of infection rates in Sweden (our treatment unit) as closely as possible prior to the policy intervention. After the policy intervention, the control unit followed a path of mandatory lockdown measures while Sweden did not. As such, the notion of policy intervention in our setting refers to the absence of mandatory lockdown measures, or no changes in government policy. Due to difficulties in picking individual countries that satisfy these criteria, we resort to a weighted average of potentially comparable countries that best resemble the characteristics of Sweden prior to the policy intervention. Any discrepancy in the infection dynamics between the two units after the policy intervention can be interpreted as an outcome of the policy or the treatment effect.

As the SCM exploits the pre-intervention data to form better counterfactual values, it is often preferred over other program evaluation methods such as difference-in-differences in comparative case studies.

[1] See Abadie (2020) for a broader overview of the methodology.

Covid Economics 35, 7 July 2020: 70-95

2.1. Data

The outcome variable of interest is the infection dynamics as measured by cumulative infection per million population. As for potential donor pool, I select 29 countries consisting of the European Union members (excluding Malta due to lack of data) as well as Iceland, Israel, Norway and Switzerland. As the infection dynamics varied across countries, we normalize the time unit such that "Day 1" refers to the day on which the infection per million exceeds one. For country-specific characteristics, epidemiological studies suggest demographic factors such as population size and the rate of urbanization to be crucial to understand the infection dynamics. I also include population density as it has been found to catalyze the spread of COVID-19 by Rocklöv and Sjödin (2020). The latest available figures for all three country-specific covariates were taken from the World Development Indicators (WDI).

Table 1: COVID-19 and demographic characteristics

Variables	Sweden	All donors (n=29)
COVID-19 dynamics		
- Day 1	29 February	4 March
- Case per million on Day 1	1.18	1.51
- Lockdown day		28 March
- Pre-lockdown duration (days)		23.9
- Case per million on Lockdown day	199.6*	471.4
- Stringency Index (SI) on Lockdown day	32.4*	82.4
Demographics		
- Population (million)	10.1	18.5
- Urban population fraction (%)	87.4	75.1
- Population density	24.7	146.7

Note: Day 1 refers to the date on which the infection per million exceeds one. Lockdown day refers to the date on which the SI index reached the maximum. For reference, * denote the numbers for Sweden on 24 March, 24 days since Day 1.

Next, for the policy intervention, lockdown measures consist of various socioeconomic measures including school and workplace closing, cancellation of public events, restrictions on gatherings, closing of public transport, stay at home requirements, restrictions on domestic/international travels, as well as public info campaigns. As these measures took place over different time with varying magnitudes, we resort to an all-inclusive index measure. The OxCGRT data[2] provides a Government Response Stringency Index (Stringency Index, SI), which ranges from 0 to 100 with each additional government response leading to a higher

[2]https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker.

index value.[3] The Stringency Index varies across jurisdictions and time on a daily basis. I pick the date at which this stringency index reached its peak in each country to pinpoint the timing of our policy intervention. Table 1 summarizes the COVID-19 dynamics as well as country specific characteristics for Sweden and the simple average of all 29 donors.

It's worth noting that the COVID-19 infection started a few days earlier in Sweden compared to the average of all donor countries. For the latter group, it took around 24 days for a full lockdown measure was in place when the Stringency Index reached 82.4. For Sweden, on the other hand, the index remained at 32.4 around the same time and never reached higher than 55 during the whole period of our analysis.[4] For the demographic covariates shown in the last three rows, we note that Sweden is characterized by a smaller population with one of the highest urbanization rate and a significantly low population density than the average of the donor group.

Next, we proceed to find the weighted average of the countries in the donor pool which will generate the synthetic control unit for Sweden. The weights are assigned by minimizing the distance between Sweden and the synthetic control unit along all three demographic covariates as well as the average infection rates in the first 20 days since Day 1. Including lagged terms of the dependent variable often helps mitigate the problem of omitting important predictor effects as suggested by Athey and Imbens (2006).

3. Results

Table 2 summarizes the predictor variables for the synthetic Sweden, which is constructed as a weighted average of Finland, France and Norway, with the largest weight assigned to Finland and followed by Norway.[5] Compared to the simple average of all countries in the donor pool (as shown in Table 1), the synthetic control unit provides a much better matched profile of Sweden along the predictors. In other words, the weighted selection of countries seem more appropriate as a control unit than taking a simple average of all countries in the donor pool. In the synthetic control approach, the root mean square prediction error (RMSPE) measures the gap between the variable of interest for the treated country and its synthetic counterpart. The last row of Table 2 reports the RMSPE for the pre-intervention period.

[3]Hale, Webster, Petherick, Phillips and Kira (2020) provides detailed information on the construction of the stringency index.

[4]No single country in the donor pool had the Stringency Index peaking below this maximum value for Sweden.

[5]In the Appendix, Table A1 breaks down the demographic and pre-intervention epidemiological profile while Figure A1 shows the dynamics of the SI index for each country comprising the synthetic unit.

Covid Economics 35, 7 July 2020: 70-95

Table 2: Predictor variables and RMSPE

	Sweden	Synthetic Sweden
		Finland (0.643), France (0.076), Norway (0.281)
Predictors		
- Population (million)	10.099	10.047
- Urban population fraction (%)	87.431	84.126
- Population density	24.718	25.041
- Case per million (first 20 days)	42.830	42.843
RMSPE		8.605

Note: For countries with positive weights, the weights are shown in brackets. All other countries in the donor pool receive zero weight.

Covid Economics 35, 7 July 2020: 70-95

Next, in the upper panel of Figure 1, I show the profile of infection dynamics for the synthetic Sweden together with the actual Sweden. I consider a period of 75 days which roughly corresponds to the entire months of March and April as well as the first half of May. As many countries started to gradually ease some of the lockdown measures in late May, the period under observation covers the full period of the initial lockdown. The policy intervention takes places on Day 24—as indicated by the dashed vertical line—which falls on the midpoint of lockdown dates of the three countries comprising the synthetic control unit.

For the first two weeks upon the policy intervention, the cumulative infection cases in the synthetic Sweden follow the actual Sweden quite closely or even higher than the latter.[6] After this period of incubation, there is a divergence in which the actual Sweden follows a much steeper path than its synthetic counterpart. By the end of our sample period on Day 75[7], or roughly 7 weeks after the lockdown intervention, the infection case in Sweden reaches around 2,700. On the other hand, the figures for the synthetic Sweden reaches slightly below 1,300. In other words, Swedish policymakers would have reduced the infection cases by more than half had they followed similar policies implemented elsewhere, which signifies the important role of the lockdown measures in limiting the spread of the COVID-19 infection.

The lower panel of Figure 1 generates a 95% confidence interval for the gap between the two profiles using a methodology proposed by Firpo and Possebom (2018). The gap becomes statistically significant approximately five weeks after the implementation of the lockdown measures. On one hand, this result is consistent with that of Born, Dietrich and Müller (2020), which looks at the first five weeks of the lockdown measures and concludes that the

[6]As an alternative, I convert the outcome variable into logs and show the profile in Figure A2 in the Appendix.

[7]This day corresponds to 13 May.

Covid Economics 35, 7 July 2020: 70-95

Figure 1: Profile of Infection Rates – Sweden vs. Synthetic Sweden

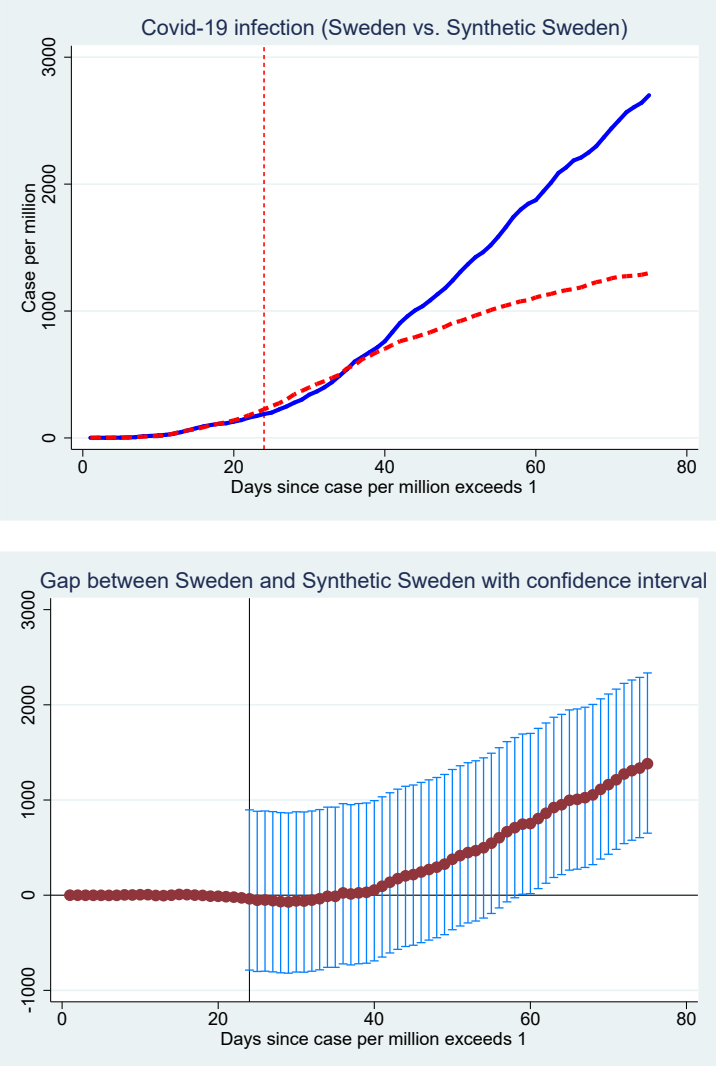

Note: Top panel shows infection case per million population for Sweden (in blue) versus synthetic Sweden (in red dash). Bottom panel shows the gap between the two units with 95% confidence intervals. Vertical line indicates the date of policy intervention.

mandatory lockdown would not have made significant differences in the infection rate in Sweden.[8] However, expanding the horizon over the entire lockdown period, I show that the epidemiological impact of lockdown measures takes places with a time lag and eventually becomes more visible in the longer horizon.

3.1. Robustness Tests and Inference

To evaluate the credibility of the baseline results, I conduct placebo (or falsification) tests based on permutation techniques, as suggested in Abadie, Diamond and Hainmueller (2010). One way the design of the study may influence the outcome comes from the choice of countries in the donor pool with positive weights assigned. If dropping one country from the donor pool creates a large effect on the results without a discernible change in pre-intervention fit, this may require a reexamination if the change in the magnitude of the estimate is caused by the effects of other interventions or by particularly large idiosyncratic shocks on the outcome of the excluded country. As such, I perform a leave-one-out analysis, where I exclude from the sample one-at-a-time each of the three countries that contributes to the synthetic control in the benchmark. For each case, the new list of donors with positive weights as well as the values of the predictors are shown in Table 3. In the benchmark, Finland was the country assigned with the largest weight followed by Norway. Dropping Finland from the donor pool generates a new set of donors consisting of Bulgaria, Croatia and Norway. On the other hand, dropping out France or Norway from the donor list produced exactly the same new donor list consisting of Belgium, Finland and Iceland.

Figure 2 shows the results of a leave-one-out re-analysis. The resulting estimates for the days after the policy intervention (in dashes) are all positive and centered around the result produced under the benchmark. The main conclusion of a positive estimate of the infection rates in Sweden over its counterfactual scenario of a mandatory lockdown is robust to the exclusion of any particular country from the donor list.

Next, I run a cross-sectional placebo test (or "placebo in-space") by sequentially applying the synthetic control algorithm to each country in the pool of potential controls, which generates a distribution of placebo estimates across 29 donors. We then can compare the benchmark estimates of the truly treated economy with this distribution. The cross-sectional placebo tests are shown in Figure 3. The gray lines show the gap in the infection rates between each country in the donor pool and its respective synthetic version. The thick red line depicts the baseline results obtained for Sweden. Visual inspection shows that Sweden

[8]The construction of the synthetic control unit in Born, Dietrich and Müller (2020) differs from mine as they do not include population density as predictors. As such, their synthetic control unit has a population density that is almost ten times larger than that of Sweden.

Covid Economics 35, 7 July 2020: 70-95

Covid Economics 35, 7 July 2020: 70-95

Table 3: Leave-one-out robustness check

	Sweden	Synthetic Sweden		
		(Benchmark)	(No Finland)	(No France/Norway)
Donors with positive weight		FIN(0.643), FRA(0.076), NOR(0.281)	BGR(0.525)*, HRV(0.088)*, NOR(0.386)	BEL(0.021)*, FIN(0.908), ISL(0.071)*
Predictors				
- Population (million)	10.099	10.047	6.102	5.299
- Urban population fraction (%)	87.431	84.126	76.138	86.246
- Population density	24.718	25.041	46.290	24.596
- Case per million (first 20 day)	42.830	42.843	43.851	42.884
RMSPE		8.605	4.304	14.117

Note: * denotes newly added countries with positive weights (shown in brackets) from each robustness check. Full list of countries in abbreviation are as follows: BEL (Belgium), BGR (Bulgaria), FIN (Finland), FRA (France), HRV (Croatia), ISL (Iceland), NOR (Norway).

Figure 2: Leave-one-out Robustness Test

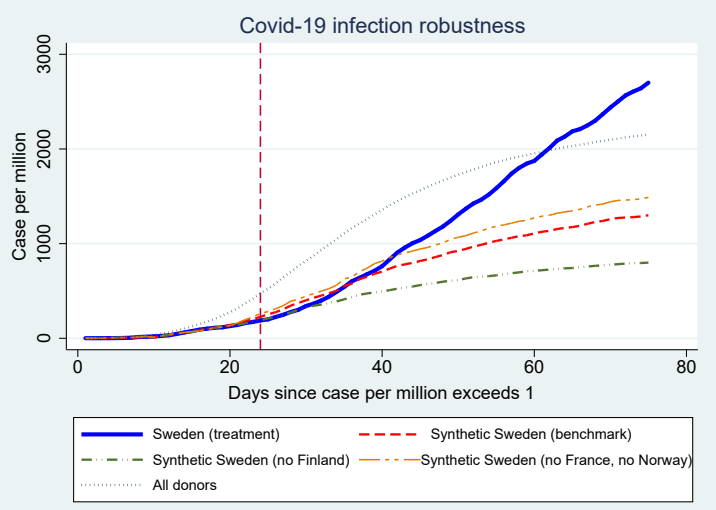

Note: New synthetic units from leave-one-out robustness check are plotted in dashes. For reference, the average profile of all donors is plotted in blue dots. Red vertical dashed line indicates the date of policy intervention.

CEPR PRESS

joins in the list of countries with positive treatment effect, but not necessarily at the right tail of the distribution of treatment effects. However, towards the end of the sample period, the treatment effect for Sweden is distinctly higher than most other countries.

Figure 3: Placebo-in-space Robustness Test

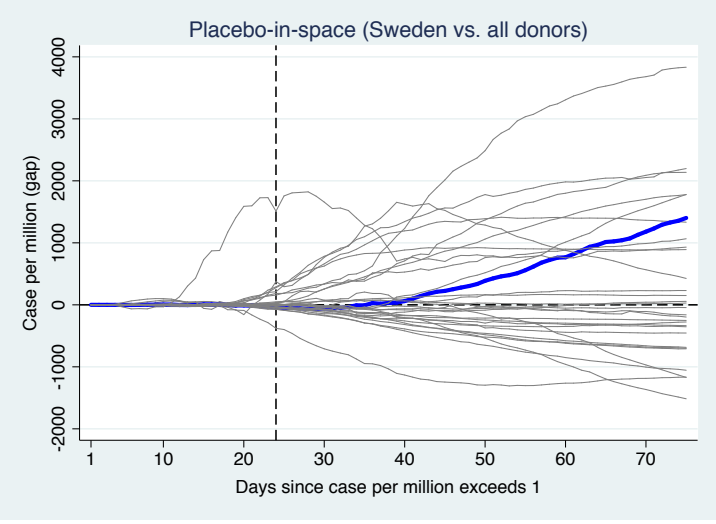

Note: Gap between the treated and synthetic control unit is plotted for Sweden (in thick blue) and each of the donors (in gray). Vertical dashed line indicates the date of policy intervention.

While the previous figure offers a visual evidence of the treatment effects over time, it does not provide a numerical measurement that quantifies the overall significance of the results. To overcome this issue, I follow Abadie, Diamond and Hainmueller (2010), who offer an alternative approach for an inference test by constructing exact p-values based on Fisher (1935). As the root mean square prediction error (RMSPE) measures the gap between the variable of interest for the treated country and its synthetic counterpart, we can calculate a set of RMSPE values for the pre- and post-treatment period for Sweden as well as each country in the cross-sectional placebo test. Countries with negative treatment effect are assigned with a minus sign to their post-treatment RMSPE value. I then compute the country-specific ratio of the post- to pre-treatment RMSPE to quantify the post-treatment divergence in the infection rate, relative to the estimated gap pre-treatment. The distribution of this RMSPE ratio (from highest to lowest) is shown in Figure 4. For Sweden, the RMSPE ratio of around 80 is far higher than those obtained for other countries in the control group. The ranking, converted into fractions, provide the basis for a p-value for Sweden, which measures the probability of observing a ratio as high as the one obtained for Sweden if one

Covid Economics 35, 7 July 2020: 70-95

were to pick a country at random from the list of potential controls. In our case, an exact p-value for Sweden is 0.1 as Sweden ranks third out of 30 countries, which falls within the conventional range of statistical significance.

Figure 4: RMSPE ratio

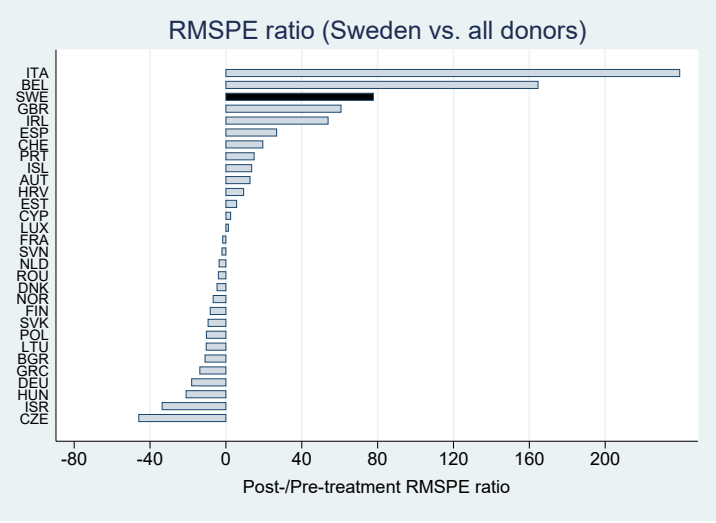

Note: Countries where the post-treatment infection cases consistently fall below those of its synthetic unit are shown with a minus sign.

4. Discussion

4.1. Infection to Mortality

So far, the focus of the analysis has been the rate of infection. One caveat of our analysis using infection cases to assess the impact on the spread of COVID-19 is that Sweden conducted very little testing compared to other countries. As the infection cases depend on the number of testing, this most likely underestimates the true treatment effect. While this issue is hard to resolve, one could take a look at the rate of mortality from COVID-19, and compare how the NPIs impacted the rate of death during the COVID-19 crisis.

While national health protection agencies report daily death counts, some jurisdictions include both confirmed and probable cases and deaths while others only report confirmed cases. As such, daily reported figures for deaths are difficult to compare across countries. Instead, I use excess mortality rate—the ratio of numbers of deaths over and above the historical average between 2015 and 2019—as a more reliable source of information for comparison. The Short-Term Mortality Fluctuation data series (STMF) from Human Mortality

Covid Economics 35, 7 July 2020: 70-95

CEPR PRESS

Covid Economics 35, 7 July 2020: 70-95

Database[9] offers weekly death counts by age groups and sex for 22 countries including Sweden as well as the countries assigned with positive weights in the construction of synthetic Sweden in Section 3: Finland, France and Norway. This allows me to generate weekly excess mortality for the synthetic Sweden and compare that with the profile of actual Sweden, which is shown in Figure 5. For reference, the top panel shows cumulative infection case per million population on a weekly basis, and Week 13 (22-28 March) is the week in which the policy responses began to diverge in the two groups. The bottom panel shows excess death rates for Sweden and its synthetic unit. Prior to Week 13, there is no visible difference in the excess mortality rates between the two groups. Leading to Week 13, however, the excess mortality rate rises much steeper in Sweden and remains consistently higher than its synthetic counterpart. At its peak, the mortality rate in Sweden is more than 40% above its historic average, while the corresponding peak for the synthetic unit is around 15%. On average, as summarized in Table 4, the excess death rate over the 10 weeks post-intervention period in Sweden is 28.5 percent higher than its historic average. In contrast, the corresponding rate in its synthetic version is 5.4 percent higher than the historic average. In other words, the excess death rate would have been more than 23 percentage points lower had the Swedish policymakers follow similar policies adopted by its parallel counterpart.

As the database provides mortality information by age, I apply the same analysis across different age groups as shown in Figure 6. A visual inspection shows that the gap in excess mortality after the lockdown becomes significantly more pronounced for older age cohorts. As summarized in Table 4, the average post-lockdown gap in excess mortality grows from around 13 percent among working age cohorts to more than 30 percent for the elderly cohort aged 85 and above.

Table 4: Excess mortality pre-lockdown vs. post-lockdown

	Pre-lockdown (Week 1–12)			Post-lockdown (Week 13–22)		
	Sweden	Synthetic Sweden	**Gap**	Sweden	Synthetic Sweden	**Gap**
Total population	0.916	0.953	-0.037	1.285	1.054	0.230
Age 15-64	0.894	0.926	-0.032	1.083	0.951	0.131
Age 65-74	0.919	0.962	-0.043	1.188	1.037	0.150
Age 75-84	0.866	0.872	-0.006	1.246	0.997	0.249
Age 85 plus	0.908	0.916	-0.008	1.330	1.029	0.301

Note: Columns labeled "Gap" measure the difference in excess mortality rates between Sweden and synthetic Sweden during each sub-period. Pre-lockdown period includes the first 12 weeks of 2020 until 21 March, while post-lockdown period includes the latter 10 weeks from 22 March until 31 May.

[9]https://www.mortality.org.

Covid Economics 35, 7 July 2020: 70-95

Figure 5: Infection to mortality

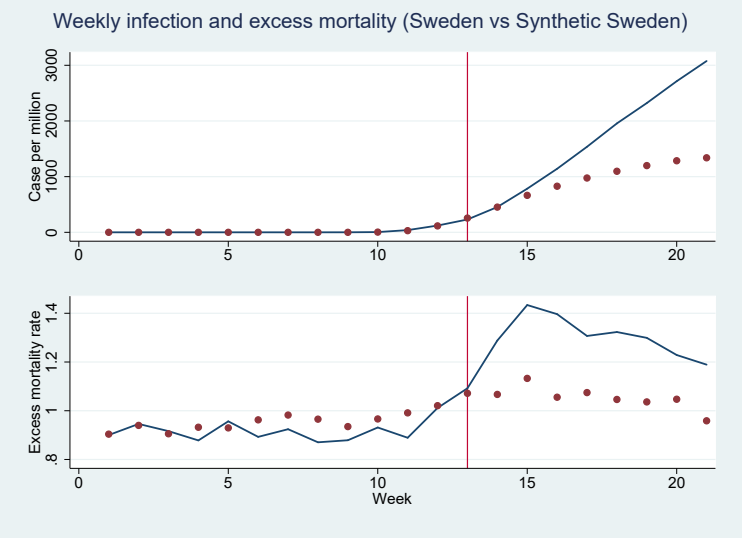

Note: Horizontal axis shows calendar weeks of 2020. Red vertical line denotes the week of policy intervention in Week 13 (22-28 March). Sweden is shown in blue line, while the synthetic Sweden is shown in red dots.

Figure 6: Excess mortality by age

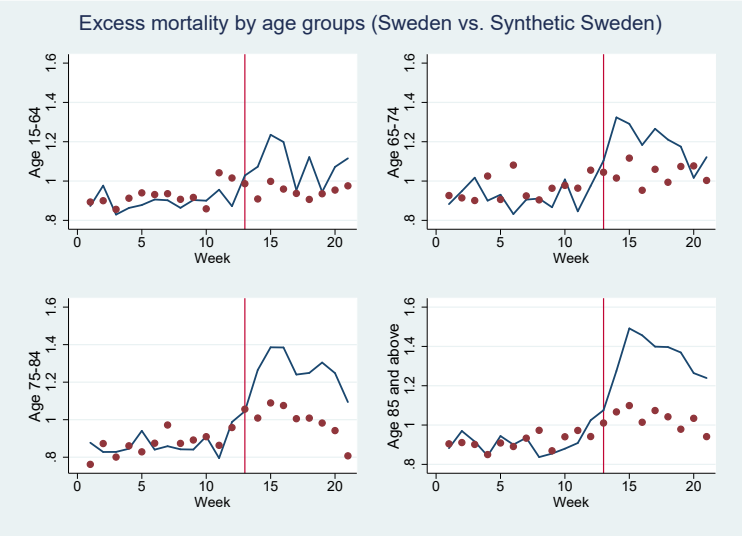

Note: Horizontal axis shows calendar weeks of 2020. Red vertical line denotes the week of policy intervention in Week 13 (22-28 March). Sweden is shown in blue line, while the synthetic Sweden is shown in red dots.

Covid Economics 35, 7 July 2020: 70-95

4.2. Voluntary Social Distancing or Involuntary Lockdown?

Naturally, infection dynamics is not only dependent on the lockdown measures. In fact, there were signs that people were already taking precautionary actions prior to various lockdown measures. For example, even before lockdown measures were announced, people made more trips to grocery stores and pharmacies to stock up on basic necessity items such as toilet papers and disinfectants. On the other hand, while the government allowed many businesses to open, most people in Sweden stayed home or followed social distancing protocols. Born, Dietrich and Müller (2020) speculate that the voluntary social distancing essentially had the same impact as a mandatory lockdown. In fact, using its location services, Google provides mobility trends by geography across different categories of places such as retail and recreation, groceries and pharmacies, parks, transit stations, workplaces, and residential areas.[10] The mobility trend for Sweden and the synthetic Sweden is shown in Figure 7 where the baseline—shown as zero in the vertical axis—is the median value, for the corresponding day of the week, during the 5-week period between 3 January and 6 February, 2020.

Figure 7: Google mobility report

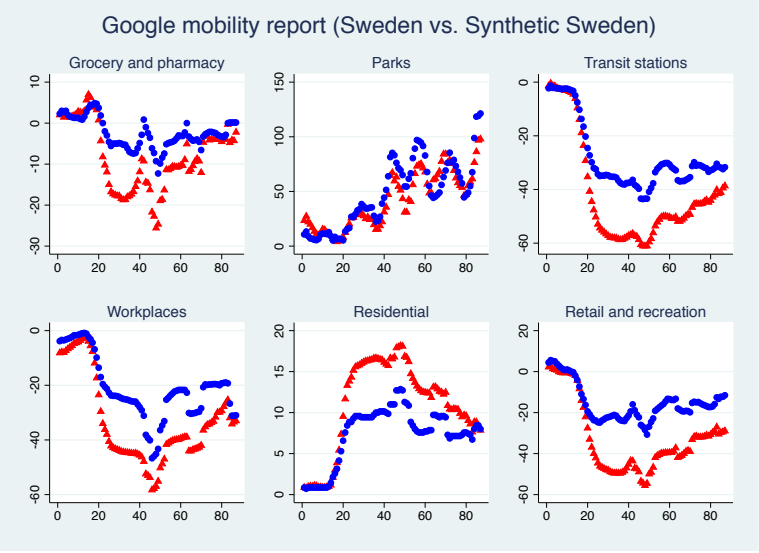

Note: Horizontal axis measures days since 29 February. 7-day moving average for Sweden and synthetic Sweden are shown in blue and red dots, respectively.

A visual inspection of 7-day moving average of the mobility for Sweden and synthetic

[10]https://www.google.com/covid19/mobility/ for direct access of the reports.

Sweden shows that before the lockdown, there is a sudden spike in trips to grocery and pharmacy in both units before they fall dramatically. Other categories negatively impacted by the Covid-19 outbreak such as visits to transit stations, workplaces and retail and recreation facilities, there is no anticipatory effects prior to the big drop. For all these categories, post-lockdown period fall is more prominent in the synthetic Sweden than in Sweden, which reflects different behavioral patterns due to voluntary and involuntary lockdown measures. An opposite pattern is shown for home stays. Finally, visits to parks show dramatic rise reflecting the lockdown (but possibly warmer weather) with no visible difference between the two units.

Given that the mobility fell in Sweden as a result of voluntary precaution, one could impute this to the delay in the divergence of infection rate in Sweden relative to its synthetic cohort. However, in the longer horizon, infection rates diverged significantly despite voluntary social distancing. In order to better control for this behavioral change, I now turn to a difference-in-differences approach.

Covid Economics 35, 7 July 2020: 70-95

5. Empirical Approach—Difference-in-differences (DD)

Taking the difference-in-differences approach pioneered by Card and Krueger (1994) to quantify an unbiased estimate of the effects of lockdown measures, I run the following two-way fixed effects specification:

$$Y_{it} = \beta_0 + b_1 L_{it} + b'\mathbf{X} + \tau_t + m_i + e_{it} \quad (1)$$

where Y_{it} denotes infection case per million in country i in time t and L_{it} is an indicator dummy for lockdown status that takes a value of 1 if there is a switch to a Swedish-style recommendation and 0 otherwise. $\mathbf{X}$ captures other additional controls; τ_t is a time fixed effect dummy; m_i controls for country fixed effects, and e_{it} is the error term. Our focus is on the coefficient b_1, which essentially captures the effect of the no-lockdown policy on the infection rate dynamics.

Our treatment country is Sweden and the control group consists of Finland, France and Norway, the three countries that collectively best approximate the synthetic control unit earlier in Section 2. Table A1 in the Appendix summarizes the demographic and epidemiological characteristics of each country. The time period covers 87 days between 29 February and 25 May. The policy intervention occurs on 23 March, when the government-mandated lockdown measures are imposed in the control group countries while Sweden moves to a soft regime switch. Post-treatment period thus covers a two-month long period during which the two groups diverged in terms of voluntary social distancing vs. involuntary lockdown

mandate. In light of the earlier findings, I anticipate that the estimate of b_1 to be positive, which implies that the infection cases in Sweden would be higher than the control group countries that went for a strict lockdown measure. Additional controls summarized in **X** include the Google mobility index for six different categories as shown in Figure 7. With the exception of visits to parks and residential places, the correlation between the mobility index and the infection case per million is negative.

5.1. Results

The regression DD estimates for different specifications of the model in equation (1) are summarized in Table 5. The first row shows the average treatment effect. The specification in column (1) considers the effect of lockdown without any mobility index as controls. The estimated coefficient shows that Sweden, on average, had additional infection of 482.42 cases per million when compared to the infection rates of the control group. Considering that the case in Sweden reached around 3,300 per million population on 25 May, this implies that the infection rate in Sweden would have been lower by around 15 percent had it followed a strict lockdown policy like the other countries did. Specifications from columns (2) to (7) consider the effect of lockdown while individually controlling for behavioral changes, while the specification in column (8) allows for all of the mobility categories combined. The main findings on the average treatment effects remain robust, and even stronger in magnitude, to the inclusion of additional country-specific mobility controls. For example, allowing for all behavioral changes (as shown in the specification in column (8)), the infection rate in Sweden could have been lower by around 21 percent had it followed a lockdown policy like the other countries did.

5.2. Leads and lags

The key identifying assumption of DD regression design is a parallel (or common) trend assumption, meaning that—in the absence of treatment—the average change for the treated group would have been identical to the observed average change for the control group. In our setup, this implies that infection trends would have been the same in both Sweden and its control group had Sweden followed the same policy intervention path as the control group did. A rigorous verification is necessary, especially since our data set covers a lengthy period. An alternative way to deal with this issue—referred to by Autor (2003) as a "placebo" test—is to include leads in the baseline regression:

$$Y_{it} = \beta_0 + \sum_{j=0}^{q} b_j L_{i,t+j} + \tau_t + m_i + e_{it} \quad (2)$$

Covid Economics 35, 7 July 2020: 70-95

Table 5: Regression DD estimates

Variables	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lockdown	482.42***	689.70***	681.79***	836.81***	828.90***	836.28***	801.26***	690.26***
	(76.96)	(74.98)	(74.72)	(89.35)	(88.49)	(79.17)	(87.85)	(78.86)
Groceries		-16.67***						12.94*
		(1.80)						(6.30)
Parks			-5.28***					6.38***
			(0.41)					(1.89)
Transit stations				-19.43***				18.14
				(2.05)				(10.35)
Workplaces					-24.74***			38.06***
					(2.80)			(10.61)
Residential						55.70***		261.94***
						(4.78)		(34.83)
Retail							-13.59***	8.13
							(1.47)	(7.79)
N	348	348	348	348	348	348	348	348
R^2	0.886	0.908	0.905	0.912	0.914	0.922	0.907	0.935

Note: Robust standard errors in parentheses. All specifications include country and time fixed effects. * $p < .05$, ** $p < .01$, *** $p < .001$

The basic idea behind the test is that if a variable of interest, say $L_{i,t}$, causes outcome variables, say Y_{it}, future values of $L_{i,t}$ should not have any effect on Y_{it}. This type of a falsification test allows us to check for any anticipatory effect in days prior to the policy implementation. In our specification, I include leads of up to seven days allowing for the notification time prior to the policy enforcement.

Figure 8 plots the coefficients and confidence intervals leading to the lockdown intervention.[11] As all the leads are very close to 0, I find no indication of any positive anticipatory effect for all seven days leading up to the lockdown measure. This provides some confidence that the parallel trend assumption is not violated and that the policy intervention occurs before its effect.

So far, we have implicitly assumed that the coefficient b_1 in equation (1) is constant, implying that we estimate the average treatment effects (ATE) for the whole post-treatment period. However, the impact of lockdowns could be immediate or lagged over time, and may possibly vary with time. In fact, earlier findings point out that during the first few weeks

[11]Table A2 in the Appendix presents the regression estimates for the all the specifications considered in Table 5.

Covid Economics 35, 7 July 2020: 70-95

Figure 8: Placebo tests

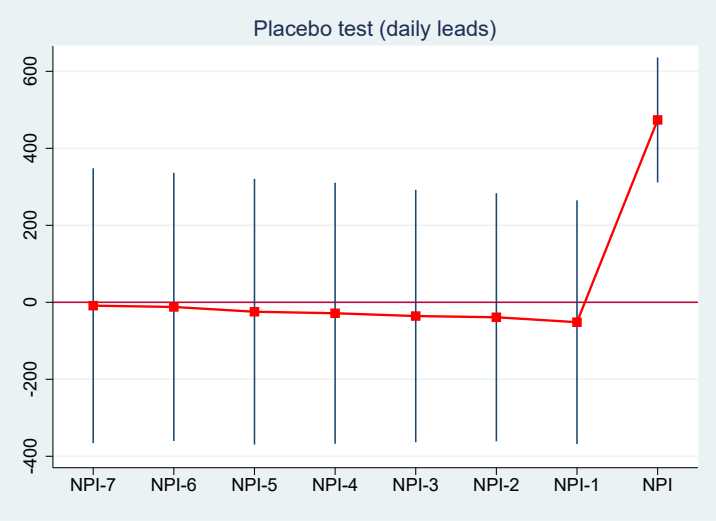

Note: NPI-n indicates n days prior to the policy intervention. NPI indicates all post-intervention period. Vertical lines represent 95 percent confidence intervals.

of post-lockdown, there was no discernible difference in the infection rates between Sweden and its synthetic counterpart. To explore the dynamic effects of the lockdown measures, I allow for lags in the regression specification as suggested by Autor (2003). More specifically, I add a dummy variable for each week up to the fifth week after the lockdown, as well as a dummy that captures all the weeks after week six since the lockdown is enforced. Each dummy variable takes the value of one in its relevant week. The modified specification with post-treatment dynamic effects is:

$$Y_{it} = \beta_0 + \sum_{j=0}^{q} b_j L_{i,t-7\times j} + \tau_t + m_i + e_{it} \tag{3}$$

Here, b_0 captures the immediate effect of lockdown in the initial week, while the b_j $(\forall j > 0)$ coefficients pick up any subsequent weekly effects. If $b_j > b_0 (> 0)$, this implies that the effect of the lockdown rises over time, while if the opposite is true then the initial impact fades with time.

Figure 9 plots the coefficients and the 95 percent confidence intervals allowing for lagged effects of the lockdown.[12] Similar to earlier findings under the synthetic control approach,

[12]Table A3 in the Appendix presents the regression estimates for the all the specifications considered in

Covid Economics 35, 7 July 2020: 70-95

Figure 9: Persistence effect of NPI

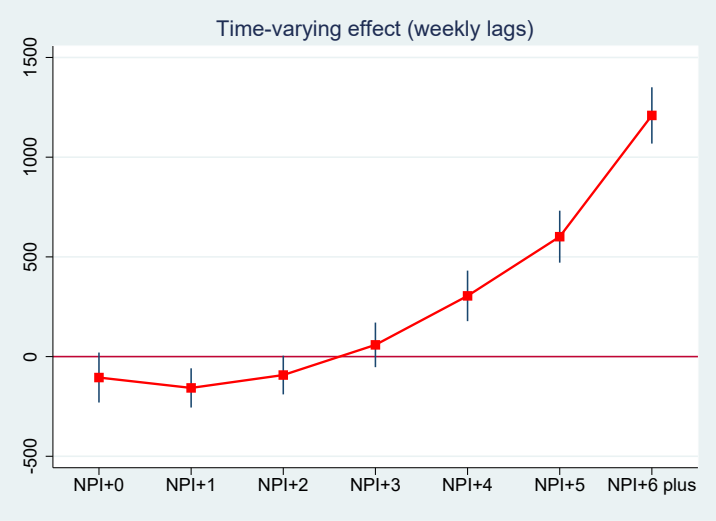

Note: NPI+n indicates n weeks after the policy intervention. NPI+6 plus indicates all post-intervention period after 6 weeks. Vertical lines represent 95 percent confidence intervals.

the estimated coefficients are not significant in the first four weeks. However, from week 5 onward, the coefficient becomes significantly positive and monotonically grows over time. This finding also coincides with the earlier outcome where the treatment effect becomes statistically significant five weeks after the implementation of the lockdown measures.

6. Conclusion

Policymakers have implemented a wide range of non-pharmaceutical interventions to fight the spread of COVID-19. Using variation in policies across countries and over time, I consider a synthetic control approach which is further complemented by a difference-in-differences (DD) research design to estimate causal effects of counter-COVID measures. I find that the lockdown measures played an important role in limiting the spread of the COVID-19 infection and that Swedish policymakers would have reduced the infection cases by more than half had they followed similar policies implemented elsewhere. I also find that as people dynamically adjust their behavior in response to information and policies, the impact of NPIs does not manifest immediately but only with a time lag of approximately 5 weeks or more.

Table 5.

Covid Economics 35, 7 July 2020: 70-95

One topic that the current study abstracts from is how each of the counter-COVID measures have different epidemiological impacts. A worthwhile project to pursue would be one that investigates the impact of individual measures along both epidemiological and economic aspects. Such explorations would better inform policymakers seeking to protect public health and facilitate an eventual economic recovery.

References

Abadie, A., 2020. Using synthetic controls: Feasibility, data requirements, and methodological aspects. Journal of Economic Literature (forthcoming) .

Abadie, A., Diamond, A., Hainmueller, J., 2010. Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. Journal of the American Statistical Association 105, 493– 505.

Abadie, A., Diamond, A., Hainmueller, J., 2015. Comparative politics and the synthetic control method. American Journal of Political Science 59, 495– 510.

Abadie, A., Gardeazabal, J., 2003. The economic costs of conflict: A case study of the Basque country. American Economic Review 93, 113–132.

Athey, S., Imbens, G.W., 2006. Identification and inference in nonlinear difference-in-differences models. Econometrica 74, 431–497.

Autor, D.H., 2003. Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. Journal of Labor Economics 21, 1–42.

Baldwin, R., di Mauro, B.W. (Eds.), 2020. Economics in the Time of COVID-19. CEPR Press.

Born, B., Dietrich, A.M., Müller, G.J., 2020. Do lockdowns work? A counterfactual for Sweden. Covid Economics 16, 1– 22.

Card, D., Krueger, A.B., 1994. Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. American Economic Review 84, 772–793.

Castex, G., Dechter, E., Lorca, M., 2020. COVID-19: Cross-country heterogeneity in effectiveness of non-pharmaceutical interventions. Covid Economics 14, 175–199.

Chen, X., Qiu, Z., 2020. Scenario analysis of non-pharmaceutical interventions on global Covid-19 transmissions. Covid Economics 7, 46– 67.

Covid Economics 35, 7 July 2020: 70-95

Chernozhukov, V., Kasahara, H., Schrimpf, P., 2020. Causal impact of masks, policies, behavior on early Covid-19 pandemic in the U.S. medRxiv .

Firpo, S., Possebom, V., 2018. Synthetic control method: Inference, sensitivity analysis and confidence sets. Journal of Causal Inference 6.

Fisher, R.A. (Ed.), 1935. The Design of Experiments. Oliver and Boyd.

Gonzalez-Eiras, M., Niepelt, D., 2020. On the optimal 'lockdown' during an epidemic. Covid Economics 7, 68– 87.

Goodman-Bacon, A., Marcus, J., 2020. Using difference-in-differences to identify causal effects of COVID-19 policies. Survey Research Methods 14, 153–158.

Hale, T., Webster, S., Petherick, A., Phillips, T., Kira, B., 2020. Variation in government responses to COVID-19. BSG Working Paper Series .

Imbens, G.W., Wooldridge, J.M., 2009. Recent developments in the econometrics of program evaluation. Journal of Economic Literature 47, 5–86.

Rocklöv, J., Sjödin, H., 2020. High population densities catalyse the spread of COVID-19. Journal of Travel Medicine 27.

Ullah, A., Ajala, O.A., 2020. Do lockdown and testing help in curbing Covid-19 transmission? Covid Economics 13, 138–156.

Covid Economics 35, 7 July 2020: 70-95

Appendix

Table A1: Composition of Synthetic Sweden

	Finland	France	Norway	Sweden
Weight for Synthetic Control	0.643	0.076	0.281	–
Population (million)	5.5	67.0	5.3	10.1
Urban population fraction (%)	85.4	80.4	82.2	87.4
Population density	18.1	122.3	14.5	24.7
Case per million (first 20 day)	25.9	46.1	80.7	42.8
Day 1	2 March	1 March	29 February	29 February
Lockdown Day	28 March	17 March	24 March	
Pre-lockdown duration (days)	26	16	24	
Government Stringency Index (SI) on Lockdown	68.5	90.7	75.9	

Figure A1: Government Stringency Index

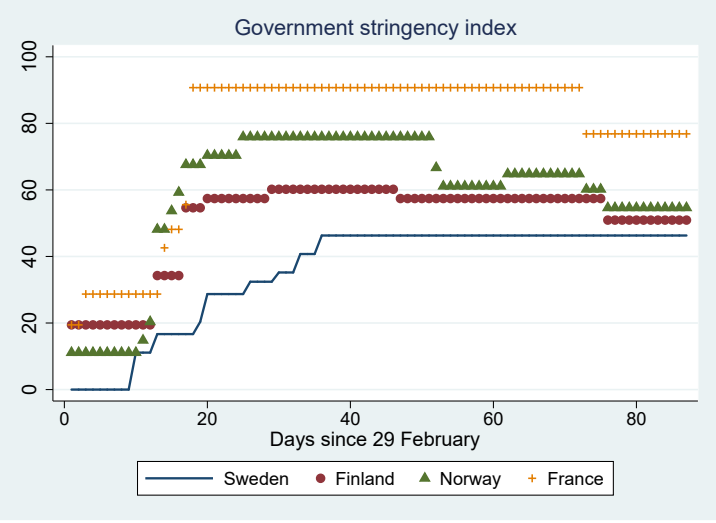

Note: Vertical axis measures the stringency index taken from Oxford COVID-19 Government Response Tracker since 29 February.

Covid Economics 35, 7 July 2020: 70-95

Covid Economics 35, 7 July 2020: 70-95

Figure A2: Profile of Infection Rates in Logs – Treatment vs. Synthetic Control

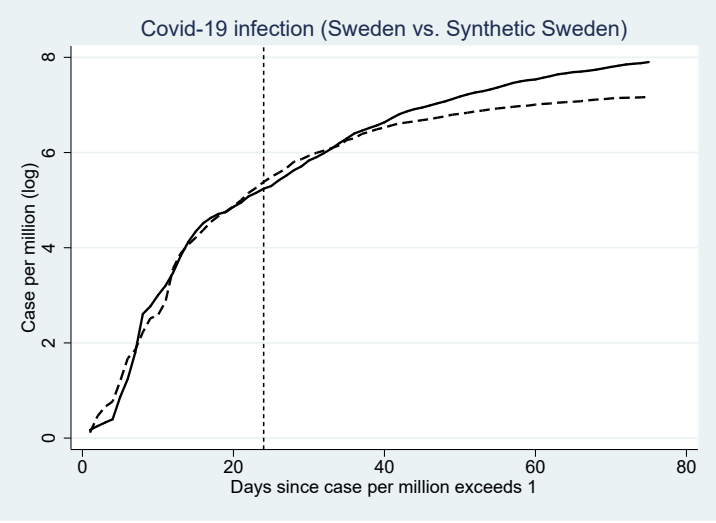

Note: Infection case per million population in logs is shown for Sweden versus synthetic Sweden (in dashed line).

CEPR PRESS

Covid Economics 35, 7 July 2020: 70-95

Table A2: Placebo test

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
NPI-7	-8.72	-4.17	8.30	65.05	5.16	14.92	10.90	-13.52
	(181.22)	(104.06)	(70.45)	(153.87)	(138.47)	(130.68)	(152.61)	(82.99)
NPI-6	-12.05	-28.31	7.24	107.97	66.34	52.72	33.46	43.89
	(176.80)	(103.93)	(65.93)	(165.85)	(158.12)	(147.74)	(158.23)	(105.72)
NPI-5	-24.47	-16.71	-0.90	151.43	143.24	111.57	66.17	134.20
	(175.18)	(109.72)	(63.97)	(167.27)	(171.14)	(146.64)	(163.75)	(102.96)
NPI-4	-28.45	28.17	-3.88	209.11	233.54	192.57	117.94	260.74*
	(172.18)	(116.06)	(74.54)	(170.29)	(181.40)	(144.90)	(171.86)	(103.88)
NPI-3	-35.67	94.64	29.67	283.79	324.32	289.52	187.71	365.27**
	(166.53)	(136.16)	(86.91)	(171.65)	(187.19)	(147.30)	(179.42)	(128.04)
NPI-2	-38.81	143.55	37.85	338.45	400.57*	376.84*	238.32	528.06*
	(163.62)	(130.20)	(104.06)	(176.68)	(192.60)	(159.41)	(185.90)	(206.78)
NPI-1	-51.65	153.94	30.05	373.79*	455.97*	402.39*	269.95	501.83*
	(160.66)	(134.97)	(122.95)	(189.25)	(220.07)	(181.22)	(196.99)	(220.04)
NPI	473.74***	707.71***	686.70***	918.05***	918.04***	910.81***	849.71***	796.83***
	(82.39)	(79.66)	(78.97)	(94.20)	(89.43)	(79.61)	(95.06)	(81.62)
Grocery_and_pharmacy		-16.82***						15.04*
		(1.84)						(6.15)
Parks			-5.28***					6.94***
			(0.42)					(1.91)
Transit_stations				-20.23***				15.39
				(2.15)				(10.65)
Workplaces					-26.05***			38.69***
					(2.93)			(10.87)
Residential						57.57***		268.51***
						(4.94)		(34.18)
Retail_and_recreation							-13.94***	8.48
							(1.56)	(7.70)
N	348	348	348	348	348	348	348	348
R^2	0.886	0.908	0.905	0.913	0.915	0.923	0.907	0.937

Note: NPI-n indicates n days prior to the policy intervention. NPI refers to the all post-treatment. Robust standard errors in parentheses. All specifications include country and time fixed effects. * $p < .05$, ** $p < .01$, *** $p < .001$

CEPR PRESS

Covid Economics 35, 7 July 2020: 70-95

Table A3: Time-varying effects of lockdown

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
NPI+0	-105.10	70.54	45.40	181.42*	212.33*	212.74**	118.84	77.95
	(63.65)	(72.26)	(63.18)	(80.93)	(87.72)	(81.10)	(74.63)	(72.30)
NPI+1	-157.16**	16.84	7.83	141.45*	159.44*	178.55**	95.43	70.17
	(49.91)	(60.16)	(46.00)	(64.10)	(69.62)	(65.20)	(61.94)	(62.48)
NPI+2	-92.50	51.98	110.46*	166.91**	168.81**	201.83***	149.49*	178.69**
	(49.27)	(58.31)	(55.55)	(58.80)	(57.57)	(56.19)	(62.96)	(54.69)
NPI+3	58.71	236.18***	292.40***	287.56***	228.25***	322.44***	290.21***	248.17***
	(56.76)	(55.10)	(54.56)	(60.41)	(61.93)	(56.46)	(59.15)	(66.82)
NPI+4	304.43***	418.25***	553.84***	556.74***	556.80***	581.13***	523.43***	617.89***
	(64.33)	(58.41)	(63.91)	(67.76)	(69.99)	(62.11)	(64.98)	(64.21)
NPI+5	601.36***	721.69***	764.95***	830.28***	836.36***	832.04***	829.47***	776.95***
	(66.14)	(58.15)	(41.91)	(61.18)	(62.72)	(55.31)	(61.85)	(57.24)
NPI+6 plus	1209.40***	1257.31***	1361.02***	1343.93***	1330.91***	1327.60***	1340.16***	1338.46***
	(71.70)	(64.50)	(69.74)	(64.67)	(61.99)	(59.89)	(64.53)	(74.89)
Grocery_and_pharmacy		-9.29***						14.87**
		(1.60)						(5.23)
Parks			-4.75***					-3.97**
			(0.37)					(1.53)
Transit_stations				-11.86***				33.93***
				(1.52)				(8.11)
Workplaces					-15.11***			-9.33
					(2.13)			(9.17)
Residential						35.99***		92.30**
						(3.63)		(30.52)
Retail_and_recreation							-8.43***	-5.54
							(1.15)	(6.47)
N	348	348	348	348	348	348	348	348
R^2	0.952	0.958	0.966	0.960	0.961	0.965	0.959	0.974

Note: NPI+n indicates n weeks after the policy intervention. NPI+6 plus indicates all post-intervention periods after 6 weeks. Robust standard errors in parentheses. All specifications include country and time fixed effects. * $p < .05$, ** $p < .01$, *** $p < .001$

Covid Economics **Issue 35, 7 July 2020**

Weighting bias and inflation in the time of Covid-19: Evidence from Swiss transaction data[1]

Pascal Seiler[2]

Date submitted: 29 June 2020; Date accepted: 1 July 2020

Sharp changes in consumer expenditure may bias inflation during the Covid-19 pandemic. Using public data from debit card transactions, I quantify these changes in consumer spending, update CPI basket weights and construct an alternative price index to measure the effect of the Covid-induced weighting bias on the Swiss consumer price index. I find that inflation was higher during the lock-down than suggested by CPI inflation. The annual inflation rate of the Covid price index was -0.4% by April 2020, compared to -1.1% of the equivalent CPI. Persistent "low-touch" consumer behaviour can further lead to inflation being underestimated by more than a quarter of a percentage point until the end of 2020.

Covid Economics 35, 7 July 2020: 96-115

1 I am thankful to Barbara Rudolf, Klaus Abberger, Martin Brown, Alexander Rathke and Samad Sarferaz for helpful comments and discussions on earlier drafts.

2 PhD student, ETH Zurich, and Researcher at the KOF Swiss Economic Institute.

Covid Economics 35, 7 July 2020: 96-115

1 Introduction

The Covid-19 pandemic and the measures enacted to contain it have led to a standstill of public life and a severe downturn of economic activity in many countries, including Switzerland. The measures implemented by the Swiss Federal Council – including lock-downs, mobility restrictions, and social-distancing rules – have greatly affected consumer expenditure patterns. Non-essential retail outlets and many service industries such as restaurants, bars as well as entertainment and leisure facilities were temporarily closed. Public transport services were reduced. Only grocery stores, pharmacies, banks and post offices were allowed to remain open (Eichenauer and Sturm, 2020). During the roughly two-month[1] lock-down period, consumer spending was thus severely restricted.

Sudden and profound changes like these can introduce significant bias in the consumer price index (CPI) used to measure inflation. The CPI is compiled on the basis of expenditure weights that are kept constant within a given year, reflecting the purpose of the index to measure changes in prices only without accounting for adjustments in consumption patterns. Most national statistical offices update their CPI expenditure weights once a year, often with lagged expenditure data.[2] While this practice is reasonable in normal times, it makes inflation indices much harder to interpret during the Covid-19 pandemic (Tenreyro, 2020; Lane, 2020), as the underlying weighting scheme is no longer representative of what is being consumed or what can be consumed at all in the lock-down period, thus introducing a *weighting bias* in inflation.

In this paper I study the effect that biases induced by such changes in spending patterns have on the measurement of inflation in Switzerland during the Covid-19 crisis. For this purpose, I use high-frequency estimates of spending based on transactional data to update CPI basket weights and compute an alternative price index based on such "Covid weights".

I find that inflation was higher during the lock-down than suggested by CPI inflation.

[1] Commercial activity was particularly affected from 13 March to 11 May 2020 due to the containment measures taken by the authorities. See Appendix A for a chronology of the events and measures taken by the Swiss government.

[2] The Swiss Federal Statistical Office (FSO) updated the CPI weights for 2020 in December 2019 using expenditure information collected back in 2017 and 2018. The main source for calculating the basket weights is the Household Budget Survey (HBS), conducted annually by the FSO among private households with permanent residence in Switzerland.

Covid Economics 35, 7 July 2020: 96-115

The annual inflation rate of the Covid price index was −0.4% by April 2020, compared to −1.1% of the equivalent CPI. This is a consequence of the relative increase in consumption of "Food & Non-Alcoholic Beverages", which are more inflationary than other spending categories.

Moreover, a persistent change in consumer behaviour – driven by "low-touch" considerations due to new working habits, prolonged uncertainty and the lifestyle adopted during the lock-down period – is likely to keep underestimating short to medium-term inflation. The inflation forecast based on an alternative "low-touch" consumption basket is over a quarter of a percentage point higher than a comparable CPI forecast for the rest of 2020.

This study is closely related to two current research trends. First, it contributes to the rapidly growing literature making use of alternative and high-frequency data sources to track consumer expenditure. In the rapidly evolving Covid-19 crisis there is great need for reliable data that is available in almost real-time. Official economic statistics, however, are usually published with a considerable lag. This has led many researchers to explore alternative and high-frequency data sources to track the pandemic and its effects. Transactions data from banks and other financial institutions have proved to be particularly fruitful sources and were used for the analysis of consumer spending, among others, by Baker et al. (2020) and Chetty et al. (2020) (for the US), Chen et al. (2020) (for China), Andersen et al. (2020) (for Denmark), Bounie et al. (2020) (for France), Carvalho et al. (2020) (for Spain) and Hacioglu et al. (2020) (for the UK).

Second and more specifically, it contributes to studies of inflation and potential biases of it during the Covid-19 crisis. My results are consistent with the analytical argument of Diewert and Fox (2020) who show that a downward bias in consumer price indices may result from current calculation methodologies. Using scanner data of fast-moving consumer goods in the UK, Jaravel and O'Connell (2020) empirically document a spike in inflation in the first month of lock-down. Using official price indices and updating CPI weights in a similar way to this study, Cavallo (2020) finds comparable results for the US but overall mixed international evidence.

In light of the fact that the CPI is an essential tool for economic policy making, my results have important implications for the crisis period and beyond. They provide evidence

Covid Economics 35, 7 July 2020: 96-115

that conventional price measures have underestimated inflation during the crisis. By unveiling this, I hope that my results contribute to the assessment of inflation in times of economic turmoil. Beyond, they raise conceptual issues concerning adequate price measurement. While most of the changes in prices and consumption expenditure can be expected to reverse once the crisis is over, some of them may be more persistent. Considering current calculation methodologies, this can make official CPI measures less informative, in particular during the transition period. In response to this challenge, the use of high-frequency and alternative data sources on both prices and consumer spending may become key to producing a more robust and informative consumer price index in the future.

The remainder of this paper is organised as follows. In Section 2, I describe how I measure changes in consumer spending, update CPI weights and construct the alternative Covid price index. Section 3 estimates the effects of the Covid-induced weighting bias on Swiss consumer price index. Section 4 assesses how a lasting "low-touch" consumer behaviour will affect inflation in the short to medium term. Section 5 concludes.

2 Data and methodology

Constructing the Covid weights To construct the Covid basket weights, I use weekly data on Swiss debit card expenditure that are publicly available as part of the Consumption Monitoring for Switzerland.[3] They are produced using transactions of debit cards issued by banks to their customers in Switzerland and include debit card payments at points of sale such as grocery stores or service providers (e.g. hairdressers, restaurants or petrol stations). Figure 1 depicts the change in Swiss consumption patterns since January 2020 based on these debit card expenditures. The vertical dotted line coincides with 16 March, when the *extraordinary situation* was declared by the Swiss Federal Council and when most shops, restaurants and leisure facilities were temporarily closed.

[3]The Consumption Monitoring for Switzerland is a project powered the University of St. Gallen (Prof. Martin Brown, Prof. Matthias Fengler) and Novalytica together with Dr. Robert Rohrkemper (Distinguished Expert, Senior Data Scientist at Worldline) and Prof. Rafael Lalive (University of Lausanne). See http://monitoringconsumption.org/switzerland.

Covid Economics 35, 7 July 2020: 96-115

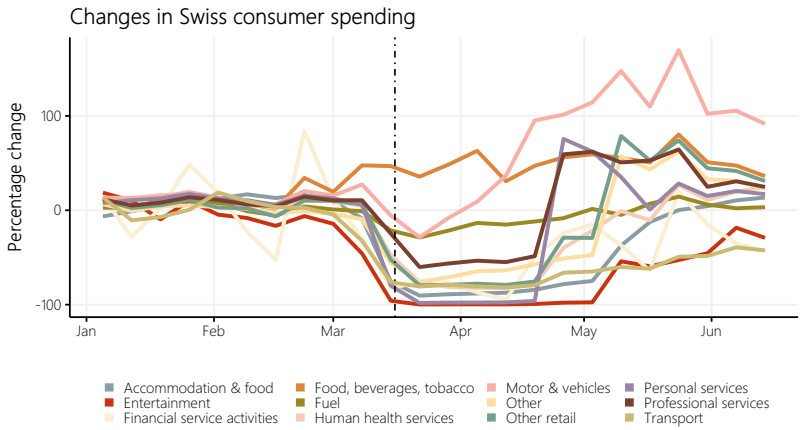

Figure 1: Changes in Swiss consumer spending before and during the Covid-19 crisis as measured by debit card transaction volumes. Cumulative expenditure change across categories of goods and services in Switzerland since January 2020 (normalisation month). The vertical dotted line marks the declaration of the *extraordinary situation* by the Swiss Federal Council on 16 March 2020.

Three observations stand out.[4] First, there was a massive drop in consumer spending which started even before the introduction of lock-down measures on 16 March. This is suggestive of uncertainty and consumer confidence being substantial drivers of the fall in consumption expenditure. Second, different expenditure categories are affected differently by the official measures. While it was virtually impossible to spend money on entertainment or personal services, consumption expenditure for groceries has increased significantly before and during the lock-down period. Third, expenditure categories recover heterogeneously after the lock-down. Expenditure on personal and professional services (including hairdressers) spiked at the end of April, when a first opening step was taken. "Other Retail" (including garden stores, clothing and furniture) is stepwise retracing the two relaxations of 27 April and 11 May. While most categories are at least partially recouping their losses, "Accommodation & Food" and "Transport" are recovering only very slowly, and are currently still well below their levels at the beginning of the year.

[4]These findings are corroborated by international studies tracking consumption similarly with high-frequency credit and debit card transactions. See Baker et al. (2020) and Chetty et al. (2020) for the US, Chen et al. (2020) for China, Andersen et al. (2020) for Denmark, Bounie et al. (2020) for France, Carvalho et al. (2020) for Spain and Hacioglu et al. (2020) for the UK. Beyond, Coibion et al. (2020) find comparable survey-based estimates.

I combine these estimates for consumer spending with official CPI data from the Swiss Federal Statistical Office. In particular, I use the sectoral CPI series that form the lowest level of disaggregation of the CPI (i.e. *expenditure items*, not seasonally adjusted), as well as the latest available weights in the official CPI basket (i.e. expenditure weights for 2020).

Matching debit card transaction categories with the CPI expenditure items requires some assumptions. Table 1 shows the correspondence table for the CPI main groups.

Table 1: Matching Swiss CPI main groups and debit card transaction categories.

	CPI main group	Debit card transaction category	Weight change (Apr vs. Jan 2020, in %)
01	Food & Non-Alcoholic Beverages	Food, beverages, tobacco	36.0
02	Alcoholic Beverages & Tobacco	Food, beverages, tobacco	36.0
03	Clothing & Footwear	Other retail	−89.7
04	Housing & Energy	Other, None	−7.93
05	Furniture & Home Maintenance	Other, Other retail, Professional services	−83.9
06	Health	Human health services	−98.2
07	Transport	Fuel, Motor & vehicles, Transport	−17.8
08	Communications	Other retail	−89.7
09	Recreation & Culture	Entertainment, Other, Other retail, Professional services, Transport	−92.6
10	Education	None	0.0
11	Restaurants & Hotels	Accommodation & food	−98.2
12	Miscellaneous Goods & Services	Financial service activities, Other, Other retail, Personal services, Professional services	−83.3

About five categories are closely matched in both datasets. In particular, there are almost one-to-one mappings for the main groups "Health", "Transport" and "Restaurants & Hotels". For both "Food & Non-Alcoholic Beverages" and "Alcoholic Beverages & Tobacco", I use the debit card transaction category "Food, beverages, tobacco". Further, "Clothing & Footwear", "Furniture & Home Maintenance" and "Communications" are contained in "Other retail". Items of "Recreation & Culture" and "Miscellaneous Goods & Services" are matched with multiple transaction categories. Finally, as I do not find any corresponding debit card transaction categories for rents in "Housing & Energy" and all items in "Education", I assume that expenditure for these items remains unchanged.

To estimate the expenditure shares in the Covid basket, I start with the latest official CPI expenditure weights $w_{i,0}$, multiply them by the average percentage change in the corresponding expenditure category each month, and normalise them as a share of the total. Formally, the Covid weight of CPI expenditure item i in month t is thus given by

Covid Economics 35, 7 July 2020: 96-115

$$w'_{i,t} = \frac{w_{i,0}\Delta e_{i,t}}{\sum_i w_{i,0}\Delta e_{i,t}}, \tag{2.1}$$

where $\Delta e_{i,t} = \frac{P_{i,t}Q_{i,t}}{P_{i,0}Q_{i,0}}$ is the change in consumer expenditure since January 2020 (base month) as measured by debit card purchases volumes $P_{i,t}Q_{i,t}$. Equation 2.1 highlights the fact that these are *relative* weights, so the importance of a category in the basket can change even when its expenditure does not.

Constructing the Covid price index Two concepts of price index calculation are most common: the Laspeyres index and the Paasche index. Conceptually, they differ in one important aspect. The Laspeyres index answers the question of how much the old basket of goods and services costs at current prices. The Paasche index, on the other hand, answers the question of how much the current basket at current prices costs in relation to the current basket at old prices.

Given the Covid-induced monthly changes in consumption expenditure, I favour the use of a Paasche index for the purpose of this study. Comparing index values with the same underlying weighting reduces fluctuations due to changes in expenditure quantities and allows to capture more precisely changes in the price level only. This approach herein differs from similar contributions of the literature, in particular from Cavallo (2020) who, through the monthly variation of both prices and weights, blurs the boundary between the CPI as a tool for pure price measurement and the CPI as a mere turnover statistic.

Consequentially, I compute the Covid price index as the weighted sum of sectoral CPI indices using two distinct weighting schemes in any month, namely

$$I_{t,w=t} = \sum_{i=1}^{12} w'_{i,t} I_{i,t} \tag{2.2}$$

and

$$I_{t-k,w=t} = \sum_{i=1}^{12} w'_{i,t} I_{i,t-k} \tag{2.3}$$

where $k \in \{1, 12\}$ for the calculation of either the monthly or annual inflation rates. These inflation rates are then calculated by comparing the indices of the same weighting scheme, i.e.

Covid Economics 35, 7 July 2020: 96-115

$$\pi_{t,t-1} = \frac{I_{t,w=t} - I_{t-1,w=t}}{I_{t-1,w=t}} \cdot 100 \tag{2.4}$$

in case of monthly inflation, and

$$\pi_{t,t-12} = \frac{I_{t,w=t} - I_{t-12,w=t}}{I_{t-12,w=t}} \cdot 100 \tag{2.5}$$

in case of annual inflation.

3 Inflation with Covid consumption

Figure 2 and Table 2 show the alternative Covid price index and thus illustrate the effect of the Covid-induced changes in consumer spending on both the monthly and annual inflation rates of the Swiss consumer price index.

Covid Economics 35, 7 July 2020: 96-115

Covid Economics 35, 7 July 2020: 96-115

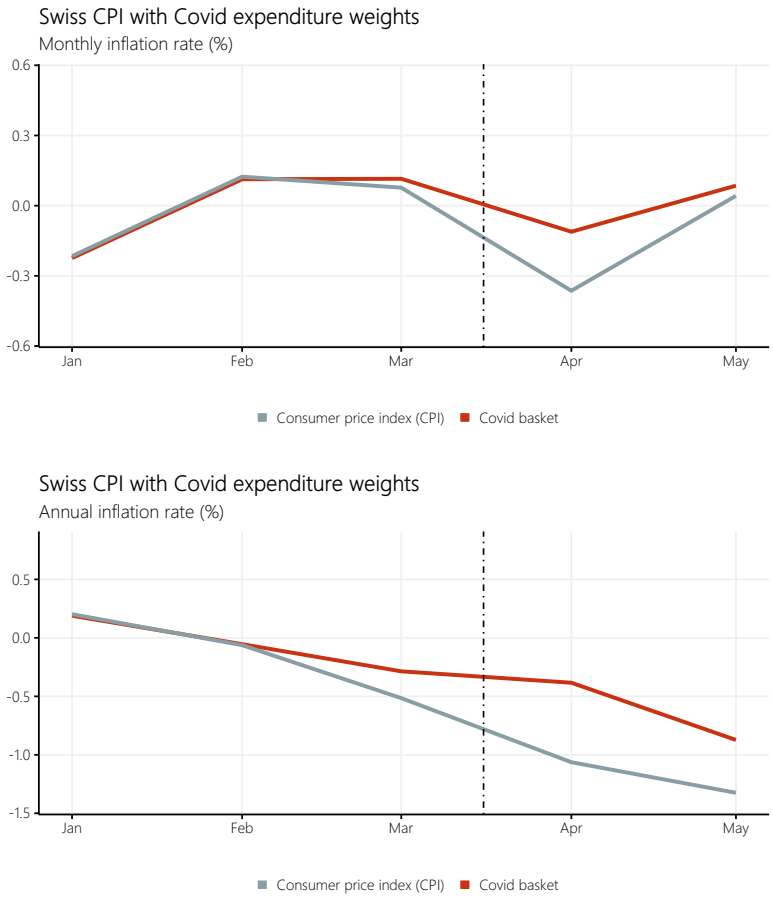

Figure 2: Swiss consumer price index (CPI, not seasonally adjusted) and the Covid index construed using estimates of the consumption expenditure shares during the Covid-19 crisis ("Covid basket"). The vertical dotted line marks the declaration of the *extraordinary situation* by the Swiss Federal Council on 16 March 2020.

The Swiss CPI was low but relatively stable in the beginning of the year, before it started showing deflation from February onward. During the Covid-19 crisis, it contracted strongly. It fell by -0.36% in April 2020 compared with the previous month. Inflation was -1.05% compared with the same month of the previous year. After lock-down in May 2020, it increased slightly by 0.05%.

The Covid price index, by contrast, was consistently higher during the crisis. It contracted by -0.11% only in April 2020 compared with the previous month, and inflation

was −0.38% compared with the same month of the previous year. After lock-down in May 2020, it increased slightly by 0.09%.

Table 2: Swiss inflation rates during the Covid-19 crisis

	Monthly inflation		Annual inflation	
	CPI	Covid basket	CPI	Covid basket
January	-0.22	-0.22	0.19	0.19
February	0.13	0.11	-0.07	-0.05
March	0.08	0.11	-0.51	-0.29
April	-0.36	-0.11	-1.05	-0.38
May	0.05	0.09	-1.32	-0.87

Monthly and annual inflation rate in the not seasonally adjusted Swiss CPI and the Covid index construed using estimates of the consumption expenditure shares during the Covid-19 crisis.

To illustrate what is driving these results, Table 3 shows the category weights and incidence details for April 2020. The second column has the monthly CPI sector inflation used in both the official and Covid price index. The third and fourth columns show the weights of the CPI and Covid basket in each category. Finally, the last two columns show the incidence that each category has on the total monthly inflation rate. The incidence is the monthly inflation rate multiplied by the weight. Therefore, the sum of all the category incidence numbers is equal to the monthly inflation rate.

Table 3: Swiss CPI weights and incidence in April 2020

	Monthly	Weight		Incidence	
Main group	CPI inflation	CPI	Covid basket	CPI	Covid basket
Food and non-alcoholic beverages	0.71	10.54	23.70	0.07	0.17
Alcoholic beverages and tobacco	-0.47	2.76	6.20	-0.01	-0.03
Clothing and footwear	0.26	3.40	1.52	0.01	0.00
Housing and energy	-0.19	24.96	40.27	-0.05	-0.08
Household goods and services	-0.80	3.79	1.69	-0.03	-0.01
Healthcare	-0.08	15.69	6.11	-0.01	-0.01
Transport	-1.74	10.97	7.90	-0.19	-0.14
Communications	-0.06	2.94	4.75	-0.00	-0.00
Recreation and culture	-0.89	8.37	0.08	-0.07	-0.00
Education	0.00	1.00	1.61	0.00	0.00
Restaurants and hotels	-0.75	9.46	2.04	-0.07	-0.02
Other goods and services	-0.09	6.12	4.13	-0.01	-0.00

The CPI weight is the share of expenditure in a given category over total expenditures. The incidence is the monthly inflation rate multiplied by the weight. The sum of all the category incidence numbers is equal to the monthly inflation rate.

Note that main groups whose spending does not change over time per assumption (e.g. "Education") can have different Covid weights due to the normalisation of the Covid basket as a share of total basket expenditure. Table 3 illustrates that the result is driven by shifts in relative basket weights. The Covid inflation rate is higher than CPI inflation because the index based on Covid weights gives more weight to main groups that have a positive inflation rate, and less weight to categories experiencing deflation.

Covid Economics 35, 7 July 2020: 96-115

Covid Economics 35, 7 July 2020: 96-115

In particular, the weight for "Food & Non-Alcoholic Beverages" rose from 10.54% to 23.70%, increasing the incidence of this category from 0.07 to 0.17. At the same time, the weight for the deflationary category "Recreation & Culture" fell from 8.37% to 0.08%, virtually eliminating the influence of this main group on the total monthly inflation rate.

Altogether, the weighting bias in official CPI statistics seems to have underestimated inflation. After adjusting for the change in consumer spending during the pandemic, the inflation rate of the Covid price index is two thirds percentage points higher and lies at -0.38% in April 2020. This result is driven by the relative weight shifts, positively reinforcing inflationary CPI main groups, and negatively reinforcing deflationary main groups.

4 Inflation with low-touch consumption

Having quantified the effect during the lock-down period raises the question of how any lasting change in consumer behaviour will affect inflation in the short to medium term.

Most of the changes in consumer spending during the Covid-19 crisis are obviously driven by the containment measures and forced closures of many retail sectors. Once these restrictions are lifted, catch-up effects can be expected, and consumer spending will gradually converge back to pre-crisis levels, as apparent in Figure 1.

Nevertheless, an immediate rebound and normalisation to pre-crisis levels seems unlikely. Rather, getting back to pre-Covid consumption patterns is likely to be a long and difficult task. Figure 3 compares the decline in consumer spending in the Covid-19 crisis with the earlier contractions of Swiss households' final consumption expenditure since 1980.

Covid Economics 35, 7 July 2020: 96-115

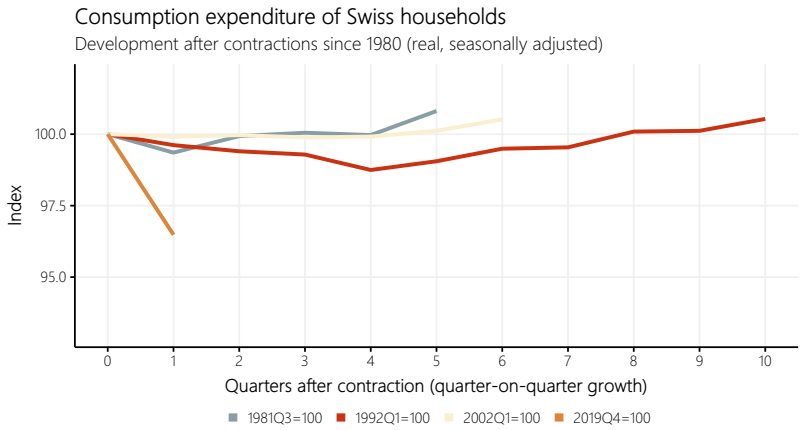

Figure 3: Development of Swiss households' final consumption expenditure (real, seasonally adjusted) after contractions (quarter-on-quarter growth) since 1980. Included are contractions with negative quarterly growth rates in at least two consecutive quarters.

So far, there have been a total of three downturns (of two consecutive quarters or more) in consumption, related to the second oil crisis (1981), the Swiss real estate crisis (1992) and the dot-com bubble (2002). Even if consumers return to normality at different speeds depending on the type of disruption that hit the economy, Figure 3 illustrates how unprecedented the current downturn in consumption is. Measured against earlier recoveries, the normalisation after the current crisis is likely to extend over several quarters.

Moreover, consumer behaviour may have changed persistently during the crisis, and henceforth be driven by "low-touch" considerations. This kind of "low-touch" consumption is characterised first and foremost by continuing uncertainty about the spread of the virus, changed working habits and the lock-down lifestyle.[5]

The lock-down has led to working from home in many professions. Once efficient conditions for teleworking are in place, they are likely to stay, which in turn creates more opportunities for eating at home, and reduces eating out. This negatively affects the hospitality and tourism sectors, which have already been severely impacted by the far-reaching travel freeze. In view of the great uncertainties regarding the spread of the

[5]For survey evidence in support of this "low-touch" consumption scenario, see for example Boston Consulting Group (2020) or McKinsey & Company (2020).

Covid Economics 35, 7 July 2020: 96-115

virus and the travel preparations which have become fairly complex (administrative effort to obtain information on the local situation and regulations), it seems plausible that expenditure for travel and tourism will remain subdued in the short and medium term. Further, the lock-down period has had a negative impact upon leisure activities. Digital activities such as online gatherings, at-home entertainment as well as remote learning and exercise have emerged. Outdoor and fitness activities have replaced going to the gym. These activities are likely to remain strong and reduce the share of spending on traditional "recreation & culture" activities reflected in the CPI.

In the following I will examine the effects this kind of prolonged change in consumer spending has on inflation in the short to medium term. I do so by producing inflation forecasts based on disaggregated CPI data, which I aggregate using two different weighting schemes: the official CPI weights, and the expenditure weights implied by the "low-touch" scenario.

Table 4 compares the latest CPI weights with the weights implied by this "low-touch" scenario. I calculate the latter by applying fixed markups (or markdowns) to the official CPI weights for 2020.

Table 4: Swiss CPI weights and "low-touch" weights

Main groups	CPI weights	Low-touch weights	Low-touch add-on
Food and non-alcoholic beverages	10.31	12.00	15%
Alcoholic beverages and tobacco	2.91	2.94	
Clothing and footwear	3.59	3.63	
Housing and energy	25.31	25.60	
Household goods and services	3.67	3.71	
Healthcare	14.33	14.49	
Transport	11.58	11.04	-7.7%
Communications	3.11	3.14	
Recreation and culture	8.43	7.45	-3.8%
Education	1.05	1.06	
Restaurants and hotels	9.45	8.61	-10%
Other goods and services	6.25	6.32	

The CPI weight is the share of expenditure in a given category over total expenditures. The low-touch weights are the assumed expenditure shares in the "low-touch" scenario. They are calculated by adding fixed markups (or markdowns) to the CPI weights for 2020. The fourth column gives the average add-on per main group. Weights are normalised and applied from June 2020 until the end of the forecast period.

I assume that spending on groceries will continue to fall and approach the original level, but then remain at +15% compared to January 2020. Conversely, out-of-home expenditure in restaurants and hotels will fall by 10%. In addition, consumption expenditure on selected transport services (air transport and package holidays in particular) and leisure activities (such as leisure courses, cinema, theatre and concert) are assumed to decrease,

Covid Economics 35, 7 July 2020: 96-115

leading to the overall differences of -7.7% and -3.8% in the corresponding main groups. As the imposed change in weight is negative overall and the weights are normalised as a share of total expenditure, the relative "low-touch" weights of the otherwise unchanged main groups end up marginally increased. These weighting schemes are applied from June 2020 until the end of the forecast period.

The disaggregated forecasts of the CPI items are based on univariate ARIMA models. For each expenditure item, the model is selected in two steps. First, the statistical properties of each price series in analysed. For simplicity, I assume that all items are integrated of order one and use them accordingly in first log-differences. Further, seasonal patterns are detected through inspecting the autocorrelation function (ACF) and considering the price collection frequency. Second, the lag selection of each model is automated using the Schwarz information criterion.

I forecast the CPI items with the monthly series from May 2000 (where available) to December 2021 and then aggregate the forecasts using the weightings schemes based on the two scenarios. For the benchmark scenario, I use the official CPI expenditure weights for 2020 throughout. For the "low-touch" scenario, I take the Covid weights from Section 3 until May 2020, and apply the low-touch weights (Table 4) from June 2020 onwards.

Both indices are shown in Figure 4 as Laspeyres indices, ensuring comparability.

Covid Economics 35, 7 July 2020: 96-115

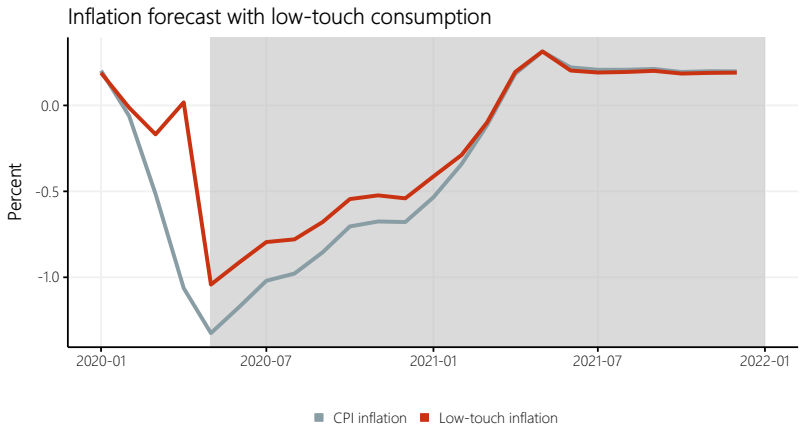

Figure 4: Comparison of inflation forecasts with official CPI weights and with "low-touch" consumption.

CPI inflation is projected to increase gradually from its low in May 2020 (−1.3%) but will remain considerably below 0% well into 2021. An average inflation rate of −0.72% is expected for 2020, and 0.12% for 2021. In contrast, inflation under the "low-touch" scenario is persistently higher. Particularly in the current year, inflation rates differ significantly. In the low-touch scenario the average annual inflation is −0.43%. Only in the course of next year the two forecasts converge.

Thus, without taking into account a sustained change in consumer behaviour based on low-touch considerations, the official CPI statistics risks underestimating inflation in the current year by more than a quarter of a percentage point.

It is difficult to say to what extent the presumed changes in consumer behaviour are really permanent. It is probable that the longer the crisis lasts, the more the new behaviours will gradually become the new normality and continue after the pandemic. The extent to which consumer behaviour and spending will change in the medium to long term after the recession is therefore likely to depend to a large extent on the further course of the pandemic, as well as on how the new work experiences are integrated into existing work habits and consumers' assessment of their future prospects.

5 Conclusion

Measuring and interpreting inflation is challenging during economic disruptions in general and the Covid-19 crisis in particular. Consumer spending is greatly affected by the pandemic containment measures, introducing a weighting bias into the measurement of CPI inflation.

Using public data from debit card transactions in Switzerland, I estimate the changes in consumption expenditure during the Covid-19 crisis, and construct an alternative price index with updated Covid consumption weights to study the effect of these sources of bias on Swiss CPI inflation.

I find that Covid inflation was higher than suggested by CPI inflation. By April 2020, the annual inflation rate of the Covid price index was -0.38%, compared to -1.05% of the CPI. This is a consequence of the relative increase in consumption of "Food & Non-Alcoholic Beverages".

Moreover, a persistent change in consumer behaviour – driven by "low-touch" considerations due to new working habits, prolonged uncertainty and the lifestyle adopted during the lock-down period – keeps underestimating short to medium-term inflation throughout the year by more than a quarter of a percentage point. In 2020, CPI inflation is projected to average -0.72%, while Covid inflation is projected to average -0.43%.

In light of the fact that the CPI is an essential tool for economic policy making, my results have important implications for the crisis period and beyond. They provide evidence that conventional price measures have underestimated inflation during the crisis. By unveiling this, I hope that my results contribute to the assessment of inflation in times of economic turmoil. Beyond, they raise conceptual issues concerning adequate price measurement. While most of the changes in prices and consumption expenditure can be expected to reverse once the crisis is over, some of them may be more persistent. Considering current calculation methodologies, this can make official CPI measures less informative, in particular during the transition period. In response to this challenge, the use of high-frequency and alternative data sources on both prices and consumer spending may become key to producing a more robust and informative consumer price index in the future.

Covid Economics 35, 7 July 2020: 96-115

Covid Economics 35, 7 July 2020: 96-115

References

Andersen, A., Hansen, E. T., and Sheridan, A. (2020). Consumer responses to the COVID-19 crisis: Evidence from bank account transaction data. CEPR Discussion Paper DP14809.

Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., and Yannelis, C. (2020). How does household spending respond to an epidemic? consumption during the 2020 COVID-19 pandemic. NBER Working Paper 26949, National Bureau of Economic Research.

Boston Consulting Group (2020). COVID-19 consumer sentiment snapshot Nr. 11: Getting to the other side. Koslow, lara and lee, jean, https://www.bcg.com/publications/2020/covid-consumer-sentiment-survey-snapshot-6-02-20.aspx [Online; Accessed: 29 June 2020].

Bounie, D., Camara, Y., and Galbraith, J. W. (2020). Consumers' mobility, expenditure and online-offline substitution response to COVID-19: Evidence from French transaction data. *Available at SSRN 3588373*.

Carvalho, V. M., Hansen, S., Ortiz, A., Garcia, J. R., Rodrigo, T., Mora, S. R., and de Aguirre, P. R. (2020). Tracking the COVID-19 crisis with high-resolution transaction data. CEPR Discussion Paper DP14642.

Cavallo, A. (2020). Inflation with Covid consumption baskets. NBER Working Paper 27352, National Bureau of Economic Research.

Chen, H., Qian, W., and Wen, Q. (2020). The impact of the COVID-19 pandemic on consumption: Learning from high frequency transaction data. *Available at SSRN 3568574*.

Chetty, R., Friedman, J. N., Hendren, N., Stepner, M., and The Opportunity Insights Team (2020). Real-time economics: A new platform to track the impacts of COVID-19 on people, businesses, and communities using private sector data. NBER Working Paper 27431, National Bureau of Economic Research.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). The cost of the Covid-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. NBER Working Paper 27141, National Bureau of Economic Research.

Diewert, E. E. and Fox, K. J. (2020). Measuring real consumption and CPI bias under lockdown conditions. NBER Working Paper 27144, National Bureau of Economic Research.

Eichenauer, V. and Sturm, J.-E. (2020). Die wirtschaftspolitischen Maßnahmen der Schweiz zu Beginn der Covid-19-Pandemie. *Perspektiven der Wirtschaftspolitik*.

Hacioglu, S., Känzig, D. R., and Surico, P. (2020). Consumption in the time of Covid-19: Evidence from UK transaction data. CEPR Discussion Paper DP14733.

Jaravel, X. and O'Connell, M. (2020). Inflation spike and falling product variety during the Great Lockdown. CEPR Discussion Paper DP14880.

Lane, T. (2020). Policies for the great global shotdown and beyond. https://www.bankofcanada.ca/2020/05/policies-great-global-shutdown-and-beyond/ [Accessed: 2020 06 16].

McKinsey & Company (2020). Consumer sentiment is evolving as countries around the world begin to reopen. Technical report, https://www.mckinsey.com/business-functions/marketing-and-sales/our-insights/a-global-view-of-how-consumer-behavior-is-changing-amid-covid-19 [Online; Accessed: 29 June 2020].

Tenreyro, S. (2020). Monetary policy during pandemics: inflation before, during, and after Covid-19. Speech, https://www.bankofengland.co.uk/speech/2020/silvana-tenreyro-speech-monetary-policy-during-pandemics [Online, Accessed: 17 June 2020].

Covid Economics 35, 7 July 2020: 96-115

Covid Economics 35, 7 July 2020: 96-115

A The stages of Swiss lock-down

28 February	The Swiss Federal Council categorises the situation in Switzerland as *special* in terms of the Epidemics Act. Events with more than 1000 persons are prohibited effective immediately. Among other things, the Basel Carnival and the Geneva Motor Show are cancelled.
13 March	The government announces the closure of the schools on Monday 16 March. Events with more than 100 people are prohibited. Restaurants, bars and discos are limited to 50 people.
16 March	Federal President declares the *extraordinary situation*, allowing the Federal Council to order the introduction of uniform measures in all cantons. All public and private events are prohibited. Shops, restaurants and leisure facilities must close. The lock-down also applies to schools and businesses at which the recommended distance cannot be maintained (e.g. hairdressers and cosmetics studios). Only grocery stores and health facilities remain open. Border controls at the borders with Germany, Austria and France were introduced, and entry bans imposed, albeit with exceptions. Border checks at the Italian border were already introduced at an earlier stage. Up to 8000 members of the armed forces were deployed to assist the cantons at hospitals and with logistics and security.
27 April	First step towards opening: hairdressers, DIY stores and garden centres may resume operations with protection concepts.
11 May	Shops, restaurants, public markets and museums may reopen. Primary and secondary schools can again teach on site.
6 June	Events with up to 300 people are permitted again. Mountain railways, camping sites, zoos and leisure facilities may open. Secondary, vocational and higher education establishments may resume teaching.
15 June	The borders to all states within the EU/EFTA area will be opened completely. Among other things, shopping tourism to Germany or Austria is permitted again.
19 June	Return from the extraordinary to the special situation. The cantons will have a greater say and more room for manoeuvre. In public spaces, the minimum distance is reduced from 2 to 1.5 metres. Restaurants will be allowed to move their tables closer together, while at the same time the Swiss midnight curfew will be lifted. Meetings and events for up to 1000 people are again permitted. Masks are compulsory at rallies. The recommendation to work from home if possible is repealed.

Table 5: Timeline of the events and measures taken by the Swiss Federal government. Compiled from the ordinances and media releases of the Swiss Federal Council, see https://www.admin.ch/gov/en/start/documentation.html. Situation as of 24 June 2020.

Covid Economics 35, 7 July 2020: 96-115

B Comparison of Paasche and Laspeyres price indices

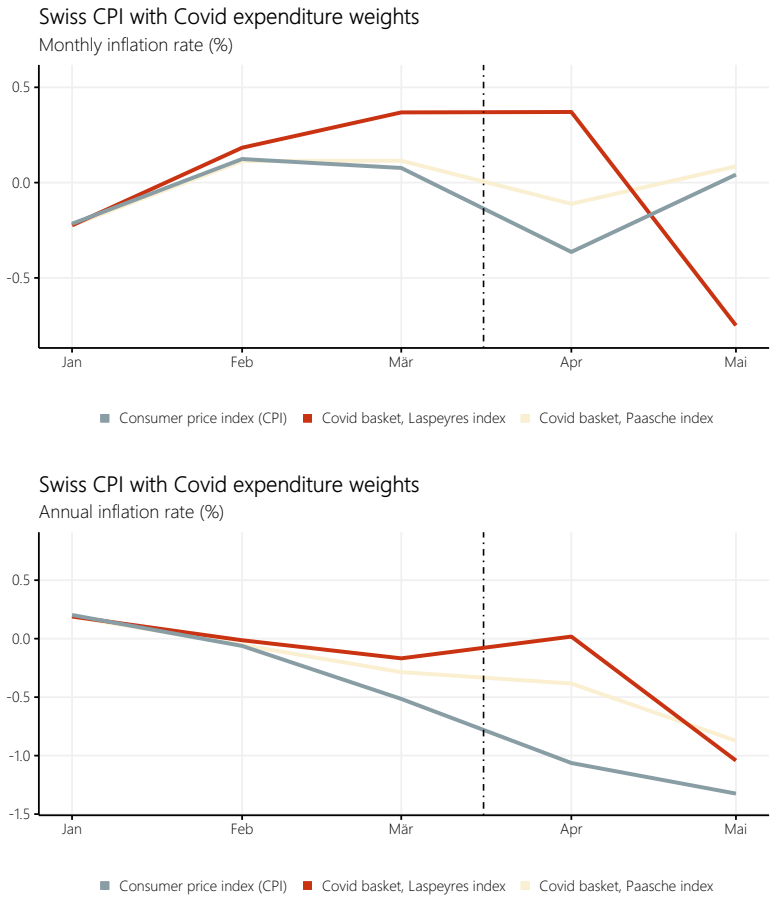

Figure B.1: Swiss consumer price index (CPI, not seasonally adjusted) and Covid price indices, calculated once as Paasche index and once as Laspeyres index. The vertical dotted line marks the declaration of the *extraordinary situation* by the Swiss Federal Council on 16 March 2020.

Covid Economics **Issue 35, 7 July 2020**

Causal impact of masks, policies, behavior on early Covid-19 pandemic in the US[1]

Victor Chernozhukov,[2] Hiroyuki Kasahara[3] and Paul Schrimpf[4]

Date submitted: 2 July 2020; Date accepted: 3 July 2020

This paper evaluates the dynamic impact of various policies adopted by US states on the growth rates of confirmed Covid-19 cases and deaths as well as social distancing behavior measured by Google Mobility Reports, where we take into consideration people's voluntarily behavioral response to new information of transmission risks. Our analysis finds that both policies and information on transmission risks are important determinants of Covid-19 cases and deaths and shows that a change in policies explains a large fraction of observed changes in social distancing behavior. Our counterfactual experiments suggest that nationally mandating face masks for employees on April 1st could have reduced the growth rate of cases and deaths by more than 10 percentage points in late April, and could have led to as much as 17 to 55 percent less deaths nationally by the end of May, which roughly translates into 17 to 55 thousand saved lives. Our estimates imply that removing non-essential business closures (while maintaining school closures, restrictions on movie theaters and restaurants) could have led to -20 to 60 percent more cases and deaths by the end of May. We also find that, without stay-at-home orders, cases would have been larger by 25 to 170 percent, which implies that 0.5 to 3.4 million more Americans could have been infected if stay-at-home orders had not been implemented. Finally, not having implemented any policies could have led to at least a 7 fold increase with an uninformative upper bound in cases (and deaths) by the end of May in the US, with considerable uncertainty over the effects of school closures, which had little cross-sectional variation.

1 We are grateful to Daron Acemoglu, V.V. Chari, Raj Chetty, Christian Hansen, Glenn Ellison, Ivan Fernandez-Val, David Green, Ido Rosen, Konstantin Sonin, James Stock, and Ivan Werning for helpful comments. We also thank Chiyoung Ahn, Joshua Catalano, Jason Chau, Samuel Gyetvay, Sev Chenyu Hou, Jordan Hutchings, and Dongxiao Zhang for excellent research assistance. All mistakes are our own.

2 Professor, Department of Economics + Center for Statistics, MIT.
3 Professor, Vancouver School of Economics.
4 Associate professor at the Vancouver School of Economics.

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

1. Introduction

Accumulating evidence suggests that various policies in the US have reduced social interactions and slowed down the growth of Covid-19 infections.[1] An important outstanding issue, however, is how much of the observed slow down in the spread is attributable to the effect of policies as opposed to a voluntarily change in people's behavior out of fear of being infected. This question is critical for evaluating the effectiveness of restrictive policies in the US relative to an alternative policy of just providing recommendations and information such as the one adopted by Sweden. More generally, understanding people's dynamic behavioral response to policies and information is indispensable for properly evaluating the effect of policies on the spread of Covid-19.

This paper quantitatively assesses the impact of various policies adopted by US states on the spread of Covid-19, such as non-essential business closure and mandatory face masks, paying particular attention to how people adjust their behavior in response to policies as well as new information on cases and deaths.

We present a conceptual framework that spells out the causal structure on how the Covid-19 spread is dynamically determined by policies and human behavior. Our approach explicitly recognizes that policies not only directly affect the spread of Covid-19 (e.g., mask requirement) but also indirectly affect its spread by changing people's behavior (e.g., stay-at-home order). It also recognizes that people react to new information on Covid-19 cases and deaths, and voluntarily adjust their behavior (e.g., voluntary social distancing and hand washing) even without any policy in place. Our casual model provides a framework to quantitatively decompose the growth of Covid-19 cases and deaths into three components: (1) direct policy effect, (2) policy effect through behavior, and (3) direct behavior effect in response to new information.[2]

Guided by the causal model, our empirical analysis examines how the weekly growth rates of confirmed Covid-19 cases and deaths are determined by (the lags of) policies and behavior using US state-level data. To examine how policies and information affect people's behavior, we also regress social distancing measures on policy and information variables. Our regression specification for case and death growths is explicitly guided by a SIR model although our causal approach does not hinge on the validity of a SIR model.

As policy variables, we consider mandatory face masks for employees in public businesses, stay-at-home orders (or shelter-in-place orders), closure of K-12 schools, closure of restaurants except take out, closure of movie theaters, and closure of non-essential businesses.

[1] See Courtemanche et al. (2020), Hsiang et al. (2020), Pei, Kandula, and Shaman (2020), Abouk and Heydari (2020), and Wright et al. (2020).

[2] The causal model is framed using the language of structural equations models and causal diagrams of econometrics (Wright (1928); Haavelmo (1944); Heckman and Vytlacil (2007); see Greenland, Pearl, and Robins (1999), Peters, Janzing, and Bernhard (2017), and Hernán and Robins (2020) for modern developments, especially in computer science and epidemiology), with natural unfolding potential outcomes representation (Rubin, 1974; Tinbergen, 1930; Neyman, 1925; Imbens and Rubin, 2015). As such it naturally converses in all languages for causal inference.

Covid Economics 35, 7 July 2020: 116-175

Our behavior variables are four mobility measures that capture the intensity of visits to "transit," "grocery," "retail," and "workplaces" from Google Mobility Reports. We take the lagged growth rate of cases and deaths and the log of lagged cases and deaths at both state-level and national-level as our measures of information on infection risks that affects people's behavior. We also consider the growth rate of tests, month dummies, and state-level characteristics (e.g., population size and total area) as confounders that have to be controlled for in order to identify the causal relationship between policy/behavior and the growth rate of cases and deaths.

Our key findings from regression analysis are as follows. We find that both policies and information on past cases and deaths are important determinants of people's social distancing behavior, where policy effects explain more than 50% of the observed decline in the four behavior variables.[3] Our estimates suggest that there are both large policy effects and large behavioral effects on the growth of cases and deaths. Except for mandatory masks, the effect of policies on cases and deaths is indirectly materialized through their impact on behavior; the effect of mandatory mask policy is direct without affecting behavior.

Using the estimated model, we evaluate the dynamic impact of the following counterfactual policies on Covid-19 cases and deaths: mandating face masks, allowing non-essential businesses to open, not implementing a stay-at-home order, and removing all policies. The counterfactual experiments show a large impact of those policies on the number of cases and deaths. They also highlight the importance of voluntary behavioral response to infection risks for evaluating the dynamic policy effects.

Figure 1 shows that nationally implementing mandatory face masks for employees in public businesses on April 1st would have reduced the growth rate of cases (top panel) and that of deaths (bottom panel) by more than 10 percentage points in late April. This leads to reductions of 25% and 35% in reported cases and deaths, respectively, by the end of May with a 90 percent confidence interval of $[10, 45]$% and $[17, 55]$%, which roughly implies that as many as 17 to 55 thousand lives could have been saved.[4] This finding is significant: given this potentially large benefit of reducing the spread of Covid-19, mandating masks is an attractive policy instrument especially because it involves relatively little economic disruption. These estimates contribute to the ongoing efforts towards designing approaches to minimize risks from reopening (Stock, 2020b).

Figure 2 illustrates how allowing non-essential businesses to remain open could have affected the growth of cases. We estimate that non-essential business closures have a small impact on growth rates, with a 90% confidence interval that includes both negative and positive effects. When this effect on growth rates is converted to a change in levels, the point estimates indicate that keeping non-essential businesses open (other than movie theaters, gyms, and keeping restaurants in the "take-out" mode) could have increased cases and

[3]The behavior accounts for the other half. This is in line with theoretical study by Gitmez, Sonin, and Wright (2020) that investigates the role of private behavior and negative external effects for individual decisions over policy compliance as well as information acquisition during pandemics.

[4]As of May 27, 2020, the US Centers for Disease Control and Prevention reports 99,031 deaths in the US.

deaths by 15% (with a 90 percent confidence interval of −20% to 60%). These estimates contribute to the ongoing efforts of evaluating various reopening approaches.

In Figure 3, we find that, without stay-at-home orders, the case growth rate would have been nearly 10 percentage points higher in late April. No stay-at-home orders could have led to 80% more cases by the start of June with a 90 precent confidence interval given by 25% to 170%. This implies that 0.5 to 3.4 million more Americans would have been infected without stay-at-home orders, providing suggestive evidence that reopening via removal of stay-at-home orders could lead to a substantial increase in cases and deaths.

In our counterfactual experiment of removing all policies, we find that the results are sensitive to whether the number of past national cases/deaths is included in a specification or not. This sensitivity arises because there is little variation across states in the timing of school closures. This makes the effect of school closures difficult to identify. In Figure 15, we show that in a specification that excludes past national cases (which allow for greater attribution of effects to school closures), the number of cases by the end of May could have increased 7-fold or more with a very large upper bound. On the other hand, as shown in Figure 16, under a specification with past national cases, our counterfactual experiment implies a 0 to 10 fold increase in cases by the end of May. This highlights the uncertainty regarding the impact of all policies versus private behavioral responses to information. Evaluation of re-opening policies needs to be aware of this uncertainty.

A growing number of other papers have examined the link between non-pharmaceutical interventions and Covid-19 cases.[5] Hsiang et al. (2020) estimate the effect of policies on the growth rate of cases using data from the United States, China, Iran, Italy, France, and South Korea. In the United States, they find that the combined effect of all policies they consider on the growth rate is −0.347 (0.061). Courtemanche et al. (2020) use US county level data to analyze the effect of interventions on case growth rates. They find that the combination of policies they study reduced growth rates by 9.1 percentage points 16-20 days after implementation, out of which 5.9 percentage points are attributed to shelter in place orders. Both Hsiang et al. (2020) and Courtemanche et al. (2020) adopted a reduced-form approach to estimate the total policy effect on case growth without using any social distancing behavior measures. In contrast, our study highlights the role of behavioral response to policies and information.

Existing evidence for the impact of social distancing policies on behavior in the US is mixed. Abouk and Heydari (2020) employ a difference-in-differences methodology to find that statewide stay-at-home orders have strong causal impacts on reducing social interactions. In contrast, using data from Google Mobility Reports, Maloney and Taskin (2020) find that the increase in social distancing is largely voluntary and driven by information.[6]

[5]We refer the reader to Avery et al. (2020) for a comprehensive review of a larger body of work researching Covid-19; here we focus on few quintessential comparisons on our work with other works that we are aware of.

[6]Specifically, they find that of the 60 percentage point drop in workplace intensity, 40 percentage points can be explained by changes in information as proxied by case numbers, while roughly 8 percentage points can be explained by policy changes.

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

FIGURE 1. Effect of nationally mandating masks for employees on April 1st in the US

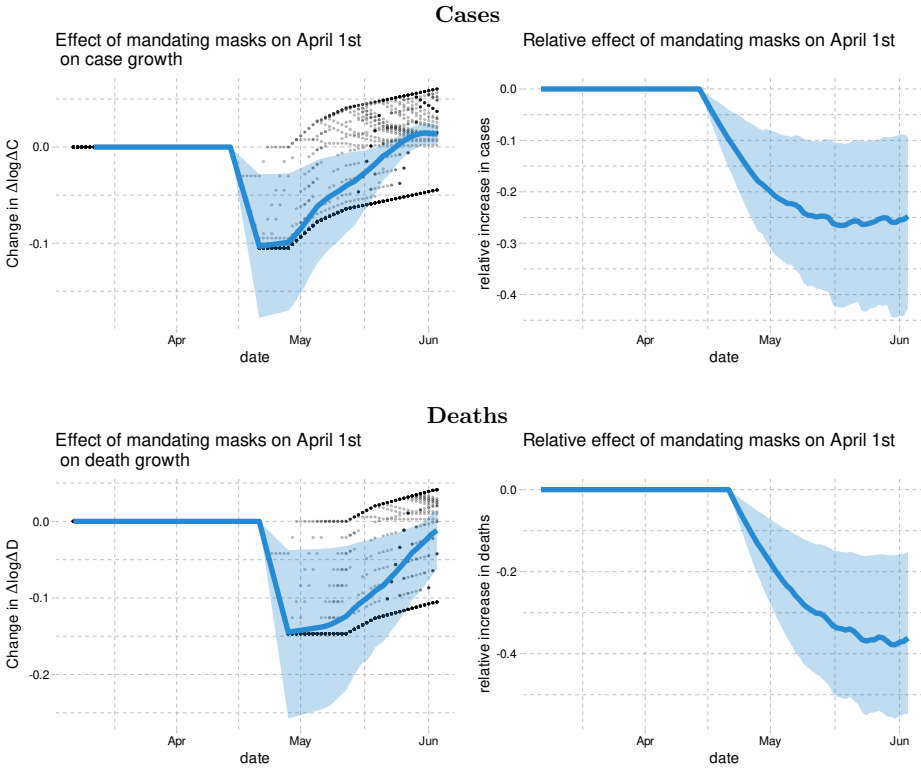

Another study by Gupta et al. (2020) also found little evidence that stay-at-home mandates induced distancing by using mobility measures from PlaceIQ and SafeGraph. Using data from SafeGraph, Andersen (2020) show that there has been substantial voluntary social distancing but also provide evidence that mandatory measures such as stay-at-home orders have been effective at reducing the frequency of visits outside of one's home.

Pei, Kandula, and Shaman (2020) use county-level observations of reported infections and deaths in conjunction with mobility data from SafeGraph to estimate how effective reproductive numbers in major metropolitan areas change over time. They conduct simulation of implementing all policies 1-2 weeks earlier and found that it would have resulted in reducing the number of cases and deaths by more than half. However, their study does not explicitly analyze how policies are related to the effective reproduction numbers.

Covid Economics 35, 7 July 2020: 116-175

FIGURE 2. Effect of leaving non-essential businesses open on cases in the US

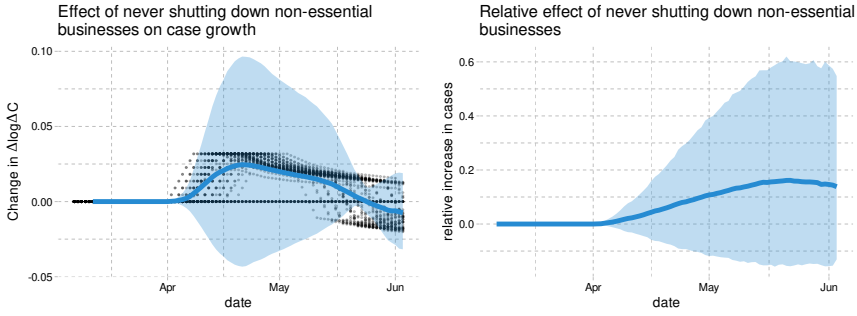

FIGURE 3. Effect of not implementing stay-at-home order on cases in the US

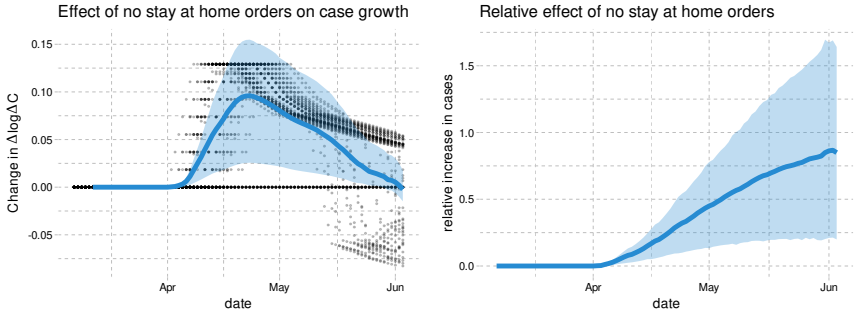

Epidemiologists use model simulations to predict how cases and deaths evolve for the purpose of policy recommendation. As reviewed by Avery et al. (2020), there exist substantial uncertainty about the values of key epidimiological parameters (see also Atkeson, 2020a; Stock, 2020a). Simulations are often done under strong assumptions about the impact of social distancing policies without connecting to the relevant data (e.g., Ferguson et al., 2020). Furthermore, simulated models do not take into account that people may limit their contact with other people in response to higher transmission risks.[7] When such a voluntary behavioral response is ignored, simulations would produce exponential spread of disease and would over-predict cases and deaths. Our counterfactual experiments illustrate the importance of this voluntary behavioral change.

[7]See Atkeson (2020b) and Stock (2020a) for the implications of the SIR model for Covid-19 in the US. Fernández-Villaverde and Jones (2020) estimate a SIRD model in which time-varying reproduction numbers depend on the daily deaths to capture feedback from daily deaths to future behavior and infections.

Covid Economics 35, 7 July 2020: 116-175

Whether wearing masks in public place should be mandatory or not has been one of the most contested policy issues with health authorities of different countries providing contradiction recommendations. Reviewing evidence, Greenhalgh et al. (2020) recognize that there is no randomized controlled trial evidence for the effectiveness of face masks, but they state "indirect evidence exists to support the argument for the public wearing masks in the Covid-19 pandemic."[8] Howard et al. (2020) also review available medical evidence and conclude that "mask wearing reduces the transmissibility per contact by reducing transmission of infected droplets in both laboratory and clinical contexts." The laboratory findings in Hou et al. (2020) suggest that the nasal cavity may be the initial site of infection followed by aspiration to the lung, supporting the argument "for the widespread use of masks to prevent aersol, large droplet, and/or mechanical exposure to the nasal passages."

Given the lack of experimental evidence on the effect of masks, conducting observational studies is useful and important. To the best of our knowledge, our paper is the first empirical study that shows the effectiveness of mask mandates on reducing the spread of Covid-19 by analyzing the US state-level data. This finding corroborates and is complementary to the medical observational evidence in Howard et al. (2020). Analyzing mitigation measures in New York, Wuhan, and Italy, Zhang et al. (2020b) conclude that mandatory face coverings substantially reduced infections. Abaluck et al. (2020) find that the growth rates of cases and of deaths in countries with pre-existing norms that sick people should wear masks are lower by 8 to 10% than those rates in countries with no pre-existing mask norms. Our finding is also corroborated by a completely different causal methodology based on synthetic control using German data in Mitze et al. (2020).[9]

Our empirical results contribute to informing the economic-epidemiological models that combine economic models with variants of SIR models to evaluate the efficiency of various economic policies aimed at gradual "reopening" of various sectors of economy.[10] For example, the estimated effects of masks, stay-home mandates, and various other policies on behavior, and of behavior on infection can serve as useful inputs and validation checks in the calibrated macro, sectoral, and micro models (see, e.g., Alvarez, Argente, and Lippi (2020); Baqaee et al. (2020); Fernández-Villaverde and Jones (2020); Acemoglu et al. (2020); Keppo et al. (2020); McAdams (2020) and references therein). Furthermore, the causal framework developed in this paper could be applicable, with appropriate extensions, to the impact of policies on economic outcomes replacing health outcomes (see, e.g., Chetty et al. (2020); Coibion, Gorodnichenko, and Weber (2020)).

[8]The virus remains viable in the air for several hours, for which surgical masks may be effective. Also, a substantial fraction of individual who are infected become infectious before showing symptom onset.

[9]Our study was first released in ArXiv on May 28, 2020 whereas Mitze et al. (2020) was released at SSRN on June 8, 2020.

[10]Adda (2016) analyzes the effect of policies reducing interpersonal contacts such as school closures or the closure of public transportation networks on the spread of influenza, gastroenteritis, and chickenpox using high frequency data from France.

Covid Economics 35, 7 July 2020: 116-175

2. The Causal Model for the Effect of Policies, Behavior, and Information on Growth of Infection

2.1. **The Causal Model and Its Structural Equation Form.** We introduce our approach through the Wright-style causal diagram shown in Figure 4. The diagram describes how policies, behavior, and information interact together:

- The *forward* health outcome, $Y_{i,t+\ell}$, is determined last, after all other variables have been determined;
- The adopted policies, P_{it}, affect health outcome $Y_{i,t+\ell}$ either directly, or indirectly by altering human behavior B_{it};
- Information variables, I_{it}, such as lagged values of outcomes can affect human behavior and policies, as well as outcomes;
- The confounding factors W_{it}, which vary across states and time, affect all other variables.

The index i denotes observational unit, the state, and t and $t+\ell$ denotes the time, where ℓ represents the time lag between infection and case confirmation or death.

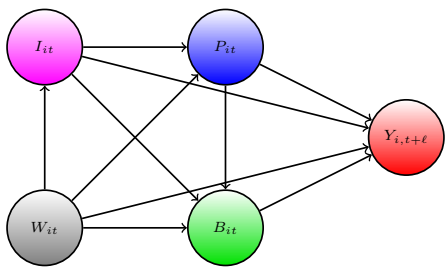

Figure 4. S. & P. Wright type causal path diagram for our model.

Our main outcomes of interest are the growth rates in Covid-19 cases and deaths, behavioral variables include proportion of time spent in transit or shopping and others, policy variables include stay-at-home orders and school and business closures, and the information variables include lagged values of outcome. We provide a detailed description of these variables and their timing in the next section.

The causal structure allows for the effect of the policy to be either direct or indirect – through-behavior or through dynamics; and all of these effects are not mutually exclusive. The structure also allows for changes in behavior to be brought by change in policies and information. These are all realistic properties that we expect from the contextual knowledge of the problem. Policies such as closures of schools, non-essential business, and restaurants, alter and constrain behavior in strong ways. In contrast, policies such as mandating employees to wear masks can potentially affect the Covid-19 transmission directly. The

Covid Economics 35, 7 July 2020: 116-175

information variables, such as recent growth in the number of cases, can cause people to spend more time at home, regardless of adopted state policies; these changes in behavior in turn affect the transmission of Covid-19.

The causal ordering induced by this directed acyclical graph is determined by the following timing sequence:

(1) information and confounders get determined at t,
(2) policies are set in place, given information and confounders at t;
(3) behavior is realized, given policies, information, and confounders at t;
(4) outcomes get realized at $t+\ell$ given policies, behavior, information, and confounders.

The model also allows for direct dynamic effects of information variables on the outcome through autoregressive structures that capture persistence in growth patterns. As highlighted below, realized outcomes may become new information for future periods, inducing dynamics over multiple periods.

Our quantitative model for causal structure in Figure 4 is given by the following econometric structural equation model:

$$\begin{aligned} Y_{i,t+\ell}(b,p,\iota) &:= \alpha' b + \pi' p + \mu' \iota + \delta_Y' W_{it} + \varepsilon_{it}^y, \\ B_{it}(p,\iota) &:= \beta' p + \gamma' \iota + \delta_B' W_{it} + \varepsilon_{it}^b, \end{aligned} \tag{SEM}$$

which is a collection of functional relations with stochastic shocks, decomposed into observable part $\delta' W$ and unobservable part ε. The terms ε_{it}^y and ε_{it}^b are the centered stochastic shocks that obey the orthogonality restrictions posed below.

The policies can be modeled via a linear form as well,

$$P_{it}(\iota) := \eta' \iota + \delta_P' W_{it} + \varepsilon_{it}^p, \tag{P}$$

although linearity is not critical.[11]

The orthogonality restrictions on the stochastic components are as follows:

$$\begin{aligned} \varepsilon_{it}^y &\perp (\varepsilon_{it}^b, \varepsilon_{it}^p, W_{it}, I_{it}), \\ \varepsilon_{it}^b &\perp (\varepsilon_{it}^p, W_{it}, I_{it}), \\ \varepsilon_{it}^p &\perp (W_{it}, I_{it}), \end{aligned} \tag{O}$$

[11] Under some additional independence conditions, this can be replaced by an arbitrary non-additive function $P_{it}(\iota) = p(\iota, W_{it}, \varepsilon_{it}^p)$, such that the unconfoundedness condition stated in the next footnote holds.

where we say that $V \perp U$ if $EVU = 0$. This is a standard way of representing restrictions on errors in structural equation modeling.[12][13]

The observed variables are generated by setting $\iota = I_{it}$ and propagating the system from the last equation to the first:

$$\begin{aligned} Y_{i,t+\ell} &:= Y_{i,t+\ell}(B_{it}, P_{it}, I_{it}), \\ B_{it} &:= B_{it}(P_{it}, I_{it}), \\ P_{it} &:= P_{it}(I_{it}). \end{aligned}$$

The system above together with orthogonality restrictions (O) implies the following collection of stochastic equations for realized variables:

$$Y_{i,t+\ell} = \alpha' B_{it} + \pi' P_{it} + \mu' I_{it} + \delta_Y' W_{it} + \varepsilon^y_{it}, \qquad \varepsilon^y_{it} \perp B_{it}, P_{it}, I_{it}, W_{it} \qquad \text{(BPI→Y)}$$

$$B_{it} = \beta' P_{it} + \gamma' I_{it} + \delta_B' W_{it} + \varepsilon^b_{it}, \qquad \varepsilon^b_{it} \perp P_{it}, I_{it}, W_{it} \qquad \text{(PI→B)}$$

$$P_{it} = \eta' I_{it} + \delta_P' W_{it} + \varepsilon^p_{it}, \qquad \varepsilon^p_{it} \perp I_{it}, W_{it} \qquad \text{(I→P)}$$

and

$$Y_{i,t+\ell} = (\alpha'\beta' + \pi') P_{it} + (\alpha'\gamma' + \mu') I_{it} + \bar\delta' W_{it} + \bar\varepsilon_{it}, \quad \bar\varepsilon_{it} \perp P_{it}, I_{it}, W_{it}. \qquad \text{(PI→Y)}$$

These equations form the basis of our empirical analysis.

As discussed below, the information variable includes case growth. Therefore, an orthogonality restriction $\varepsilon^y_{it} \perp P_{it}$ holds if the government does not have knowledge on future case growth beyond what is predicted by today's case growth, policies, behavior, and confounders; even when the government has some knowledge on ε^y_{it}, the orthogonality restriction may hold if there is a time lag for the government to implement its policies based on ε^y_{it}.

The orthogonality condition in (PI→Y) is weaker than the orthogonality conditions in (BPI→Y)-(PI→B) in that the former is implied by the latter but not vice versa. The system over-identifies the regression coefficients because $(\alpha'\beta' + \pi')$ and $(\alpha'\gamma' + \mu')$ in (PI→Y) can be also identified from α', π', μ', β', and γ' in (BPI→Y)-(PI→B). Comparing the estimates of $(\alpha'\beta' + \pi')$ and $(\alpha'\gamma' + \mu')$ from (PI→Y) with those implied by the estimates of α', π', μ', β', and γ' from (BPI→Y)-(PI→B) provides a useful specification test.

Identification and Parameter Estimation. The orthogonality equations imply that these are all projection equations, and the parameters of the SEM are identified by the parameters of these regression equation, provided the latter are identified by sufficient joint variation of these variables across states and time.

[12] An alternative useful starting point is to impose the Rubin-Rosenbaum type unconfoudedness condition:

$$Y_{i,t+\ell}(\cdot,\cdot,\cdot) \perp\!\!\!\perp (P_{it}, B_{it}, I_{it}) \mid W_{it}, \; B_{it}(\cdot,\cdot) \perp\!\!\!\perp (P_{it}, I_{it}) \mid W_{it}, \; P_{it}(\cdot) \perp\!\!\!\perp I_{it} \mid W_{it},$$

which imply, with treating stochastic errors as independent additive components, the orthogonal conditions stated above.

[13] The structural equations of this form are connected to triangular structural equation models, appearing in microeconometrics and macroeconometrics (SVARs), going back to the work of Strotz and Wold (1960).

.

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

The last point can be stated formally as follows. Consider the previous system of equations, after partialling out the confounders:

$$\begin{aligned} \tilde{Y}_{i,t+\ell} &= \alpha' \tilde{B}_{it} + \pi' \tilde{P}_{it} + \mu' \tilde{I}_{it} + \varepsilon^y_{it}, & \varepsilon^y_{it} &\perp \tilde{B}_{it}, \tilde{P}_{it}, \tilde{I}_{it}, \\ \tilde{B}_{it} &= \beta' \tilde{P}_{it} + \gamma' \tilde{I}_{it} + \varepsilon^b_{it}, & \varepsilon^b_{it} &\perp \tilde{P}_{it}, \tilde{I}_{it}, \\ \tilde{P}_{it} &= \eta' \tilde{I}_{it} + \varepsilon^p_{it}, & \varepsilon^p_{it} &\perp \tilde{I}_{it} \end{aligned} \tag{1}$$

where $\tilde{V}_{it} = V_{it} - W_{it}' \mathrm{E}[W_{it} W_{it}']^{-} \mathrm{E}[W_{it} V_{it}]$ denotes the residual after removing the orthogonal projection of V_{it} on W_{it}. The residualization is a linear operator, implying that (1) follows immediately from the above. The parameters of (1) are identified as projection coefficients in these equations, provided that residualized vectors appearing in each of the equations have non-singular variance, that is

$$\mathrm{Var}(\tilde{P}_{it}', \tilde{B}_{it}', \tilde{I}_{it}') > 0, \ \mathrm{Var}(\tilde{P}_{it}', \tilde{I}_{it}') > 0, \quad \text{and } \mathrm{Var}(\tilde{I}_{it}') > 0. \tag{2}$$

Our main estimation method is the standard correlated random effects estimator, where the random effects are parameterized as functions of observable characteristic, W_{it}, which include both state-level and time random effects. The state-level random effects are modeled as a function of state level characteristics, and the time random effects are modeled by including month dummies and their interactions with state level characteristics. The stochastic shocks $\{\varepsilon_{it}\}_{t=1}^{T}$ are treated as independent across states i and can be arbitrarily dependent across time t within a state.

A secondary estimation method is the fixed effects estimator, where W_{it} includes latent (unobserved) state level effects W_i and and time level effects W_t, which must be estimated from the data. This approach is much more demanding of the data and relies on long cross-sectional and time histories. When histories are relatively short, large biases emerge and they need to be removed using debiasing methods. In our context, debiasing changes the estimates substantially, often changing the sign of coefficients.[14] However, we find the debiased fixed effect estimates are qualitatively and quantitatively similar to the correlated random effects estimates. Given this finding, we chose to focus on the latter, as it is a more standard and familiar method, and report the former estimates in the supplementary materials for this paper.[15]

2.2. **Information Structures and Induced Dynamics.** We consider three examples of information structures: Information variable is a function of time:

$$I_{it} = g(t);$$

Information variable is lagged value of outcome:

$$I_{it} = Y_{it};$$

[14]This is a pre-cautionary message that may be useful for other researchers using fixed effects estimators in the context of Covid-19 analysis. We recommend using debiased fixed effects estimators, see e.g., Chen, Chernozhukov, and Fernández-Val (2019) for expository treatment.

[15]The similarity of the debiased fixed effects and correlated random effects served as a useful specification check. Moreover, using the fixed effects estimators only yielded minor gains in predictive performances, as judging by the adjusted R^2's, providing another useful specification check.

and finally:

(I) Information variables include time, lagged and integrated values of outcome:

$$I_{it} = \left(g(t), Y_{it}, \sum_{m=0}^{t/\ell} Y_{i,t-\ell m} \right)',$$

with the convention that $Y_{it} = 0$ for $t \leq 0$.

The first information structure captures the basic idea that, as individuals discover more information about covid over time, they adapt to safer modes of behavior (stay-at-home, wear masks, wash hands). Under this structure, information is common across states and exogenously evolves over time, independent of the number of cases. The second structure arises from considering autoregressive components and captures people's behavioral response to information on cases in the state they reside. Specifically, we model persistence in growth rates, $Y_{i,t+\ell}$, through an AR(1) model, which leads to $I_{it} = Y_{it}$. This provides useful *local*, state-specific, information about the forward growth rate and people may adjust their behavior to safer modes when they see a high value. We model this adjustment via the term $\gamma' I_t$ in the behavior equation. The third information structure is the combination of the first two structures plus an additional term representing the log of the total number of new cases in the state. We use this information structure in our empirical specification. In this structure, people respond to both global information, captured by a function of time such as month dummies, and local information sources, captured by the local growth rate and the total number of cases. The last element of the information set can be thought of as a local stochastic trend in cases.

All of these examples fold into a specification of the form:

$$I_{it} := I_{it}(I_{i,t-\ell}, Y_{it}, t), \quad t = 1, ..., T, \tag{I}$$

with the initialization $I_{i0} = 0$ and $Y_{i0} = 0$.[16]

With any structure of this form, realized outcomes may become new information for future periods, inducing a dynamical system over multiple periods. We show the resulting dynamical system in a causal diagram of Figure 5. Specification of this system is useful for studying delayed effects of policies and behaviors and in considering the counterfactual policy analysis.

2.3. **Outcome and Key Confounders via SIR model.** Letting C_{it} denote the cumulative number of confirmed cases in state i at time t, our outcome

$$Y_{it} = \Delta \log(\Delta C_{it}) := \log(\Delta C_{it}) - \log(\Delta C_{i,t-7}) \tag{3}$$

[16]This initialization is appropriate in our context for analyzing pandemics from the very beginning, but other initializations could be appropriate in other contexts. The lagged values of behavior variable may be also included in the information set.

Covid Economics 35, 7 July 2020: 116-175

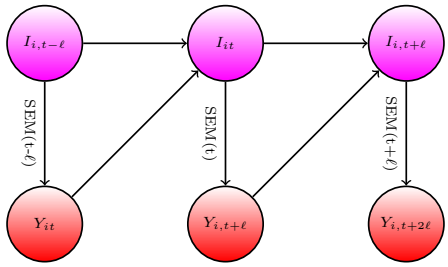

FIGURE 5. Dynamic System Induced by Information Structure and SEM

approximates the weekly growth rate in new cases from $t-7$ to t.[17] Here Δ denotes the differencing operator over 7 days from t to $t-7$, so that $\Delta C_{it} := C_{it} - C_{i,t-7}$ is the number of new confirmed cases in the past 7 days.

We chose this metric as this is the key metric for policy makers deciding when to relax Covid mitigation policies. The U.S. government's guidelines for state reopening recommend that states display a "downward trajectory of documented cases within a 14-day period" (White House, 2020). A negative value of Y_{it} is an indication of meeting this criteria for reopening. By focusing on weekly cases rather than daily cases, we smooth idiosyncratic daily fluctuations as well as periodic fluctuations associated with days of the week.

Our measurement equation for estimating equations (BPI→Y) and (PI→Y) will take the form:

$$\Delta \log(\Delta C_{it}) = X'_{i,t-14}\theta - \gamma + \delta_T \Delta \log(T_{it}) + \epsilon_{it}, \tag{M-C}$$

where i is state, t is day, C_{it} is cumulative confirmed cases, T_{it} is the number of tests over 7 days, Δ is a 7-days differencing operator, ϵ_{it} is an unobserved error term. $X_{i,t-14}$ collects other behavioral, policy, and confounding variables, depending on whether we estimate (BPI→Y) or (PI→Y), where the lag of 14 days captures the time lag between infection and confirmed case (see the Appendix A.6). Here

$$\Delta \log(T_{it}) := \log(T_{it}) - \log(T_{i,t-7})$$

is the key confounding variable, derived from considering the SIR model below. We describe other confounders in the empirical section.

Our main estimating equation (M-C) is motivated by a variant of SIR model, where we add confirmed cases and infection detection via testing. Let S, $\mathcal{I}$, R, and D denote the number of susceptible, infected, recovered, and dead individuals in a given state. Each of

[17] We may show that $\log(\Delta C_{it}) - \log(\Delta C_{i,t-7})$ approximates the average growth rate of cases from $t-7$ to t.

Covid Economics 35, 7 July 2020: 116-175

these variables are a function of time. We model them as evolving as

$$\dot{S}(t) = -\frac{S(t)}{N}\beta(t)\mathcal{I}(t) \tag{4}$$

$$\dot{\mathcal{I}}(t) = \frac{S(t)}{N}\beta(t)\mathcal{I}(t) - \gamma\mathcal{I}(t) \tag{5}$$

$$\dot{R}(t) = (1-\kappa)\gamma\mathcal{I}(t) \tag{6}$$

$$\dot{D}(t) = \kappa\gamma\mathcal{I}(t) \tag{7}$$

where N is the population, $\beta(t)$ is the rate of infection spread, γ is the rate of recovery or death, and κ is the probability of death conditional on infection.

Confirmed cases, $C(t)$, evolve as

$$\dot{C}(t) = \tau(t)\mathcal{I}(t), \tag{8}$$

where $\tau(t)$ is the rate that infections are detected.

Our goal is to examine how the rate of infection $\beta(t)$ varies with observed policies and measures of social distancing behavior. A key challenge is that we only observed $C(t)$ and $D(t)$, but not $\mathcal{I}(t)$. The unobserved $\mathcal{I}(t)$ can be eliminated by differentiating (8) and using (5) as

$$\frac{\ddot{C}(t)}{\dot{C}(t)} = \frac{S(t)}{N}\beta(t) - \gamma + \frac{\dot{\tau}(t)}{\tau(t)}. \tag{9}$$

We consider a discrete-time analogue of equation (9) to motivate our empirical specification by relating the detection rate $\tau(t)$ to the number of tests T_{it} while specifying $\frac{S(t)}{N}\beta(t)$ as a linear function of variables $X_{i,t-14}$. This results in

$$\underbrace{\Delta\log(\Delta C_{it})}_{\frac{\ddot{C}(t)}{\dot{C}(t)}} = \underbrace{X'_{i,t-14}\theta + \epsilon_{it}}_{\frac{S(t)}{N}\beta(t)-\gamma} + \underbrace{\delta_T \Delta\log(T)_{it}}_{\frac{\dot{\tau}(t)}{\tau(t)}}$$

which is equation (M-C), where $X_{i,t-14}$ captures a vector of variables related to $\beta(t)$.

> Structural Interpretation. Early in the pandemic, when the number of susceptibles is approximately the same as the entire population, i.e. $S_{it}/N_{it} \approx 1$, the component $X'_{i,t-14}\theta$ is the projection of infection rate $\beta_i(t)$ on $X_{i,t-14}$ (policy, behavioral, information, and confounders other than testing rate), provided the stochastic component ϵ_{it} is orthogonal to $X_{i,t-14}$ and the testing variables:
>
> $$\beta_i(t)S_{it}/N_{it} - \gamma = X'_{i,t-14}\theta + \epsilon_{it}, \quad \epsilon_{it} \perp X_{i,t-14}.$$

2.4. **Growth Rate in Deaths as Outcome.** By differentiating (7) and (8) with respect to t and using (9), we obtain

$$\frac{\ddot{D}(t)}{\dot{D}(t)} = \frac{\ddot{C}(t)}{\dot{C}(t)} - \frac{\dot{\tau}(t)}{\tau(t)} = \frac{S(t)}{N}\beta(t) - \gamma. \tag{10}$$

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

Our measurement equation for the growth rate of deaths is based on equation (10) but account for a 21 day lag between infection and death as

$$\Delta \log(\Delta D_{it}) = X'_{i,t-21}\theta + \epsilon_{it}, \tag{M-D}$$

where

$$\Delta \log(\Delta D_{it}) := \log(\Delta D_{it}) - \log(\Delta D_{i,t-7}) \tag{11}$$

approximates the weekly growth rate in deaths from $t-7$ to t in state i.

3. Decomposition and Counterfactual Policy Analysis

3.1. **Total Change Decomposition.** Given the SEM formulation above, we can carry out the following decomposition analysis, after removing the effects of confounders. For example, we can decompose the total change $\mathrm{E}\tilde{Y}_{i,t+\ell} - \mathrm{E}\tilde{Y}_{io}$ in the expected outcome, measured at two different time points $t+\ell$ and o into the sum of three components:

$$\begin{aligned}
\underbrace{\mathrm{E}\tilde{Y}_{i,t+\ell} - \mathrm{E}\tilde{Y}_{io}}_{\text{Total Change}} = & \underbrace{\alpha'\beta'\left(\mathrm{E}\tilde{P}_{it} - \mathrm{E}\tilde{P}_{io}\right)}_{\text{Policy Effect via Behavior}} \\
& + \underbrace{\pi'\left(\mathrm{E}\tilde{P}_{it} - \mathrm{E}\tilde{P}_{io}\right)}_{\text{Direct Policy Effect}} \\
& + \underbrace{\alpha'\gamma'\left(\mathrm{E}\tilde{I}_{it} - \mathrm{E}\tilde{I}_{io}\right) + \mu'\left(\mathrm{E}\tilde{I}_{it} - \mathrm{E}\tilde{I}_{io}\right)}_{\text{Dynamic Effect}} \\
=: & \ \mathrm{PEB}_t + \mathrm{PED}_t + \mathrm{DynE}_t,
\end{aligned} \tag{12}$$

where the first two components capture the immediate effect and the third represents the delayed or dynamic effect.

In the three examples of information structure given earlier, we have the following forms for the dynamic effect: for the trend model,

$$\mathrm{DynE}_t = (\gamma\alpha + \mu)\Delta g_t, \quad \Delta g_t = (g(t) - g(t-\ell))$$

and for the lag model,

$$\mathrm{DynE}_t = \sum_{m=1}^{t/\ell} (\gamma\alpha + \mu)^m \left(\mathrm{PEB}_{t-m\ell} + \mathrm{PED}_{t-m\ell}\right),$$

interpreting t/ℓ as $\lfloor t/\ell \rfloor$. For the general model we use, the dynamic effect is

$$\begin{aligned}
\text{(I)} \quad \mathrm{DynE}_t = & \sum_{m=0}^{t/\ell} \left(((\gamma\alpha)_2 + \mu_2 + (\gamma\alpha)_3 + \mu_3)^m \quad ((\gamma\alpha)_1 + \mu_1)\Delta g_t \right. \\
& + \sum_{m=1}^{t/\ell} ((\gamma\alpha)_2 + \mu_2 + (\gamma\alpha)_3 + \mu_3)^m \left(\mathrm{PEB}_{t-m\ell} + \mathrm{DPE}_{t-m\ell}\right) \\
& \left. + \sum_{m=1}^{t/\ell-1} ((\gamma\alpha)_3 + \mu_3)^m \left(\mathrm{PEB}_{t-(m+1)\ell} + \mathrm{DPE}_{t-(m+1)\ell}\right) \right.
\end{aligned}$$

The effects can be decomposed into (a) delayed policy effects via behavior by summing terms containing PEB, (b) delayed policy effects via direct impact by summing terms containing DPE, (c) pure behavior effects, and (d) pure dynamic feedback effects.

Covid Economics 35, 7 July 2020: 116-175

3.2. **Counterfactuals.** We also consider simple counterfactual exercises, where we examine the effects of setting a sequence of counterfactual policies for each state:

$$\{P^\star_{it}\}_{t=1}^T, \quad i = 1, \dots N.$$

We assume that the SEM remains invariant, except of course for the policy equation.[18] Given the policies, we propagate the dynamic equations:

$$\begin{aligned} Y^\star_{i,t+\ell} &:= Y_{i,t+\ell}(B^\star_{it}, P^\star_{it}, I^\star_{it}), \\ B^\star_{it} &:= B_{it}(P^\star_{it}, I^\star_{it}), \\ I^\star_{it} &:= I_{it}(I^\star_{i,t-\ell}, Y^\star_{it}, t), \end{aligned} \tag{CEF-SEM}$$

with the initialization $I^\star_{i0} = 0$, $Y^\star_{i0} = 0$, $B^\star_{i0} = 0$, $P^\star_{i0} = 0$. In stating this counterfactual system of equations, we make the following invariance assumption

> Invariance Assumption. The equations of (CF-SEM) remain exactly of the same form as in the (SEM) and (I). That is, under the policy intervention $\{P^\star_{it}\}$, parameters and stochastic shocks in (SEM) and (I) remain the same as under the original policy intervention $\{P_{it}\}$.

Let $\mathcal{P}Y^\star_{i,t+\ell}$ and $\mathcal{P}Y_{i,t+\ell}$ denote the predicted values produced by working with the counterfactual system (CEF-SEM) and the factual system (SEM):

$$\begin{aligned} \mathcal{P}Y^\star_{i,t+\ell} &= (\alpha'\beta' + \pi')P^\star_{it} + (\alpha'\gamma' + \mu')I^\star_{it} + \bar{\delta}'W_{it}, \\ \mathcal{P}Y_{i,t+\ell} &= (\alpha'\beta' + \pi')P_{it} + (\alpha'\gamma' + \mu')I_{it} + \bar{\delta}'W_{it}. \end{aligned}$$

In generating these predictions, we make the assumption of invariance stated above.

Then we can write the difference into the sum of three components:

$$\begin{aligned} \underbrace{\mathcal{P}Y^\star_{i,t+\ell} - \mathcal{P}Y_{i,t+\ell}}_{\text{Predicted CF Change}} = & \underbrace{\alpha'\beta'(P^\star_{it} - P_{it})}_{\text{CF Policy Effect via Behavior}} + \underbrace{\pi'(P^\star_{it} - P_{it})}_{\text{CF Direct Effect}} \\ & + \underbrace{\alpha'\gamma'(I^\star_{it} - I_{it}) + \mu'(I^\star_{it} - I_{it})}_{\text{CF Dynamic Effect}} \\ =: & \ \mathrm{PEB}^\star_{it} + \mathrm{PED}^\star_{it} + \mathrm{DynE}^\star_{it}. \end{aligned} \tag{13}$$

Similarly to what we had before, the counterfactual dynamic effects take the form:

$$\begin{aligned} \text{(I)} \quad \mathrm{DynE}^\star_{it} = & \sum_{m=1}^{t/\ell} ((\gamma\alpha)_2 + \mu_2 + (\gamma\alpha)_3 + \mu_3)^m \left(\mathrm{PEB}^\star_{i,t-m\ell} + \mathrm{DPE}^\star_{i,t-m\ell}\right) \\ & + \sum_{m=1}^{t/\ell-1} ((\gamma\alpha)_3 + \mu_3)^m \left(\mathrm{PEB}^\star_{i,t-(m+1)\ell} + \mathrm{DPE}^\star_{i,t-(m+1)\ell}\right), \end{aligned}$$

interpreting t/ℓ as $\lfloor t/\ell \rfloor$. The effects can be decomposed into (a) delayed policy effects via behavior by summing terms containing PEB, (b) delayed policy effects via direct impact by

[18] It is possible to consider counterfactual exercises in which policy responds to information through the policy equation if we are interested in endogenous policy responses to information. Although this is beyond the scope of the current paper, counterfactual experiments with endogenous government policy would be important, for example, to understand the issues related to the lagged response of government policies to higher infection rates due to incomplete information.

summing terms containing DPE, (c) pure behavior effects, and (d) pure dynamic feedback effects.

4. Empirical Analysis

4.1. **Data.** Our baseline measures for daily Covid-19 cases and deaths are from The New York Times (NYT). When there are missing values in NYT, we use reported cases and deaths from JHU CSSE, and then the Covid Tracking Project. The number of tests for each state is from Covid Tracking Project. As shown in Figure 21 in the appendix, there was a rapid increase in testing in the second half of March and then the number of tests increased very slowly in each state in April.

We use the database on US state policies created by Raifman et al. (2020). In our analysis, we focus on 6 policies: stay-at-home, closed nonessential businesses, closed K-12 schools, closed restaurants except takeout, closed movie theaters, and mandate face mask by employees in public facing businesses. We believe that the first four of these policies are the most widespread and important. Closed movie theaters is included because it captures common bans on gatherings of more than a handful of people. We also include mandatory face mask use by employees because its effectiveness on slowing down Covid-19 spread is a controversial policy issue (Howard et al., 2020; Greenhalgh et al., 2020; Zhang et al., 2020b). Table 1 provides summary statistics, where N is the number of states that have ever implemented the policy. We also obtain information on state-level covariates from Raifman et al. (2020), which include population size, total area, unemployment rate, poverty rate, and a percentage of people who are subject to illness.These confounders are motivated by Wheaton and Thompson (2020) who finds that case growth is associated with residential density and per capita income.

	N	Min	Median	Max
Date closed K 12 schools	51	2020-03-13	2020-03-17	2020-04-03
Stay at home shelter in place	42	2020-03-19	2020-03-28	2020-04-07
Closed movie theaters	49	2020-03-16	2020-03-21	2020-04-06
Closed restaurants except take out	48	2020-03-15	2020-03-17	2020-04-03
Closed non essential businesses	43	2020-03-19	2020-03-25	2020-04-06
Mandate face mask use by employees	39	2020-04-03	2020-05-01	2020-06-01

Table 1. State Policies

We obtain social distancing behavior measures from"Google COVID-19 Community Mobility Reports" (LLC, 2020). The dataset provides six measures of "mobility trends" that report a percentage change in visits and length of stay at different places relative to a baseline computed by their median values of the same day of the week from January 3

Covid Economics 35, 7 July 2020: 116-175

FIGURE 6. The Evolution of Google Mobility Measures: Transit stations and Workplaces

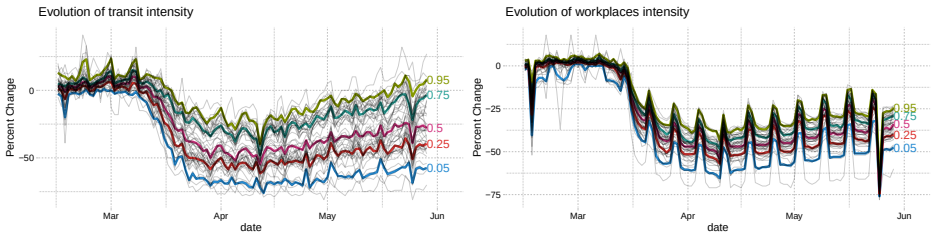

This figure shows the evolution of "Transit stations" and "Workplaces" of Google Mobility Reports. Thin gray lines are the value in each state and date. Thicker colored lines are quantiles of the variables conditional on date.

to February 6, 2020. Our analysis focuses on the following four measures: "Grocery & pharmacy," "Transit stations," "Retail & recreation," and "Workplaces."[19]

Figure 6 shows the evolution of "Transit stations" and "Workplaces," where thin lines are the value in each state and date while thicker colored lines are quantiles conditional on date. The figures illustrate a sharp decline in people's movements starting from mid-March as well as differences in their evolutions across states. They also reveal periodic fluctuations associated with days of the week, which motivates our use of weekly measures.

In our empirical analysis, we use weekly measures for cases, deaths, and tests by summing up their daily measures from day t to $t-6$. We focus on weekly cases and deaths because daily new cases and deaths are affected by the timing of reporting and testing, and are quite volatile as shown in Figure 17 in the appendix. Aggregating to weekly new cases/deaths/tests smooths out idiosyncratic daily noises as well as periodic fluctuations associated with days of the week. We also construct weekly policy and behavior variables by taking 7 day moving averages from day $t-14$ to $t-21$ for case growth, where the delay reflects the time lag between infection and case confirmation. The four weekly behavior variables are referred as "Transit Intensity," "Workplace Intensity," "Retail Intensity," and "Grocery Intensity." Consequently, our empirical analysis uses 7 day moving averages of all variables recorded at daily frequencies. Our sample period is from March 7, 2020 to June 3, 2020.

Table 2 reports that weekly policy and behavior variables are highly correlated with each other, except for the"masks for employees" policy. High correlations may cause multicolinearity problems and could limit our ability to separately identify the effect of each policy or behavior variable on case growth, but this does not prevent us from identifying the aggregate effect of all policies and behavior variables on case or death growth.

[19]The other two measures are "Residential" and "Parks." We drop "Residential" because it is highly correlated with both "Workplaces" and "Retail & recreation" at correlation coefficients of -0.98. We also drop "Parks" because it does not have clear implication on the spread of Covid-19.

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

TABLE 2. Correlations among Policies and Behavior

	workplaces	retail	grocery	transit	masks for employees	closed K-12 schools	stay at home	closed movie theaters	closed restaurants	closed businesses
workplaces	1.00									
retail	0.94	1.00								
grocery	0.75	0.82	1.00							
transit	0.90	0.92	0.83	1.00						
masks for employees	-0.32	-0.19	-0.16	-0.30	1.00					
closed K-12 schools	-0.92	-0.81	-0.58	-0.75	0.46	1.00				
stay at home	-0.70	-0.69	-0.71	-0.72	0.31	0.65	1.00			
closed movie theaters	-0.82	-0.77	-0.65	-0.72	0.40	0.85	0.75	1.00		
closed restaurants	-0.79	-0.83	-0.69	-0.77	0.26	0.77	0.74	0.84	1.00	
closed businesses	-0.66	-0.68	-0.68	-0.66	0.12	0.59	0.77	0.69	0.73	1.00

Each off-diagonal entry reports a correlation coefficient of a pair of policy and behavior variables.

Figure 7 shows the portion of states that have each policy in place at each date. For most policies, there is considerable variation across states in the time in which the policies are active. The one exception is K-12 school closures. About 80% of states closed schools within a day or two of March 15th, and all states closed schools by early April. This makes the effect of school closings difficult to separate from aggregate time series variation.

FIGURE 7. Portion of states with each policy

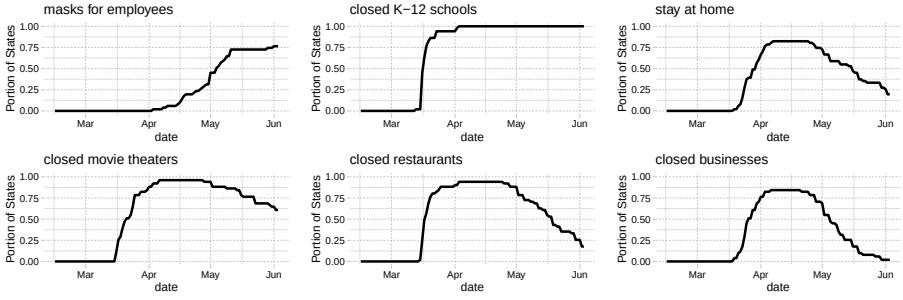

Covid Economics 35, 7 July 2020: 116-175

4.2. **The Effect of Policies and Information on Behavior.** We first examine how policies and information affect social distancing behaviors by estimating a version of (PI→B):

$$B_{it}^{j} = (\beta^{j})' P_{it} + (\gamma^{j})' I_{it} + (\delta_{B}^{j})' W_{it} + \varepsilon_{it}^{bj},$$

where B_{it}^{j} represents behavior variable j in state i at time t. P_{it} collects the Covid related policies in state i at t. Confounders, W_{it}, include state-level covariates, month indicators, and their interactions. I_{it} is a set of information variables that affect people's behaviors at t. As information, we include each state's growth of cases (in panel 3a) or deaths (in panel 3b), and log cases or deaths. Additionally, in columns (5)-(8) of each panel, we include national growth and log of cases or deaths.

Table 3 reports the estimates with standard errors clustered at the state level. Across different specifications, our results imply that policies have large effects on behavior. Comparing columns (1)-(4) with columns (5)-(8), the magnitude of policy effects are sensitive to whether national cases or deaths are included as information. The coefficient on school closures is particularly sensitive to the inclusion of national information variables. As shown in Figure 7, there is little variation across states in the timing of school closures. Consequently, it is difficult to separate the effect of school closures from a behavioral response to the national trend in cases and deaths.

The other policy coefficients are less sensitive to the inclusion of national case/death variables. After school closures, stay-at-home orders and restaurant closures have the next largest effect. Somewhat surprisingly, closure of nonessential businesses appears to have a modest effect on behavior. Closing movie theaters has a similar, small effect on behavior. The effect of masks for employees is also small. The comparison of effects across policies should be interpreted with caution. Differences between policy effects are often statistically insignificant; except for masks for employees, the policies are highly correlated and it is difficult to separate their effects.

The row "$\sum_j \text{Policy}_j$" reports the sum of the estimated effect of all policies, which is substantial and can account for a large fraction of the observed declines in behavior variables. For example, in Figure 6, transit intensity for a median state was approximately -50% at its lowest point in early April. The estimated policy coefficients in columns imply that imposing all six policies would lead to roughly 85% (in columns 1-4) or roughly 50% (in columns 5-8) of the observed decline. The large impact of policies on transit intensity suggests that the policies may have reduced the Covid-19 infection by reducing people's use of public transportation.[20]

In panel 3b, estimated coefficients of deaths and death growth are generally negative. This suggests that the higher number of deaths reduces social interactions measured by Google Mobility Reports perhaps because people are increasingly aware of prevalence of Covid-19 (Maloney and Taskin, 2020). The coefficients on log cases and case growth in panel 3a are more mixed. In columns (5)-(8) of both panels, we see that national case/death variables have large, negative coefficients. This suggests that behavior responded to national

[20]Analyzing the New York City's subway ridership, Harris (2020) finds a strong link between public transit and spread of infection.

Covid Economics 35, 7 July 2020: 116-175

TABLE 3. The Effect of Policies and Information on Behavior ($PI \rightarrow B$)

(A) Cases as Information

	Dependent variable:							
	workplaces	retail	grocery	transit	workplaces	retail	grocery	transit
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
masks for employees	−0.011	−1.207	−2.178**	−3.104	−0.812	−2.422*	−2.422***	−4.044*
	(0.873)	(1.513)	(0.952)	(2.213)	(0.660)	(1.347)	(0.902)	(2.094)
closed K-12 schools	−19.678***	−21.898***	−13.021***	−22.694***	−4.908***	−1.873	−7.923***	−5.147
	(2.830)	(4.409)	(2.536)	(5.597)	(1.526)	(1.979)	(2.944)	(4.868)
stay at home	−2.943***	−5.625***	−5.598***	−8.577***	−3.222***	−6.306***	−5.620***	−8.881***
	(1.045)	(1.346)	(1.361)	(2.366)	(0.957)	(1.154)	(1.356)	(2.347)
closed movie theaters	−1.975*	−3.444**	−2.897**	1.129	−1.464*	−3.061**	−2.643**	1.764
	(1.103)	(1.607)	(1.200)	(2.359)	(0.820)	(1.310)	(1.150)	(2.252)
closed restaurants	−3.151***	−7.682***	−1.431*	−7.969***	−1.435**	−5.095***	−0.903	−5.954**
	(1.012)	(1.500)	(0.756)	(2.557)	(0.698)	(1.002)	(0.623)	(2.365)
closed businesses	−1.942*	−1.742	−2.390**	−1.300	−2.131**	−2.147*	−2.418**	−1.510
	(1.116)	(1.362)	(1.044)	(2.039)	(0.908)	(1.125)	(0.981)	(1.917)
$\Delta \log \Delta C_{it}$	1.791***	1.046**	1.870***	1.857***	1.596***	1.155***	1.710***	1.591***
	(0.356)	(0.532)	(0.376)	(0.553)	(0.221)	(0.378)	(0.403)	(0.601)
$\log \Delta C_{it}$	−2.107***	−1.934**	0.225	−1.092	−0.366	0.210	0.880	0.997
	(0.493)	(0.900)	(0.481)	(1.175)	(0.340)	(0.784)	(0.542)	(1.285)
$\Delta \log \Delta C_{it}$.national					−2.998***	−6.952***	−0.319	−3.294***
					(0.452)	(0.759)	(0.680)	(1.187)
$\log \Delta C_{it}$.national					−6.610***	−8.957***	−2.283***	−7.854***
					(0.440)	(0.853)	(0.826)	(1.396)
state variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month × state variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
$\sum_j \text{Policy}_j$	-29.699***	-41.597***	-27.515***	-42.515***	-13.972***	-20.904***	-21.931***	-23.772***
	(3.296)	(5.343)	(3.246)	(6.813)	(1.953)	(2.859)	(3.325)	(5.127)
Observations	4,284	4,284	4,284	4,284	4,284	4,284	4,284	4,284
R^2	0.912	0.854	0.788	0.812	0.945	0.902	0.794	0.836
Adjusted R^2	0.912	0.853	0.786	0.810	0.945	0.901	0.793	0.835

Note: *p<0.1; **p<0.05; ***p<0.01

Dependent variables are "Transit Intensity," "Workplace Intensity," "Retail Intensity," and "Grocery Intensity" defined as 7 days moving averages of corresponding daily measures obtained from Google Mobility Reports. All specifications include state-level characteristics (population, area, unemployment rate, poverty rate, and a percentage of people subject to illness) as well as their interactions with the log of days since Jan 15, 2020. The row "$\sum_j \text{Policy}_j$" reports the sum of six policy coefficients. Standard errors are clustered at the state level.

Covid Economics 35, 7 July 2020: 116-175

(B) Deaths as Information

	Dependent variable:							
	workplaces	retail	grocery	transit	workplaces	retail	grocery	transit
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
masks for employees	−0.477	−2.217	−2.720**	−3.914*	−1.335**	−3.487**	−3.156***	−4.857**
	(0.753)	(1.415)	(1.059)	(2.320)	(0.642)	(1.389)	(0.989)	(2.270)
closed K-12 schools	−24.156***	−26.171***	−12.250***	−24.946***	−5.355***	−1.900	−3.859	−5.245
	(2.253)	(3.220)	(1.771)	(3.818)	(1.703)	(1.934)	(2.378)	(4.737)
stay at home	−2.579***	−5.589***	−6.090***	−8.761***	−2.799***	−5.998***	−6.229***	−9.024***
	(0.985)	(1.347)	(1.523)	(2.513)	(0.959)	(1.188)	(1.518)	(2.557)
closed movie theaters	−2.298**	−4.148**	−3.102**	0.658	−1.032	−2.661*	−2.585**	1.945
	(1.140)	(1.693)	(1.229)	(2.364)	(0.820)	(1.379)	(1.144)	(2.321)
closed restaurants	−3.479***	−7.579***	−1.317*	−7.934***	−1.507**	−4.919***	−0.400	−5.838**
	(1.104)	(1.559)	(0.752)	(2.583)	(0.707)	(1.016)	(0.660)	(2.437)
closed businesses	−2.106**	−2.351*	−2.516**	−1.656	−1.072	−0.977	−2.042*	−0.563
	(1.055)	(1.343)	(1.126)	(2.077)	(0.896)	(1.160)	(1.050)	(2.023)
$\Delta \log \Delta D_{it}$	−0.922**	−2.050***	−0.469	−1.263**	0.115	−0.278	0.136	−0.061
	(0.407)	(0.595)	(0.418)	(0.619)	(0.237)	(0.438)	(0.422)	(0.578)
$\log \Delta D_{it}$	−1.077***	−0.185	0.057	−0.262	−0.644	0.155	0.179	0.134
	(0.389)	(0.741)	(0.565)	(1.195)	(0.409)	(0.790)	(0.609)	(1.284)
$\Delta \log \Delta D_{it}$.national					−4.066***	−6.883***	−2.351***	−4.695***
					(0.353)	(0.619)	(0.449)	(0.833)
$\log \Delta D_{it}$.national					−6.322***	−7.884***	−2.731***	−6.551***
					(0.420)	(0.594)	(0.561)	(0.997)
state variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month × state variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
$\sum_j$ Policy$_j$	-35.094***	-48.055***	-27.995***	-46.554***	-13.100***	-19.941***	-18.270***	-23.581***
	(2.253)	(3.604)	(2.982)	(5.781)	(2.119)	(3.144)	(3.258)	(6.007)
Observations	4,284	4,284	4,284	4,284	4,284	4,284	4,284	4,284
R^2	0.902	0.850	0.778	0.810	0.943	0.905	0.792	0.834
Adjusted R^2	0.902	0.849	0.776	0.809	0.943	0.904	0.791	0.833

Note: *p<0.1; **p<0.05; ***p<0.01

Dependent variables are "Transit Intensity," "Workplace Intensity," "Retail Intensity," and "Grocery Intensity" defined as 7 days moving averages of corresponding daily measures obtained from Google Mobility Reports. All specifications include state-level characteristics (population, area, unemployment rate, poverty rate, and a percentage of people subject to illness) as well as their interactions with the log of days since Jan 15, 2020. The row "$\sum_j$ Policy$_j$" reports the sum of six policy coefficients. Standard errors are clustered at the state level.

conditions although it is also likely that national case/death variables capture unobserved aggregate time effects beyond information (e.g., latent policy variables and time-varying confounders that are common across states) which are not fully controlled by month dummies.

FIGURE 8. Case and death growth conditional on mask mandates

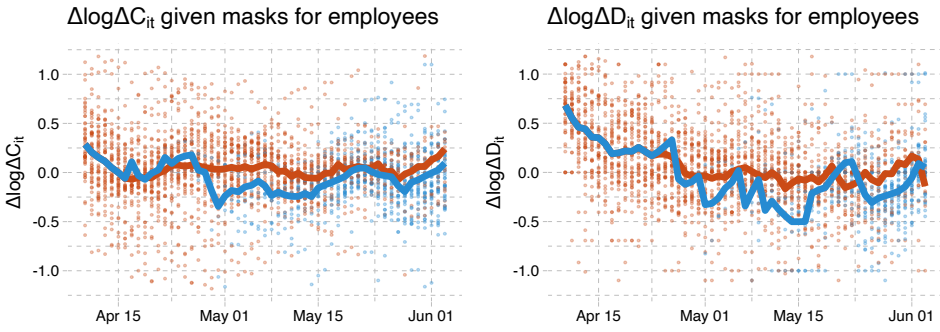

In these figures, red points are the case or death growth rate in states without a mask mandate. Blue points are states with a mask mandate 14 (21 for deaths) days prior. The red line is the average across states without a mask mandate 14 (21 for deaths) days earlier. The blue line is the average across states with a mask mandate 14 (21 for deaths) earlier.

4.3. **The Direct Effect of Policies and Behavior on Case and Death Growth.** We now analyze how behavior and policies together influence case and death growth rates. We begin with some simple graphical evidence of the effect of policies on case and death growth. Figure 8 shows average case and death growth conditional on date and whether masks are mandatory for employees.[21] The left panel of the figure shows that states with a mask mandate consistently have 0-0.2 lower case growth than states without. The right panel also illustrates that states with a mask mandate tend to have lower average death growth than states without a mask mandate.

Similar plots are shown for other policies in Figures 23 and 24 in the appendix. The figures for stay-at-home orders and closure of nonessential businesses are qualitatively similar to that for masks. States with these two policies appear to have about 0.1 percentage point lower case growth than states without. The effects of school closures, movie theater closures, and restaurant closures are not clearly visible in these figures. These figures are merely suggestive; the patterns observed in them may be driven by confounders.

We more formally analyze the effect of policies by estimating regressions. We first look at the direct effect of policies on case and death growth conditional on behavior by estimating

[21]We take 14 and 21 day lags of mask policies for case and death growths, respectively, to identify the states with a mask mandate because policies affect cases and deaths with time lags. See our discussion in the Appendix A.6.

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

TABLE 4. The Direct Effect of Behavior and Policies on Case and Death Growth ($BPI \rightarrow Y$)

	Dependent variable:			
	$\Delta \log \Delta C_{it}$			
	(1)	(2)	(3)	(4)
lag(masks for employees, 14)	−0.084**		−0.097***	
	(0.035)		(0.032)	
lag(masks*April, 14)		−0.098*		−0.111**
		(0.051)		(0.051)
lag(masks*May, 14)		−0.080**		−0.094***
		(0.038)		(0.034)
lag(closed K-12 schools, 14)	−0.095	−0.096	0.025	0.024
	(0.089)	(0.089)	(0.103)	(0.103)
lag(stay at home, 14)	−0.041	−0.042	−0.064	−0.065
	(0.047)	(0.048)	(0.048)	(0.049)
lag(closed movie theaters, 14)	0.049	0.047	0.053	0.052
	(0.049)	(0.049)	(0.048)	(0.049)
lag(closed restaurants, 14)	0.020	0.020	0.021	0.021
	(0.048)	(0.048)	(0.046)	(0.045)
lag(closed businesses, 14)	−0.004	−0.003	−0.016	−0.015
	(0.042)	(0.042)	(0.042)	(0.042)
lag(workplaces, 14)	0.010*	0.010*	0.003	0.003
	(0.006)	(0.006)	(0.006)	(0.006)
lag(retail, 14)	0.005*	0.005*	0.003	0.003
	(0.003)	(0.003)	(0.003)	(0.003)
lag(grocery, 14)	−0.004	−0.004	−0.002	−0.002
	(0.003)	(0.003)	(0.003)	(0.003)
lag(transit, 14)	0.003	0.003	0.003	0.003
	(0.003)	(0.003)	(0.003)	(0.003)
lag($\Delta \log \Delta C_{it}$, 14)	0.017	0.017	0.023	0.023
	(0.025)	(0.025)	(0.028)	(0.028)
lag($\log \Delta C_{it}$, 14)	−0.110***	−0.110***	−0.089***	−0.089***
	(0.019)	(0.019)	(0.021)	(0.021)
lag($\Delta \log \Delta C_{it}$.national, 14)			−0.090**	−0.089**
			(0.044)	(0.044)
lag($\log \Delta C_{it}$.national, 14)			−0.184***	−0.184***
			(0.048)	(0.048)
$\Delta \log T_{it}$	0.153***	0.153***	0.158***	0.158***
	(0.044)	(0.044)	(0.042)	(0.042)
state variables	Yes	Yes	Yes	Yes
Month × state variables	Yes	Yes	Yes	Yes
$\sum_j$ Policy$_j$	-0.155	-0.252	-0.078	-0.189
	(0.136)	(0.156)	(0.160)	(0.178)
$\sum_k w_k$Behavior$_k$	-0.756***	-0.753***	-0.372**	-0.368**
	(0.143)	(0.144)	(0.153)	(0.152)
Observations	3,823	3,823	3,823	3,823
R^2	0.759	0.759	0.765	0.765
Adjusted R^2	0.757	0.757	0.762	0.762
Note:				*p<0.1; **p<0.05; ***p<0.01

	Dependent variable:			
	$\Delta \log \Delta D_{it}$			
	(1)	(2)	(3)	(4)
lag(masks for employees, 21)	−0.145***		−0.148***	
	(0.052)		(0.051)	
lag(masks*April, 21)		−0.159*		−0.166**
		(0.082)		(0.082)
lag(masks*May, 21)		−0.138**		−0.138**
		(0.061)		(0.059)
lag(closed K-12 schools, 21)	−0.271***	−0.272***	−0.199**	−0.201**
	(0.096)	(0.096)	(0.094)	(0.094)
lag(stay at home, 21)	−0.040	−0.041	−0.047	−0.048
	(0.065)	(0.065)	(0.063)	(0.064)
lag(closed movie theaters, 21)	0.039	0.038	0.054	0.052
	(0.091)	(0.091)	(0.090)	(0.090)
lag(closed restaurants, 21)	0.085	0.085	0.081	0.081
	(0.067)	(0.067)	(0.067)	(0.067)
lag(closed businesses, 21)	−0.003	−0.003	−0.003	−0.002
	(0.056)	(0.056)	(0.058)	(0.058)
lag(workplaces, 21)	0.014***	0.014***	0.009	0.009
	(0.005)	(0.005)	(0.005)	(0.006)
lag(retail, 21)	0.006	0.006	0.006	0.006
	(0.004)	(0.004)	(0.004)	(0.004)
lag(grocery, 21)	−0.010***	−0.010***	−0.010***	−0.010***
	(0.004)	(0.004)	(0.004)	(0.004)
lag(transit, 21)	0.003	0.003	0.003	0.003
	(0.003)	(0.003)	(0.003)	(0.003)
lag($\Delta \log \Delta D_{it}$, 21)	0.016	0.016	0.017	0.017
	(0.034)	(0.034)	(0.037)	(0.037)
lag($\log \Delta D_{it}$, 21)	−0.051**	−0.051**	−0.049**	−0.049**
	(0.025)	(0.025)	(0.024)	(0.024)
lag($\Delta \log \Delta D_{it}$.national, 21)			−0.046	−0.047
			(0.045)	(0.045)
lag($\log \Delta D_{it}$.national, 21)			−0.060	−0.060
			(0.039)	(0.039)
state variables	Yes	Yes	Yes	Yes
Month × state variables	Yes	Yes	Yes	Yes
$\sum_j$ Policy$_j$	-0.334**	-0.489**	-0.262	-0.422**
	(0.164)	(0.198)	(0.176)	(0.208)
$\sum_k w_k$Behavior$_k$	-0.871***	-0.870***	-0.646***	-0.642***
	(0.164)	(0.163)	(0.181)	(0.179)
Observations	3,468	3,468	3,468	3,468
R^2	0.521	0.521	0.522	0.522
Adjusted R^2	0.516	0.516	0.517	0.517
Note:				*p<0.1; **p<0.05; ***p<0.01

Dependent variable is the weekly growth rate of confirmed cases (in the left panel) or deaths (in the right panel) as defined in equation (3). The covariates include lagged policy and behavior variables, which are constructed as 7 day moving averages between t to $t-7$ of corresponding daily measures. The row "$\sum_j$ Policies$_j$" reports the sum of six policy coefficients. The row "$\sum_k w_k$Behavior$_k$" reports the sum of four coefficients of behavior variables weighted by the average of each behavioral variable from April 1st-10th. Standard errors are clustered at the state level.

equation (BPI→Y):

$$Y_{i,t+\ell} = \alpha' B_{it} + \pi' P_{it} + \mu' I_{it} + \delta_Y' W_{it} + \varepsilon_{it}^y, \tag{14}$$

where the outcome variable, $Y_{i,t+\ell}$, is either case growth or death growth.

For case growth as the outcome, we choose a lag length of $\ell = 14$ days for behavior, policy, and information variables to reflect the delay between infection and confirmation of case.[22] $B_{it} = (B_{it}^1, ..., B_{it}^4)'$ is a vector of four behavior variables in state i. P_{it} includes the Covid-related policies in state i that directly affect the spread of Covid-19 after controlling for behavior variables (e.g., masks for employees). We include information variables, I_{it}, that include the past cases and case growths because the past cases may be correlated with (latent) government policies or people's behaviors that are not fully captured by our observed policy and behavior variables. We also consider a specification that includes the past cases and case growth at the national level as additional information variables. W_{it} is a set of confounders that include month dummies, state-level covariates, and the interaction terms between month dummies and state-level covariates.[23] For case growth, W_{it} also includes the test rate growth $\Delta \log(T)_{it}$ to capture the effect of changing test rates on confirmed cases. Equation (14) corresponds to (M-C) derived from the SIR model.

For death growth as the outcome, we take a lag length of $\ell = 21$ days. The information variables I_{it} include past deaths and death growth rates; W_{it} is the same as that of the case growth equation except that the growth rate of test rates is excluded from W_{it} as implied by equation (M-D).

Table 4 shows the results of estimating (14) for case and death growth rates. Column (1) represents our baseline specification while column (2) allows the effect of masks to be different before and after May 1st. Columns (3) and (4) include past cases/deaths and growth rates at national level as additional regressors.

The estimates indicate that mandatory face masks for employees reduce the growth rate of infections and deaths by 8-15 percent, while holding behavior constant. This suggests that requiring masks for employees in public-facing businesses may be an effective preventive measure.[24] The estimated effect of masks on death growth is larger than the effect on case growth, but this difference between the two estimated effects is not statistically significant.

[22] As we review in the Appendix A.6, a lag length of 14 days between exposure and case reporting, as well as a lag length of 21 days between exposure and deaths, is broadly consistent with currently available evidence.

[23] Month dummies also represent the latent information that is not fully captured by the past cases and growths.

[24] Note that we are *not* evaluating the effect of *universal* mask-wearing for the public but that of mask-wearing for employees. The effect of *universal* mask-wearing for the public could be larger if people comply with such a policy measure. Tian et al. (2020) considered a model in which mask wearing reduces the reproduction number by a factor $(1 - e \cdot pm)^2$, where e is the efficacy of trapping viral particles inside the mask and pm is the percentage of mask-wearing population. Given an estimate of $R_0 = 2.4$, Howard et al. (2020) argue that 50% mask usage and a 50% mask efficacy level would reduce the reproduction number from 2.4 to 1.35, an order of magnitude impact.

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

Except for mask requirements, policies appear to have little direct effect on case or death growth when behavior is held constant. The one exception is that closing schools has a large and statistically significant coefficient in the death growth regressions. As discussed above, there is little cross-state variation in the timing of school closures, making estimates of its effect less reliable.

The row "$\sum_k w_k \text{Behavior}_k$" reports the sum of estimated coefficients weighted by the average of the behavioral variables from April 1st-10th. The estimates of -0.76 and -0.87 for "$\sum_k w_k \text{Behavior}_k$" in column (1) imply that a reduction in mobility measures relative to the baseline in January and February have induced a decrease in case and death growth rates by 76 and 83 percent, respectively, suggesting an importance of social distancing for reducing the spread of Covid-19. When including national cases and deaths in information, as shown in columns (3) and (4), the estimated aggregate impact of behavior is substantially smaller, but remains large and statistically significant.

A useful practical implication of these results are that Google Mobility Reports and similar data might be useful as a leading indicator of potential case or death growth. This should be done with caution, however, because other changes in the environment might alter the relationship between behavior and infections. Preventative measures, including mandatory face masks, and changes in habit that are not captured in our data might alter the future relationship between Google Mobility Reports and case/death growth.

The negative coefficients of past cases or deaths in Table 4 is consistent with a hypothesis that higher reported cases and deaths change people's behavior to reduce transmission risks. Such behavioral changes in response to new information are partly captured by Google mobility measures, but the negative estimated coefficient of past cases or deaths imply that other latent behavioral changes that are not fully captured by Google mobility measures (e.g., frequent hand-washing, wearing masks, and keeping 6ft/2m distancing) are also important for reducing future cases and deaths.

If policies are enacted and behavior changes, then future cases/deaths and information will change, which will induce further behavior changes. However, since the model includes lags of cases/deaths as well as their growth rates, computing a long-run effect is not completely straightforward. We investigate dynamic effects that incorporate feedback through information in section 5.

4.4. **The Total Effect of Policies on Case Growth.** In this section, we focus our analysis on policy effects when we hold information constant. The estimated effect of policy on behavior in Table 3 and those of policies and behavior on case/death growth in Table 4 can be combined to calculate the total effect of policy as well as its decomposition into direct and indirect effects.

The first three columns of Table 6 show the direct (holding behavior constant) and indirect (through behavior changes) effects of policy under a specification that excludes national information variables. These are computed from the specification with national cases or deaths included as information (columns (1)-(4) of Table 3 and column (1) of Table

Covid Economics 35, 7 July 2020: 116-175

TABLE 5. The Total Effect of Policies on Case and Death Growth ($PI \rightarrow Y$)

	Dependent variable: $\Delta \log \Delta C_{it}$			
	(1)	(2)	(3)	(4)
lag(masks for employees, 14)	−0.081**		−0.105***	
	(0.041)		(0.037)	
lag(masks*April, 14)		−0.157**		−0.146**
		(0.067)		(0.061)
lag(masks*May, 14)		−0.062		−0.094***
		(0.039)		(0.036)
lag(closed K-12 schools, 14)	−0.240**	−0.241**	0.009	0.007
	(0.097)	(0.097)	(0.109)	(0.108)
lag(stay at home, 14)	−0.126**	−0.128**	−0.117**	−0.118**
	(0.055)	(0.055)	(0.052)	(0.052)
lag(closed movie theaters, 14)	0.030	0.023	0.058	0.054
	(0.052)	(0.052)	(0.047)	(0.047)
lag(closed restaurants, 14)	−0.042	−0.039	−0.010	−0.009
	(0.049)	(0.048)	(0.045)	(0.044)
lag(closed businesses, 14)	−0.048	−0.041	−0.035	−0.031
	(0.050)	(0.050)	(0.044)	(0.044)
lag($\Delta \log \Delta C_{it}$, 14)	0.040*	0.039*	0.033	0.032
	(0.024)	(0.024)	(0.028)	(0.028)
lag($\log \Delta C_{it}$, 14)	−0.138***	−0.138***	−0.091***	−0.091***
	(0.024)	(0.023)	(0.026)	(0.026)
lag($\Delta \log \Delta C_{it}$.national, 14)			−0.123***	−0.121***
			(0.043)	(0.042)
lag($\log \Delta C_{it}$.national, 14)			−0.241***	−0.239***
			(0.044)	(0.044)
$\Delta \log T_{it}$	0.157***	0.158***	0.161***	0.161***
	(0.044)	(0.044)	(0.042)	(0.042)
state variables	Yes	Yes	Yes	Yes
Month × state variables	Yes	Yes	Yes	Yes
$\sum_j$ Policy$_j$	-0.508***	-0.644***	-0.199	-0.336*
	(0.162)	(0.198)	(0.164)	(0.187)
Observations	3,823	3,823	3,823	3,823
R^2	0.748	0.749	0.761	0.762
Adjusted R^2	0.746	0.747	0.759	0.759
Note:				*p<0.1; **p<0.05; ***p<0.01

	Dependent variable: $\Delta \log \Delta D_{it}$			
	(1)	(2)	(3)	(4)
lag(masks for employees, 21)	−0.133**		−0.161***	
	(0.053)		(0.052)	
lag(masks*April, 21)		−0.174*		−0.193**
		(0.089)		(0.091)
lag(masks*May, 21)		−0.112*		−0.145**
		(0.057)		(0.057)
lag(closed K-12 schools, 21)	−0.641***	−0.641***	−0.250**	−0.252**
	(0.117)	(0.117)	(0.103)	(0.104)
lag(stay at home, 21)	−0.080	−0.082	−0.075	−0.076
	(0.065)	(0.065)	(0.062)	(0.062)
lag(closed movie theaters, 21)	0.018	0.015	0.065	0.063
	(0.089)	(0.089)	(0.084)	(0.084)
lag(closed restaurants, 21)	−0.015	−0.013	0.031	0.033
	(0.059)	(0.059)	(0.055)	(0.055)
lag(closed businesses, 21)	−0.038	−0.035	−0.012	−0.010
	(0.066)	(0.064)	(0.063)	(0.062)
lag($\Delta \log \Delta D_{it}$, 21)	−0.0002	0.0002	0.019	0.019
	(0.033)	(0.033)	(0.036)	(0.036)
lag($\log \Delta D_{it}$, 21)	−0.078***	−0.078***	−0.062**	−0.063**
	(0.027)	(0.027)	(0.027)	(0.027)
lag($\Delta \log \Delta D_{it}$.national, 21)			−0.160***	−0.160***
			(0.057)	(0.057)
lag($\log \Delta D_{it}$.national, 21)			−0.120***	−0.119***
			(0.030)	(0.030)
state variables	Yes	Yes	Yes	Yes
Month × state variables	Yes	Yes	Yes	Yes
$\sum_j$ Policy$_j$	-0.889***	-1.042***	-0.402**	-0.580***
	(0.171)	(0.213)	(0.184)	(0.222)
Observations	3,468	3,468	3,468	3,468
R^2	0.504	0.504	0.515	0.515
Adjusted R^2	0.499	0.499	0.510	0.510
Note:				*p<0.1; **p<0.05; ***p<0.01

Dependent variable is the weekly growth rate of confirmed cases (in the left panel) or deaths (in the right panel) as defined in equation (3). The covariates include lagged policy variables, which are constructed as 7 day moving averages between t to $t-7$ of corresponding daily measures. The row "$\sum_j$ Policies$_j$" reports the sum of six policy coefficients. Standard errors are clustered at the state level.

Covid Economics 35, 7 July 2020: 116-175

TABLE 6. Direct and Indirect Policy Effects without national case/death variables

	Direct	Indirect	Total	PI→Y Coef.	Average	Difference
Cases						
masks for employees	-0.084**	-0.008	-0.092**	-0.081**	-0.086**	-0.011
	(0.034)	(0.024)	(0.044)	(0.040)	(0.041)	(0.015)
closed K-12 schools	-0.095	-0.337***	-0.432***	-0.240**	-0.336***	-0.192***
	(0.093)	(0.091)	(0.118)	(0.095)	(0.105)	(0.047)
stay at home	-0.041	-0.065**	-0.106**	-0.126**	-0.116**	0.020
	(0.046)	(0.031)	(0.053)	(0.055)	(0.054)	(0.013)
closed movie theaters	0.049	-0.024	0.024	0.030	0.027	-0.005
	(0.048)	(0.025)	(0.055)	(0.050)	(0.052)	(0.016)
closed restaurants	0.020	-0.091***	-0.071	-0.042	-0.057	-0.029*
	(0.046)	(0.029)	(0.044)	(0.048)	(0.045)	(0.016)
closed businesses	-0.004	-0.024	-0.028	-0.048	-0.038	0.020*
	(0.041)	(0.019)	(0.049)	(0.050)	(0.049)	(0.011)
$\sum_j \text{Policy}_j$	-0.155	-0.550***	-0.704***	-0.508***	-0.606***	-0.196***
	(0.136)	(0.140)	(0.188)	(0.157)	(0.171)	(0.052)
$\Delta \log \Delta C_{it}$	0.017	0.023**	0.040*	0.040*	0.040*	0.000
	(0.025)	(0.010)	(0.023)	(0.024)	(0.023)	(0.006)
$\log \Delta C_{it}$	-0.110***	-0.036**	-0.146***	-0.138***	-0.142***	-0.008
	(0.019)	(0.014)	(0.026)	(0.023)	(0.024)	(0.007)
Deaths						
	Direct	Indirect	Total	PI→Y Coef.	Average	Difference
masks for employees	-0.145***	-0.004	-0.149***	-0.133***	-0.141***	-0.016
	(0.050)	(0.023)	(0.055)	(0.051)	(0.052)	(0.015)
closed K-12 schools	-0.271***	-0.451***	-0.722***	-0.641***	-0.681***	-0.081***
	(0.092)	(0.082)	(0.111)	(0.107)	(0.108)	(0.026)
stay at home	-0.040	-0.034	-0.074	-0.080	-0.077	0.006
	(0.064)	(0.035)	(0.064)	(0.064)	(0.064)	(0.015)
closed movie theaters	0.039	-0.025	0.014	0.018	0.016	-0.004
	(0.091)	(0.030)	(0.089)	(0.089)	(0.088)	(0.018)
closed restaurants	0.085	-0.105**	-0.020	-0.015	-0.018	-0.005
	(0.065)	(0.042)	(0.056)	(0.057)	(0.056)	(0.016)
closed businesses	-0.003	-0.024	-0.027	-0.038	-0.032	0.011
	(0.055)	(0.021)	(0.061)	(0.063)	(0.062)	(0.013)
$\sum_j \text{Policy}_j$	-0.334**	-0.644***	-0.979***	-0.889***	-0.934***	-0.090**
	(0.160)	(0.154)	(0.171)	(0.165)	(0.167)	(0.035)
$\Delta \log \Delta D_{it}$	0.016	-0.025**	-0.009	-0.000	-0.004	-0.009*
	(0.034)	(0.011)	(0.031)	(0.032)	(0.031)	(0.005)
$\log \Delta D_{it}$	-0.051**	-0.018*	-0.069**	-0.078***	-0.073***	0.009*
	(0.024)	(0.010)	(0.028)	(0.026)	(0.027)	(0.005)

Direct effects capture the effect of policy on case growth holding behavior, information, and confounders constant. Direct effects are given by π in equation (BPI→Y). Indirect effects capture how policy changes behavior and behavior shift case growth. They are given by α from (BPI→Y) times β from (PI→B). The total effect is $\pi + \beta\alpha$. Column "PI→Y Coefficients" shows the coefficient estimates from PI→Y. Columns "Difference" are the differences between the estimates from (PI→Y) and the combination of (BPI→Y) and (PI→B) while column "Average" are their averages. Standard errors are computed by bootstrap and clustered on state.

Covid Economics 35, 7 July 2020: 116-175

TABLE 7. Direct and Indirect Policy Effects with national case/death variables

Cases

	Direct	Indirect	Total	PI→Y Coef.	Average	Difference
masks for employees	-0.097***	-0.019	-0.116***	-0.105***	-0.111***	-0.011
	(0.033)	(0.017)	(0.040)	(0.038)	(0.039)	(0.011)
closed K-12 schools	0.025	-0.021	0.004	0.009	0.007	-0.005
	(0.103)	(0.040)	(0.110)	(0.108)	(0.109)	(0.015)
stay at home	-0.064	-0.047**	-0.112**	-0.117**	-0.114**	0.005
	(0.047)	(0.023)	(0.049)	(0.049)	(0.049)	(0.009)
closed movie theaters	0.053	-0.002	0.051	0.058	0.055	-0.006
	(0.048)	(0.017)	(0.048)	(0.046)	(0.047)	(0.011)
closed restaurants	0.021	-0.038*	-0.017	-0.010	-0.013	-0.008
	(0.045)	(0.020)	(0.041)	(0.043)	(0.041)	(0.011)
closed businesses	-0.016	-0.013	-0.028	-0.035	-0.032	0.006
	(0.042)	(0.012)	(0.044)	(0.044)	(0.044)	(0.008)
$\sum_j \text{Policy}_j$	-0.078	-0.140**	-0.218	-0.199	-0.209	-0.019
	(0.160)	(0.065)	(0.168)	(0.166)	(0.167)	(0.018)
$\Delta \log \Delta C_{it}$	0.023	0.010	0.033	0.033	0.033	-0.000
	(0.028)	(0.007)	(0.028)	(0.028)	(0.028)	(0.003)
$\log \Delta C_{it}$	-0.089***	0.001	-0.088***	-0.091***	-0.090***	0.003
	(0.021)	(0.011)	(0.028)	(0.027)	(0.027)	(0.005)
$\Delta \log \Delta C_{it}$.national	-0.090**	-0.040**	-0.130***	-0.123***	-0.126***	-0.006
	(0.044)	(0.016)	(0.044)	(0.042)	(0.042)	(0.013)
$\log \Delta C_{it}$.national	-0.184***	-0.068***	-0.252***	-0.241***	-0.247***	-0.010
	(0.047)	(0.022)	(0.044)	(0.044)	(0.044)	(0.010)

Deaths

	Direct	Indirect	Total	PI→Y Coef.	Average	Difference
masks for employees	-0.148***	-0.018	-0.166***	-0.161***	-0.164***	-0.005
	(0.048)	(0.023)	(0.053)	(0.050)	(0.051)	(0.016)
closed K-12 schools	-0.199**	-0.038	-0.238**	-0.250**	-0.244**	0.012
	(0.091)	(0.038)	(0.100)	(0.099)	(0.099)	(0.020)
stay at home	-0.047	-0.030	-0.077	-0.075	-0.076	-0.002
	(0.065)	(0.032)	(0.063)	(0.063)	(0.063)	(0.014)
closed movie theaters	0.054	0.007	0.061	0.065	0.063	-0.004
	(0.090)	(0.021)	(0.086)	(0.083)	(0.084)	(0.016)
closed restaurants	0.081	-0.058**	0.023	0.031	0.027	-0.008
	(0.064)	(0.024)	(0.053)	(0.054)	(0.053)	(0.014)
closed businesses	-0.003	0.003	-0.000	-0.012	-0.006	0.012
	(0.056)	(0.016)	(0.059)	(0.060)	(0.059)	(0.012)
$\sum_j \text{Policy}_j$	-0.262	-0.135	-0.397**	-0.402**	-0.399**	0.005
	(0.167)	(0.085)	(0.179)	(0.174)	(0.176)	(0.024)
$\Delta \log \Delta D_{it}$	0.017	-0.002	0.015	0.019	0.017	-0.004
	(0.037)	(0.005)	(0.036)	(0.036)	(0.036)	(0.004)
$\log \Delta D_{it}$	-0.049**	-0.006	-0.055**	-0.062**	-0.059**	0.007
	(0.024)	(0.009)	(0.028)	(0.027)	(0.027)	(0.005)
$\Delta \log \Delta D_{it}$.national	-0.046	-0.069***	-0.115**	-0.160***	-0.137***	0.045***
	(0.046)	(0.021)	(0.050)	(0.057)	(0.053)	(0.013)
$\log \Delta D_{it}$.national	-0.060	-0.097***	-0.157***	-0.120***	-0.138***	-0.037***
	(0.038)	(0.029)	(0.032)	(0.029)	(0.030)	(0.012)

Direct effects capture the effect of policy on case growth holding behavior, information, and confounders constant. Direct effects are given by π in equation (BPI→Y). Indirect effects capture how policy changes behavior and behavior shift case growth. They are given by α from (BPI→Y) times β from (PI→B). The total effect is $\pi + \beta\alpha$. Column "PI→Y Coefficients" shows the coefficient estimates from PI→Y. Columns "Difference" are the differences between the estimates from (PI→Y) and the combination of (BPI→Y) and (PI→B) while column "Average" are their averages. Standard errors are computed by bootstrap and clustered on state.

Covid Economics 35, 7 July 2020: 116-175

4). The estimates imply that all policies combined would reduce the growth rate of cases and deaths by 0.70 and 0.98, respectively, out of which about two-third to three-fourth is attributable to the indirect effect through their impact on behavior. The estimate also indicates that the effect of mandatory masks for employees is mostly direct.

We can also examine the total effect of policies and information on case or death growth, by estimating (PI→Y). The coefficients on policy in this regression combine both the direct and indirect effects.

Table 5 shows the full set of coefficient estimates for (PI→Y). The results are broadly consistent with what we found above. As in Table 3, the effect of school closures is sensitive to the inclusion of national information variables. Also as above, mask mandates have a significant negative effect on growth rates.

In columns (2) and (4) of Table 5, we find that the estimated effect of mask mandates in April is larger than that in May for both case and death regressions. This may reflect a wider *voluntary* adoption of masks in May than in April — if more people wear masks even without mandatory mask policy, the policy effect of mandating masks for employees becomes weaker.

Table 7 presents the estimates for the specification with past national case/death variables. The effects of school closures and the sum of policies are estimated substantially smaller in Table 7 when national case/death variables are included than in Table 6. This sensitivity reflects the difficulty in identifying the aggregate time effect—which is largely captured by national cases/deaths—given little cross-sectional variation in the timing of school closures across states. On the other hand, the estimated effects of policies other than school closures are similar between Table 6 and Table 7; the effect of other policies are well-identified from cross-sectional variations.

Column "Difference" in Tables 6 and 7 show the difference between the estimate of (PI→Y) in column "PI→Y Coefficient" and the implied estimate from (BPI→Y)-(PI→B) in column "Total." Differences are generally small and statistically insignificant, broadly supporting the validity of extra orthogonality condition in (BPI→Y). The difference for school closures as well as the sum of all policies in Table 6 is significantly different from zero, which may be due to the aforementioned difficulty in identifying the effect of school closures. There is substantial external epidemiological evidence that suggests that schooling closures may have substantial effects on the spread of the virus: studies like Jones et al. (2020) and Davies et al. (2020) establish that children carry substantial amounts of viral loads and can contribute to the transmission (due to higher contact rate than other age groups).[25] The US data does not allow us to pint down the effect of closing schools reliably due to their approximate collinearity with trends in national cases.

Column "Average" of Tables 6 and 7 reports the average of "Total" and "PI→Y Coefficient" columns. The average is an appealing and simple way to combine the two estimates

[25]The evidence presented in Jones et al. (2020) has lead German to make the decision to close schools early.

of the total effect: one relying on the causal structure and another inferred from a direct estimation of equation (PI $\rightarrow$ Y).[26] We shall be using the average estimate in generating the counterfactuals in the next section. Turning to the results, the estimates of Tables 6 and 7 imply that all policies combined would reduce $\Delta \log \Delta D$ by 0.97 and 0.40, respectively. For comparison, the median of $\Delta \log \Delta D_{it}$ reached its peak in mid-March of about 1.3 (see Figure 20 in the appendix). Since then it has declined to near 0. Therefore, -0.97 and -0.40 imply that policy changes can account for roughly one-third to two-third of the observed decrease in death growth. The remainder of the decline is likely due to changes in behavior from information.

5. Empirical Evaluation of Counterfactual Policies

We now turn our focus to dynamic feedback effects. Policy and behavior changes that reduce case and death growth today can lead to a more optimistic, riskier behavior in the future, attenuating longer run effects. We perform the main counterfactual experiments using the average of two estimated coefficients as reported in column "Average" of Table 6 under a specification that excludes the number of past national cases and deaths from information variables. In the appendix, we also report additional counterfactual experiment results with the specification that includes the national information variables, and find that they are very similar. The results on mask policies, business closures, stay-at-home orders are robust with respect to this variation (see Figures 10-13 in the appendix). On the other hand, the results on removing all policies, particularly closure of schools, reported in the next section, are sensitive to the inclusion of national information variables, highlighting the large uncertainty regarding the size of the effect. In Figures 9-13 below, the top panel presents the result on cases while the bottom panel presents the result on deaths.

5.1. **Business Mask Mandate.** We first consider the impact of a nationwide mask mandate for employees beginning on April 1st. As discussed earlier, we find that mask mandates reduce case and death growth even when holding behavior constant. In other words, mask mandates may reduce infections with relatively little economic disruption. This makes mask mandates a particularly attractive policy instrument. In this section we examine what would have happened to the number of cases if all states had imposed a mask mandate on April 1st.[27]

For illustrative purpose, we begin by focusing on Washington State. The left column of Figure 9 shows the observed, estimated average, and counterfactual average of $\Delta \log \Delta C$ (top panel) and $\Delta \log \Delta D$ (bottom panel). To compute the estimated and counterfactual

[26] Averaging the two estimates theoretically reduces noise, albeit in our case the reductions are small. Another approach would be to use precision averaging, which would give similar result. Finally, another approach would be to use generalized method of moments to estimate all of the equations jointly. We don't pursue this approach since it is likely to be non-robust under local deviations from correct specification; simple model averaging is more appealing in this case.

[27] We feel this is a very plausible counterfactual policy. In a paper made publicly available on April 1st, Abaluck et al. (2020) argued for mask usage based on comparisons between countries with and without pre-existing norms of widespread mask usage.

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

FIGURE 9. Effect of mandating masks on April 1st in Washington State

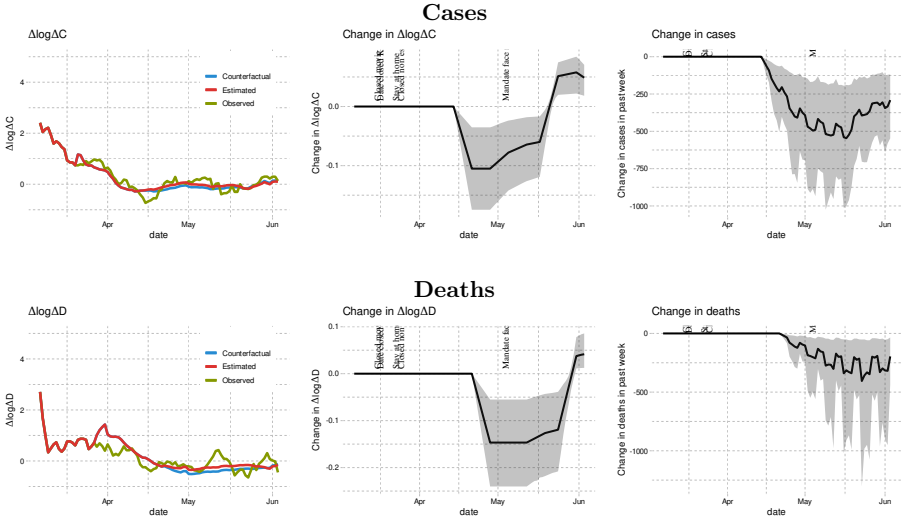

To compute the estimated and counterfactual paths we use the average of two estimated coefficients as reported in column "Average" of Table 6. We set initial $\Delta \log \Delta C$ and $\log \Delta C$ to their values first observed in the state we are simulating. We hold all other regressors at their observed values. Error terms are drawn with replacement from the residuals. We do this many times and report the average over draws of the residuals. The shaded region is a point-wise 90% confidence interval.

paths, we use the estimate in column "Average" of Table 6. We set initial $\Delta \log \Delta C$ and $\log \Delta C$ to their values first observed in the state we are simulating. We hold all other regressors at their observed values. Error terms are drawn with replacement from the residuals. We do this many times and report the average over draws of the residuals. The shaded region is a point-wise 90% confidence interval. The left column shows that the fit of the estimated and observed growth rate is quite good.

The middle column of Figure 9 shows the change in growth rate from mandating masks on April 1st. The shaded region is a 90% pointwise confidence interval. As shown, mandating masks on April 1st lowers the growth of cases or deaths 14 or 21 days later by 0.1 to 0.15. This effect then gradually declines due to information feedback. Mandatory masks reduce past cases or deaths, which leads to less cautious behavior, attenuating the impact of the policy. The reversal of the decrease in growth in late April is due to our comparison of a mask mandate on April 1st with Washington's actual mask mandate in early May. By late April, the counterfactual mask effect has decayed through information feedback, and we are comparing it the undecayed impact of Washington's actual, later mask mandate.

Covid Economics 35, 7 July 2020: 116-175

The right column of Figure 9 shows how the changes in case and death growth translate into changes in cases and deaths. The estimates imply that mandating masks on April 1st would have led to 500 fewer cases and 250 fewer deaths in Washington by the start of June.

FIGURE 10. Effect of nationally mandating masks for employees on April 1st in the US

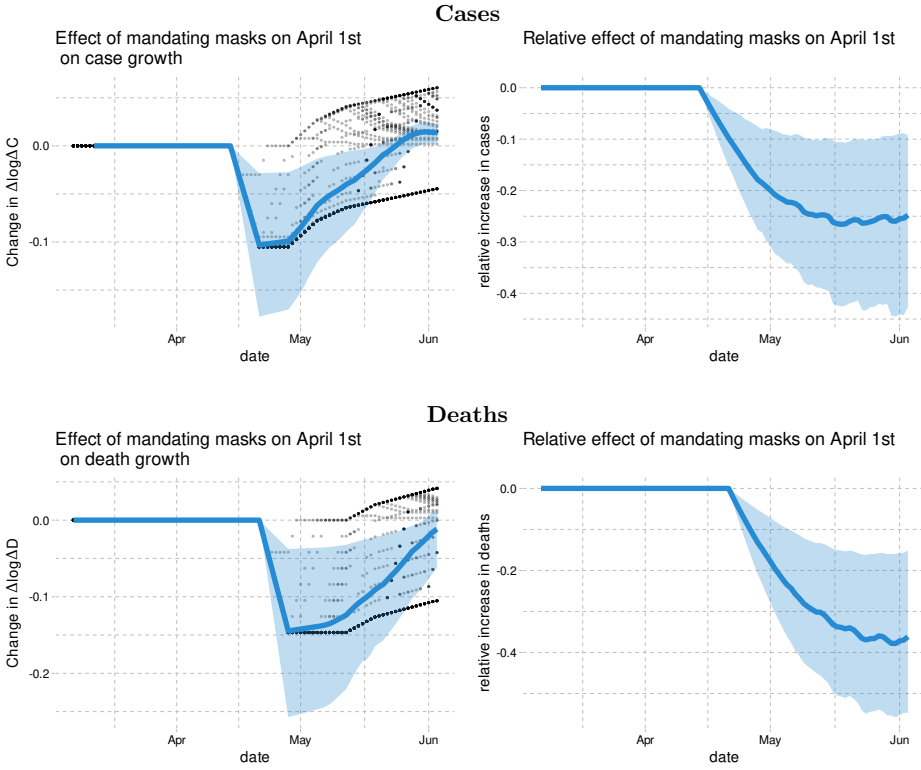

In the left column, the dots are the average change in growth in each state. The blue line is the average across states of the change in growth. The shaded region is a point-wise 90% confidence interval. The right column shows the change in cases or deaths relative to the baseline of actual policies.

The results for other states are similar to those for Washington. In the appendix, Figures 25 and 26 display similar results for Massachusetts and Illinois. Figure 10 shows the average change in cases and deaths across states, where the top panel shows the effect on cases and the bottom panel shows the effect on deaths. The point estimates indicate that mandating masks on April 1st could have led to 25% fewer cumulative cases and 37% fewer cumulative deaths by the end of May with their 90 percent intervals given by $[10, 47]\%$ and $[18, 55]\%$, respectively. The result roughly translates into 18 to 55 thousand saved lives.

Covid Economics 35, 7 July 2020: 116-175

5.2. **Non-essential Business Closures.** A particularly controversial policy is the closure of non-essential businesses. We now examine a counterfactual where non-essential businesses are never closed. Figure 11 shows the effect of leaving non-essential businesses open in Washington. The point estimate implies that the closure of non-essential businesses reduced cases and deaths by a small amount. However, this estimate is relatively imprecise; 90% confidence intervals for the change in cases and deaths from leaving non-essential businesses open by the end of May are [-250,700] and [-100,1200], respectively.

FIGURE 11. Effect of leaving non-essential businesses open in Washington

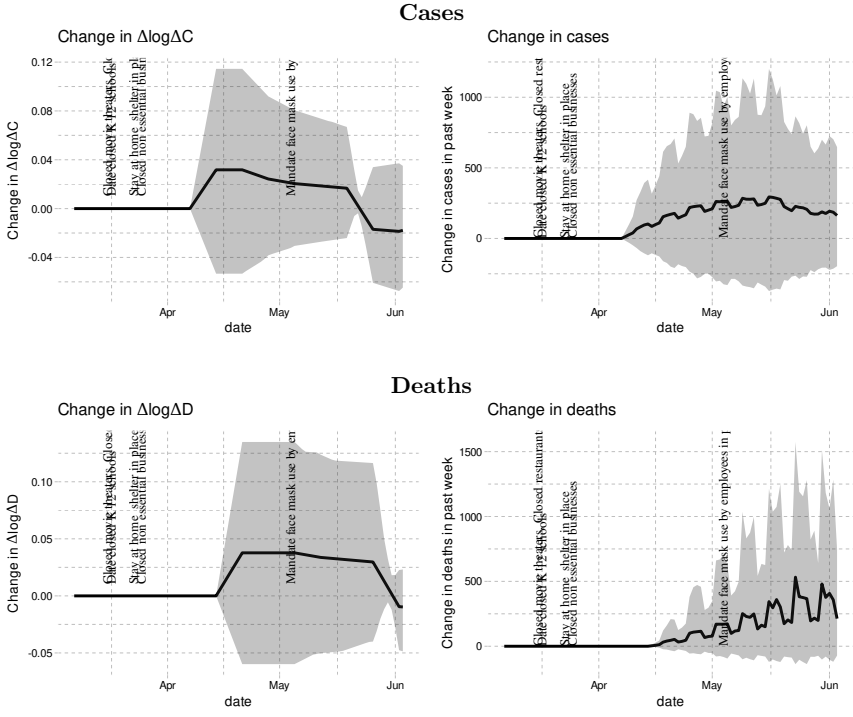

Figure 12 shows the national effect of leaving non essential businesses open on cases and deaths. For cases, the estimates imply that with non-essential businesses open, cases would be about -15 to 60% higher in late May. The results for deaths are similar but less precise.

5.3. **Stay-at-home orders.** We next examine a counterfactual where stay-at-home orders had been never issued. Figure 13 shows the average effect of no stay-at-home orders. On average, without stay-at-home orders, case growth rate would have been nearly 0.1 higher in late April. This translates to 80% [25%,170%] more cases by the start of June. The

FIGURE 12. Effect of leaving non-essential businesses open in the US

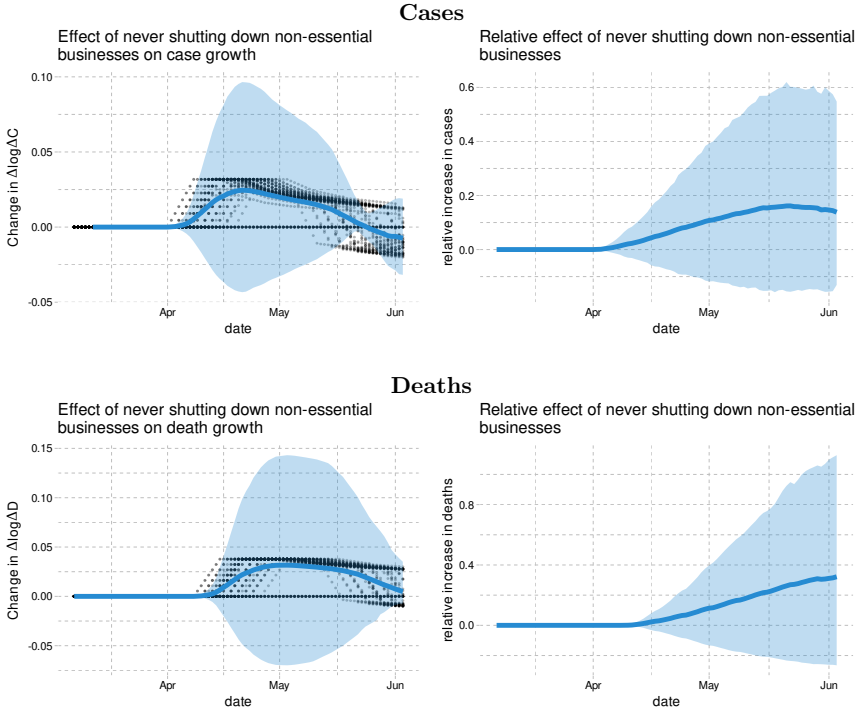

results for deaths are similar, but slightly less precise, with no increase included in a 90 percent confidence interval.

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

FIGURE 13. Effect of having no stay-at-home orders in the US

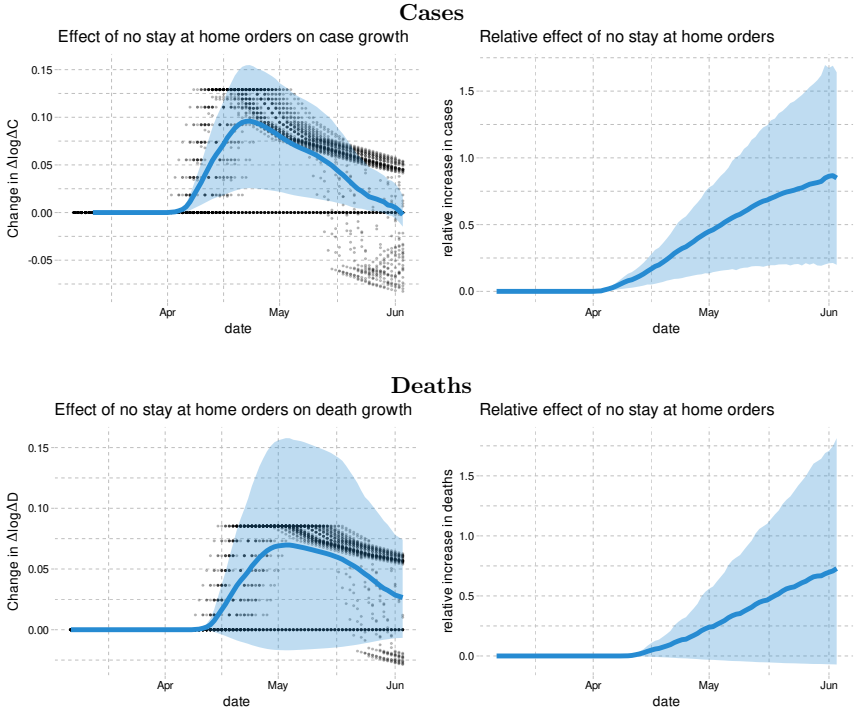

6. Counterfactual Effect of Removing All Policies and its Sensitivity

We now consider the impact of changing from the observed policies to none. Figure 15 shows the average across states of the change in case growth and relative increase in cases under a specification without past national case variables. Removing policies leads to an increase of above 0.2 in case growth throughout April and May. The confidence interval is fairly wide, and its upper bound includes a very large increase in cases by the end of May. The right panel displays the national increase in aggregate cases without any policy intervention. The estimates imply at least a 7 fold increase in cases with a large upper bound by the end of May, or at least 14 million additional cases. The estimated impact on deaths is larger than cases, and even more imprecise.

The effect of removing all policies includes the effect of school closures. The visual evidence on growth rates for states with and without school closures, presented blow, suggest that there may be a potentially large effect, though the history is very short. The main results presented in Section 3 also support the hypothesis that the school closures were

FIGURE 14. Case and death growth conditional on policies

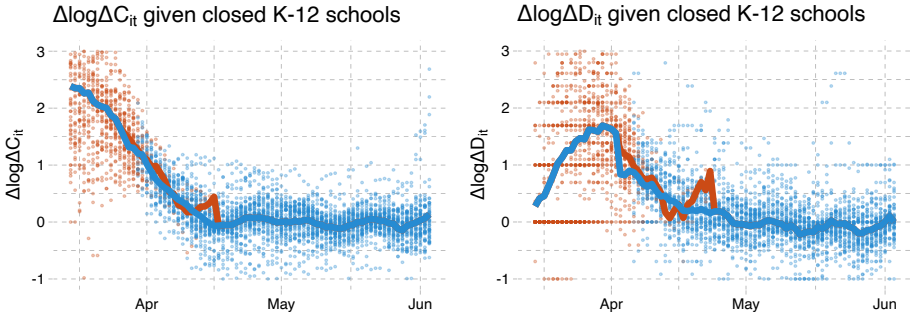

In these figures, red points are the case or death growth rate in states without each policy 14 (or 21 for deaths) days earlier. Blue points are states with each policy 14 (or 21 for deaths) days earlier. The red line is the average across states without each policy. The blue line is the average across states with each policy.

important at lowering the growth rates. This evidence is consistent with the emerging evidence of prevalence of Covid-19 among children aged 10-17. Davies et al. (2020) find that although children's transmission and susceptibility rates are half that of ages 20-30, children's contact rates are much higher. This type of evidence, as well as, evidence that children carry viral loads similar to older people (Jones et al. (2020)), led Germany to make the early decision of closing schools.

As discussed above, there is little variation across states in the timing of school closures. Consequently, the effect of school closures is difficult to identify statistically, because it is hard to separate it from aggregate time effect, and its estimate is sensitive to an inclusion of some aggregate variables such as national cases. To support this point, Figure 16 shows the effect of removing all policies on cases based on the estimates with national cases included as information. When national case variables are included in the specification, the estimated effect of school closures, and hence that of removing all policies, is much smaller with a 90% confidence interval of [0,10] fold increases.

Given this sensitivity, we conclude that there still exists a lot of uncertainty as to the effect of removing all policies, especially schooling. The impact of not implementing any policies on cases and deaths can be quite large, but the effect of school closures, hence that of removing all policies, is not well identified statistically from the US state-level data alone, because of the lack of cross-sectional variations. Any analyses of re-opening plans need to be aware of this uncertainty. An important research question is how to resolve this uncertainty using additional data sources.

Covid Economics 35, 7 July 2020: 116-175

FIGURE 15. Effect of removing policies on cases in the US under a specification with only state-level cases/deaths as information

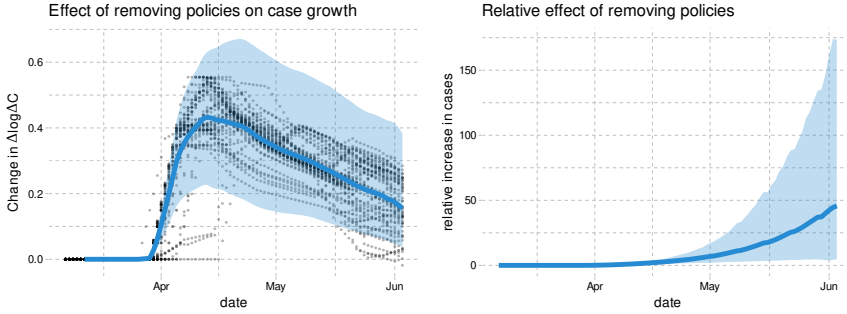

FIGURE 16. Effect of removing policies on cases in the US under a specification with both state-level cases/deaths and national-level cases/deaths as information

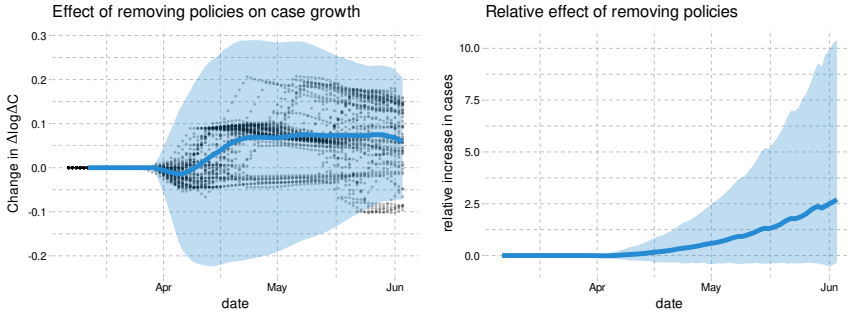

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

7. Conclusion

This paper assesses the effects of policies on the spread of Covid-19 in the US using state-level data on cases, tests, policies, and social distancing behavior measures from Google Mobility Reports. Our findings are summarized as follows.

First, our empirical analysis indicates that mandating face masks has reduced the spread of Covid-19 without affecting people's social distancing behavior measured by Google Mobility Reports. Our counterfactual experiment based on the estimated model suggests that if all states had have adopted mandatory face mask policies on April 1st of 2020, then the number of deaths by the end of May would have been smaller by as much as 17 to 55%, which roughly translates to 17 to 55 thousand saved lives.

Second, we find that keeping non-essential businesses open would have led to -20 to 60% more cases while not implementing stay-at-home orders would have increased cases by 25 to 170 % by the start of June.

Third, we find considerable uncertainty over the impact of all policies combined on case or death growth because it is difficult to identify the effect of school closures from the US state-level data due to the lack of variation in the timing of school closures across states.

Fourth, our analysis shows that people voluntarily reduce their visits to workplace, retails, grocery stores, and limit their use of public transit when they receive information on a higher number of new cases and deaths. This suggests that individuals make decisions to voluntarily limit their contact with others in response to greater transmission risks, leading to an important feedback mechanism that affects future cases and deaths. Model simulations that ignore this voluntary behavioral response to information on transmission risks would over-predict the future number of cases and deaths.

Beyond these findings, our paper presents a useful conceptual framework to investigate the relative roles of policies and information on determining the spread of Covid-19 through their impact on people's behavior. Our causal model allows us to explicitly define counterfactual scenarios to properly evaluate the effect of alternative policies on the spread of Covid-19. More broadly, our causal framework can be useful for quantitatively analyzing not only health outcomes but also various economic outcomes (Bartik et al., 2020; Chetty et al., 2020).

References

Abaluck, Jason, Judith A. Chevalier, Nicholas A. Christakis, Howard Paul Forman, Edward H. Kaplan, Albert Ko, and Sten H. Vermund. 2020. "The Case for Universal Cloth Mask Adoption and Policies to Increase Supply of Medical Masks for Health Workers." *Covid Economics,* 5.

Abouk, Rahi and Babak Heydari. 2020. "The Immediate Effect of COVID-19 Policies on Social Distancing Behavior in the United States." *medRxiv* URL `https://www.medrxiv.org/content/early/2020/04/28/2020.04.07.20057356`.

Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D Whinston. 2020. "Optimally Targeted Lockdowns in a Multi-Group SIR model." Tech. rep., National Bureau of Economic Research.

Adda, Jérôme. 2016. "Economic Activity and the Spread of Viral Diseases: Evidence from High Frequency Data*." *The Quarterly Journal of Economics* 131 (2):891–941.

Alvarez, Fernando E, David Argente, and Francesco Lippi. 2020. "A simple planning problem for covid-19 lockdown." *Covid Economics,* 14.

Andersen, Martin. 2020. "Early Evidence on Social Distancing in Response to COVID-19 in the United States." Tech. rep., UNC Greensboro.

Atkeson, Andrew. 2020a. "How Deadly is COVID-19? Understanding the Difficulties with Estimation of its Fatality Rate." Tech. rep., National Bureau of Economic Research.

———. 2020b. "What Will be the Economic Impact of COVID-19? Rough Estimates of Disease Scenarios." Tech. rep., National Bureau of Economic Research.

Avery, Christopher, William Bossert, Adam Clark, Glenn Ellison, and Sara Fisher Ellison. 2020. "Policy implications of models of the spread of Coronavirus: Perspectives and opportunities for economists." *Covid Economics,* 12.

Baqaee, David, Emmanuel Farhi, Michael J Mina, and James H Stock. 2020. "Reopening Scenarios." Tech. rep., National Bureau of Economic Research.

Bartik, Alexander, Marianne Bertrand, Feng Lin, Jesse Rothstein, and Matthew Unrath. 2020. "Measuring the Labor Market at the Onset of the COVID-19 Crisis." Tech. rep. URL `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3633053`.

Cereda, D, M Tirani, F Rovida, V Demicheli, M Ajelli, P Poletti, F Trentini, G Guzzetta, V Marziano, A Barone, M Magoni, S Deandrea, G Diurno, M Lombardo, M Faccini, A Pan, R Bruno, E Pariani, G Grasselli, A Piatti, M Gramegna, F Baldanti, A Melegaro, and S Merler. 2020. "The early phase of the COVID-19 outbreak in Lombardy, Italy."

Chen, Shuowen, Victor Chernozhukov, and Iván Fernández-Val. 2019. "Mastering panel metrics: causal impact of democracy on growth." In *AEA Papers and Proceedings*, vol. 109. 77–82.

Chetty, Raj, John N Friedman, Nathaniel Hendren, and Michael Stepner. 2020. "Real-Time Economics: A New Platform to Track the Impacts of COVID-19 on People, Businesses, and Communities Using Private Sector Data." Tech. rep., Mimeo.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2020. "Labor markets during the covid-19 crisis: A preliminary view." *Covid Economics,* 21.

Courtemanche, Charles, Joseph Garuccio, Anh Le, Joshua Pinkston, and Aaron Yelowitz. 2020. "Strong Social Distancing Measures In The United States Reduced The COVID-19 Growth Rate." *Health Affairs* :10.1377/hlthaff.2020.00608URL `https://doi.org/10.1377/hlthaff.2020.00608`.

Davies, Nicholas G, Petra Klepac, Yang Liu, Kiesha Prem, Mark Jit, Rosalind M Eggo, CMMID COVID-19 working group et al. 2020. "Age-dependent effects in the transmission and control of COVID-19 epidemics." *MedRxiv* .

Ferguson, Neil, Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, Adhiratha Boonyasiri, Zulma Cucunubá, Gina Cuomo-Dannenburg, Amy Dighe, Ilaria Dorigatti, Han Fu, Katy Gaythorpe, Will Green, Arran Hamlet, Wes Hinsley, Lucy C Okell, Sabine van Elsland, Hayley Thompson, Robert Verity, Erik Volz, Haowei Wang, Yuanrong Wang, Patrick GT Walker, Caroline Walters, Peter Winskill, Charles Whittaker, Christl A Donnelly, Steven Riley, and Azra C Ghani. 2020. "Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand." Tech. rep., Imperial College London.

Fernández-Villaverde, Jesús and Charles I Jones. 2020. "Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities." Working Paper 27128, National Bureau of Economic Research. URL `http://www.nber.org/papers/w27128`.

Gitmez, Arda, Konstantin Sonin, and Austin L. Wright. 2020. "Political Economy of Crisis Response." Tech. rep., University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2020-68.

Greenhalgh, Trisha, Manuel B Schmid, Thomas Czypionka, Dirk Bassler, and Laurence Gruer. 2020. "Face masks for the public during the covid-19 crisis." *BMJ* 369. URL `https://www.bmj.com/content/369/bmj.m1435`.

Greenland, Sander, Judea Pearl, and James M Robins. 1999. "Causal diagrams for epidemiologic research." *Epidemiology* :37–48.

Gupta, Sumedha, Thuy D Nguyen, Felipe Lozano Rojas, Shyam Raman, Byungkyu Lee, Ana Bento, Kosali I Simon, and Coady Wing. 2020. "Tracking Public and Private Responses to the COVID-19 Epidemic:

Covid Economics 35, 7 July 2020: 116-175

Evidence from State and Local Government Actions." Working Paper 27027, National Bureau of Economic Research. URL http://www.nber.org/papers/w27027.

Haavelmo, Trygve. 1944. "The probability approach in econometrics." *Econometrica: Journal of the Econometric Society* :iii–115.

Harris, Jeffrey E. 2020. "The Subways Seeded the Massive Coronavirus Epidemic in New York City." Working Paper 27021, National Bureau of Economic Research. URL http://www.nber.org/papers/w27021.

Heckman, James J. and Edward J. Vytlacil. 2007. "Econometric Evaluation of Social Programs, Part I: Causal Models, Structural Models and Econometric Policy Evaluation." In *Handbook of Econometrics*, *Handbook of Econometrics*, vol. 6, edited by J.J. Heckman and E.E. Leamer, chap. 70. Elsevier. URL https://ideas.repec.org/h/eee/ecochp/6b-70.html.

Hernán, M.A. and J.M. Robins. 2020. *Causal Inference: What If*. Chapman & Hall/CRC.

Hou, Yixuan J., Kenichi Okuda, Caitlin E. Edwards, David R. Martinez, Takanori Asakura, Kenneth H. Dinnon III, Takafumi Kato, Rhianna E. Lee, Boyd L. Yount, Teresa M. Mascenik, Gang Chen, Kenneth N. Olivier, Andrew Ghio, Longping V. Tse, Sarah R. Leist, Lisa E. Gralinski, Alexandra Schäfer, Hong Dang, Rodney Gilmore, Satoko Nakano, Ling Sun, M. Leslie Fulcher, Alessandra Livraghi-Butrico, Nathan I. Nicely, Mark Cameron, Cheryl Cameron, David J. Kelvin, Aravinda de Silva, David M. Margolis, Alena Markmann, Luther Bartelt, Ross Zumwalt, Fernando J. Martinez, Steven P. Salvatore, Alain Borczuk, Purushothama R. Tata, Vishwaraj Sontake, Adam Kimple, Ilona Jaspers, Wanda K. O'Neal, Scott H. Randell, Richard C. Boucher, and Ralph S. Baric. 2020. "SARS-CoV-2 Reverse Genetics Reveals a Variable Infection Gradient in the Respiratory Tract." *Cell* URL https://doi.org/10.1016/j.cell.2020.05.042.

Howard, Jeremy, Austin Huang, Zhiyuan Li, Zeynep Tufekci, Vladimir Zdimal, Helene-Mari van der Westhuizen, Arne von Delft, Amy Price, Lex Fridman, Lei-Han Tang, Viola Tang, Gregory Watson, Christina Bax, Reshama Shaikh, Frederik Questier, Danny Hernandez, Larry Chu, Christina Ramirez, and Anne Rimoin. 2020. "Face Masks Against COVID-19: An Evidence Review." URL https://doi.org/10.20944/preprints202004.0203.v1.

Hsiang, Solomon, Daniel Allen, Sebastien Annan-Phan, Kendon Bell, Ian Bolliger, Trinetta Chong, Hannah Druckenmiller, Andrew Hultgren, Luna Yue Huang, Emma Krasovich, Peiley Lau, Jaecheol Lee, Esther Rolf, Jeanette Tseng, and Tiffany Wu. 2020. "The Effect of Large-Scale Anti-Contagion Policies on the Coronavirus (COVID-19) Pandemic." *medRxiv* URL https://www.medrxiv.org/content/early/2020/04/29/2020.03.22.20040642.

Imbens, Guido W. and Donald B. Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press.

Jones, Terry C, Barbara Mühlemann, Talitha Veith, Guido Biele, Marta Zuchowski, Jörg Hoffmann, Angela Stein, Anke Edelmann, Victor Max Corman, and Christian Drosten. 2020. "An analysis of SARS-CoV-2 viral load by patient age." *medRxiv* URL https://www.medrxiv.org/content/early/2020/06/09/2020.06.08.20125484.

Keppo, Juusi, Elena Quercioli, Mariana Kudlyak, Lones Smith, and Andrea Wilson. 2020. "The behavioral SIR model, with applications to the Swine Flu and COVID-19 pandemics." In *Virtual Macro Seminar*.

Li, Qun, Xuhua Guan, Peng Wu, Xiaoye Wang, Lei Zhou, Yeqing Tong, Ruiqi Ren, Kathy S.M. Leung, Eric H.Y. Lau, Jessica Y. Wong, Xuesen Xing, Nijuan Xiang, Yang Wu, Chao Li, Qi Chen, Dan Li, Tian Liu, Jing Zhao, Man Liu, Wenxiao Tu, Chuding Chen, Lianmei Jin, Rui Yang, Qi Wang, Suhua Zhou, Rui Wang, Hui Liu, Yinbo Luo, Yuan Liu, Ge Shao, Huan Li, Zhongfa Tao, Yang Yang, Zhiqiang Deng, Boxi Liu, Zhitao Ma, Yanping Zhang, Guoqing Shi, Tommy T.Y. Lam, Joseph T. Wu, George F. Gao, Benjamin J. Cowling, Bo Yang, Gabriel M. Leung, and Zijian Feng. 2020. "Early Transmission Dynamics in Wuhan, China, of Novel Coronavirus–Infected Pneumonia." *New England Journal of Medicine* 382 (13):1199–1207. URL https://doi.org/10.1056/NEJMoa2001316. PMID: 31995857.

Linton, Natalie M., Tetsuro Kobayashi, Yichi Yang, Katsuma Hayashi, Andrei R. Akhmetzhanov, Sungmok Jung, Baoyin Yuan, Ryo Kinoshita, and Hiroshi Nishiura. 2020. "Incubation Period and Other Epidemiological Characteristics of 2019 Novel Coronavirus Infections with Right Truncation: A Statistical Analysis of Publicly Available Case Data." *Journal of Clinical Medicine* 9 (2):538. URL http://dx.doi.org/10.3390/jcm9020538.

LLC, Google. 2020. ""Google COVID-19 Community Mobility Reports"." URL https://www.google.com/covid19/mobility/.

Covid Economics 35, 7 July 2020: 116-175

Maloney, William F. and Temel Taskin. 2020. "Determinants of Social Distancing and Economic Activity during COVID-19: A Global View." *Covid Economics,* 13.

McAdams, David. 2020. "Nash SIR: An Economic-Epidemiological Model of Strategic Behavior During a Viral Epidemic." *Covid Economics,* 16.

MIDAS. 2020. "MIDAS 2019 Novel Coronavirus Repository: Parameter Estimates." URL `https://github.com/midas-network/COVID-19/tree/master/parameter_estimates/2019_novel_coronavirus`.

Mitze, Timo, Reinhold Kosfeld, Johannes Rode, and Klaus Wälde. 2020. "Face Masks Considerably Reduce Covid-19 Cases in Germany." *Covid Economics,* 27.

Neyman. 1925. "On the application of probability theory to agricultural experiments. Essay on principles. Section 9." *Statistical Science (1990)* :465–472.

Pei, Sen, Sasikiran Kandula, and Jeffrey Shaman. 2020. "Differential Effects of Intervention Timing on COVID-19 Spread in the United States." *medRxiv* URL `https://www.medrxiv.org/content/early/2020/05/20/2020.05.15.20103655`.

Peters, Jonas, Dominik Janzing, and Sch'olkopf Bernhard. 2017. *Elements of causal inference: foundations and learning algorithms.* Mass.

Raifman, Julia, Kristen Nocka, David Jones, Jacob Bor, Sarah Ketchen Lipson, Jonathan Jay, Philip Chan, Megan Cole Brahim, Carolyn Hoffman, Claire Corkish, Elizabeth Ferrara, Elizabeth Long, Emily Baroni, Faith Contador, Hannah Simon, Morgan Simko, Rachel Scheckman, Sarah Brewer, Sue Kulkarni, Felicia Heykoop, Manish Patel, Aishwarya Vidyasagaran, Andrew Chiao, Cara Safon, and Samantha Burkhart. 2020. "COVID-19 US state policy database." URL `https://tinyurl.com/statepolicies`.

Rubin, Donald B. 1974. "Estimating causal effects of treatments in randomized and nonrandomized studies." *Journal of educational Psychology* 66 (5):688.

Sanche, Steven, Yen Ting Lin, Chonggang Xu, Ethan Romero-Severson, Nick Hengartner, and Ruian Ke. 2020. "High Contagiousness and Rapid Spread of Severe Acute Respiratory Syndrome Coronavirus 2." *Emerging Infectious Diseases* 26 (7).

Siordia, Juan A. 2020. "Epidemiology and clinical features of COVID-19: A review of current literature." *Journal of Clinical Virology* 127:104357. URL `http://www.sciencedirect.com/science/article/pii/S1386653220300998`.

Stock, James H. 2020a. "Data Gaps and the Policy Response to the Novel Coronavirus." *Covid Economics,* 3.

———. 2020b. "Reopening the Coronavirus-Closed Economy." Tech. rep., Hutchins Center Working Paper #60.

Strotz, Robert H and Herman OA Wold. 1960. "Recursive vs. nonrecursive systems: An attempt at synthesis (part i of a triptych on causal chain systems)." *Econometrica: Journal of the Econometric Society* :417–427.

Tian, Liang, Xuefei Li, Fei Qi, Qian-Yuan Tang, Viola Tang, Jiang Liu, Zhiyuan Li, Xingye Cheng, Xuanxuan Li, Yingchen Shi, Haiguang Liu, and Lei-Han Tang. 2020. "Calibrated Intervention and Containment of the COVID-19 Pandemic."

Tinbergen, J. 1930. "Determination and interpretation of supply curves: an example." Tech. rep., Zeitschrift fur Nationalokonomie.

Wheaton, William C. and Anne Kinsella Thompson. 2020. "The Geography of Covid-19 Growth in the US: Counties and Metropolitan Areas." Tech. rep., MIT.

White House, The. 2020. "Guidelines for Opening Up America Again." URL `https://www.whitehouse.gov/openingamerica/`.

Wright, Austin L., Konstantin Sonin, Jesse Driscoll, and Jarnickae Wilson. 2020. "Poverty and Economic Dislocation Reduce Compliance with COVID-19 Shelter-in-Place Protocols." *SSRN Electronic Journal* .

Wright, Philip G. 1928. *Tariff on animal and vegetable oils.* Macmillan Company, New York.

Zhang, Juanjuan, Maria Litvinova, Wei Wang, Yan Wang, Xiaowei Deng, Xinghui Chen, Mei Li, Wen Zheng, Lan Yi, Xinhua Chen, Qianhui Wu, Yuxia Liang, Xiling Wang, Juan Yang, Kaiyuan Sun, Ira M Longini, M Elizabeth Halloran, Peng Wu, Benjamin J Cowling, Stefano Merler, Cecile Viboud, Alessandro Vespignani, Marco Ajelli, and Hongjie Yu. 2020a. "Evolving epidemiology and transmission dynamics of coronavirus disease 2019 outside Hubei province, China: a descriptive and modelling study." *The Lancet Infectious Diseases* URL `http://www.sciencedirect.com/science/article/pii/S1473309920302309`.

Covid Economics 35, 7 July 2020: 116-175

Zhang, Renyi, Yixin Li, Annie L. Zhang, Yuan Wang, and Mario J. Molina. 2020b. "Identifying airborne transmission as the dominant route for the spread of COVID-19." *Proceedings of the National Academy of Sciences* URL https://www.pnas.org/content/early/2020/06/10/2009637117.

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

Appendix A. Data Construction

A.1. **Measuring ΔC and $\Delta \log \Delta C$.** We have three data sets with information on daily cumulative confirmed cases in each state. As shown in Table 8, these cumulative case numbers are very highly correlated. However, Table 9 shows that the numbers are different more often than not.

	NYT	JHU	CTP
NYT	1.00000	0.99995	0.99991
JHU	0.99995	1.00000	0.99987
CTP	0.99991	0.99987	1.00000

Table 8. Correlation of cumulative cases

	1	2	3
NYT	1.00	0.28	0.37
JHU	0.28	1.00	0.33
CTP	0.37	0.33	1.00

Table 9. Portion of cumulative cases that are equal between data sets

Figure 17 shows the evolution of new cases in each of these three datasets. In all cases, daily changes in cumulative cases displays some excessive volatility. This is likely due to delays and bunching in testing and reporting of results. Table 10 shows the variance of log new cases in each data set, as well as their correlations. As shown, the correlations are approximately 0.9. The NYT new case numbers have the lowest variance.[28] In our subsequent results, we will primarily use the case numbers from The New York Times.

	NYT	JHU	CTP
NYT	1.00	0.88	0.87
JHU	0.88	1.00	0.80
CTP	0.87	0.80	1.00
Variance	5.63	7.02	6.64

Table 10. Correlation and variance of log daily new cases

For most of our results, we focus on new cases in a week instead of in a day. We do this for two reasons as discussed in the main text. First, a decline of new cases over two weeks has become a key metric for decision makers. Secondly, aggregating to weekly new cases smooths out the noise associated with the timing of reporting and testing.

Table 11 reports the correlation and variance of weekly log new cases across the three data sets. Figure 18 shows the evolution of weekly new cases in each state over time.

[28]This comparison is somewhat sensitive to how you handle negative and zero cases when taking logs. Here, we replaced $\log(0)$ with -1. In our main results, we work with weekly new cases, which are very rarely zero.

Covid Economics 35, 7 July 2020: 116-175

FIGURE 17. Daily cases

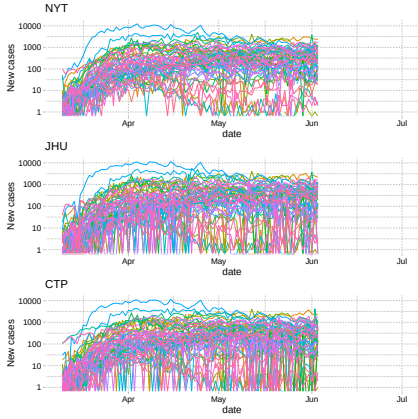

Each line shows daily new cases in a state.

	NYT	JHU	CTP
NYT	1.00	0.99	0.99
JHU	0.99	1.00	0.99
CTP	0.99	0.99	1.00
Variance	4.15	4.33	4.20

TABLE 11. Correlation and variance of log weekly new cases

A.2. **Deaths.** Table 12 reports the correlation and variance of weekly deaths in the three data sets. Figure 19 shows the evolution of weekly deaths in each state. As with cases, we use death data from The New York Times in our main results.

	NYT	JHU	CTP
NYT	1.00	0.99	0.99
JHU	0.99	1.00	0.98
CTP	0.99	0.98	1.00
Variance	293262.32	288818.77	204037.51

TABLE 12. Correlation and variance of weekly deaths

A.3. **Tests.** Our test data comes from The Covid Tracking Project. Figure 21 shows the evolution of tests over time.

Covid Economics 35, 7 July 2020: 116-175

FIGURE 18. Weekly Cases

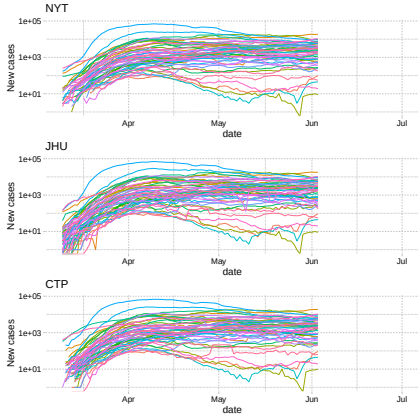

Each line shows weekly new cases in a state.

FIGURE 19. Weekly Deaths

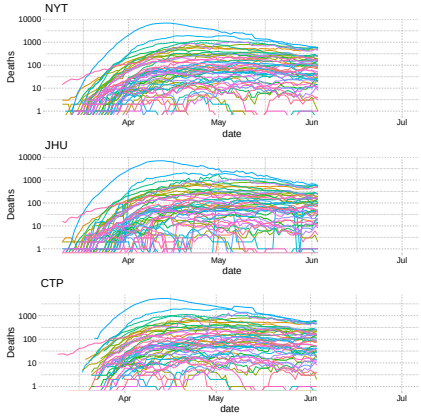

Each line shows weekly deaths in a state.

Covid Economics 35, 7 July 2020: 116-175

FIGURE 20. Case and death growth

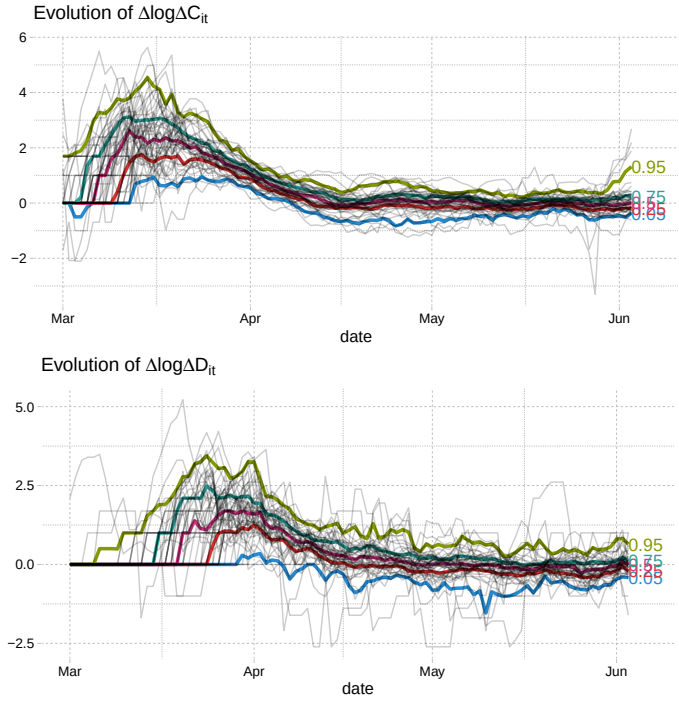

Thin gray lines are case or death growth in each state and date. Thicker colored lines are quantiles of case or death growth conditional on date.

FIGURE 21. Number of Tests

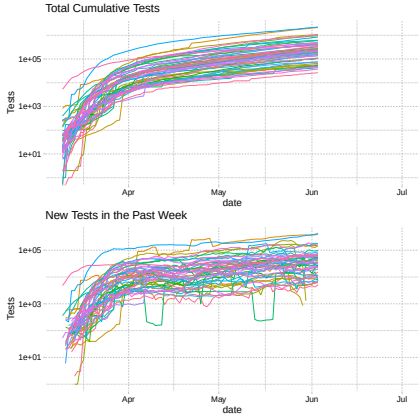

These figures use the "total test results" reported by The Covid Tracking Project. This is meant to reflect the number of people tested (as opposed to the number of specimens tested).

Covid Economics 35, 7 July 2020: 116-175

A.4. **Social Distancing Measures.** In measuring social distancing, we focus on Google Mobility Reports. This data has international coverage and is publicly available. Figure 22 shows the evolution of the four Google Mobility Reports variables that we use in our analysis.

FIGURE 22. Evolution of Google Mobility Reports

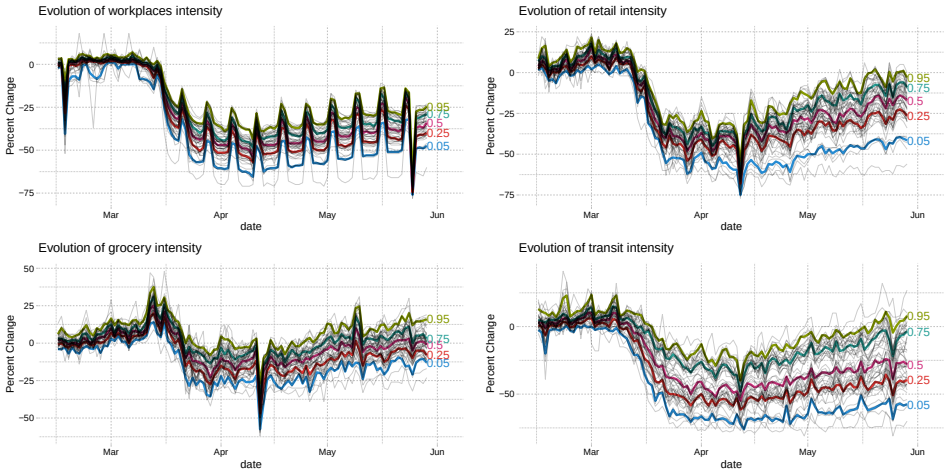

This figure shows the evolution of Google Mobility Reports over time. Thin gray lines are the value of the variables in each state and date. Thicker colored lines are quantiles of the variables conditional on date.

A.5. **Policy Variables.** We use the database on US state policies created by Raifman et al. (2020). As discussed in the main text, our analysis focuses on seven policies. For stay-at-home orders, closed nonessential businesses, closed K-12 schools, closed restaurants except takeout, and closed movie theaters, we double-checked any state for which Raifman et al. (2020) does not record a date. We filled in a few missing dates. Our modified data is available here. Our modifications fill in 1 value for school closures, 2 for stay-at-home orders, 3 for movie theater closure, and 4 for non-essential business closures. Table 13 displays all 25 dated policy variables in Raifman et al. (2020)'s database with our modifications described above.

A.6. **Timing.** There is a delay between infection and when a person is tested and appears in our case data. MIDAS (2020) maintain a list of estimates of the duration of various stages of Covid-19 infections. The incubation period, the time from infection to symptom onset, is widely believed to be 5 days. For example, using data from Wuhan, Li et al. (2020) estimate a mean incubation period of 5.2 days. Siordia (2020) reviews the literature and concludes the mean incubation period is 3-9 days.

Covid Economics 35, 7 July 2020: 116-175

FIGURE 23. Case and death growth conditional on policies

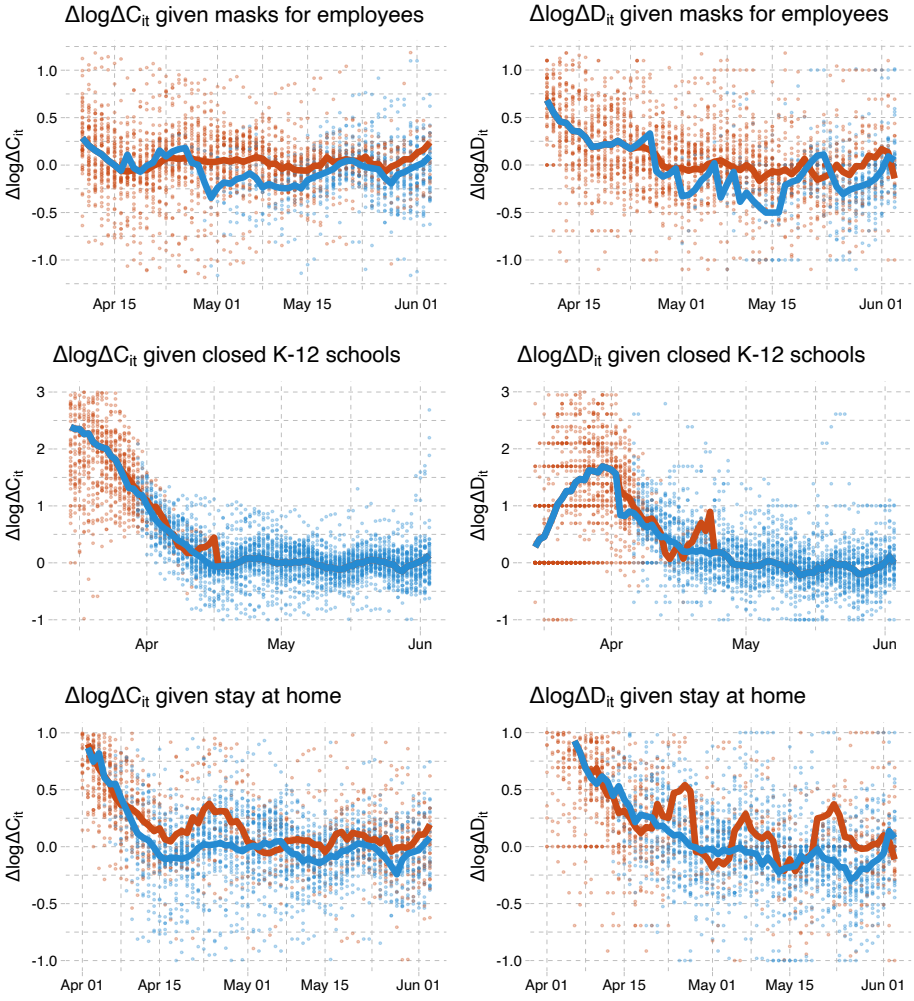

In these figures, red points are the case or death growth rate in states without each policy 14 (or 21 for deaths) days earlier. Blue points are states with each policy 14 (or 21 for deaths) days earlier. The red line is the average across states without each policy. The blue line is the average across states with each policy.

Covid Economics 35, 7 July 2020: 116-175

FIGURE 24. Case and death growth conditional on policies

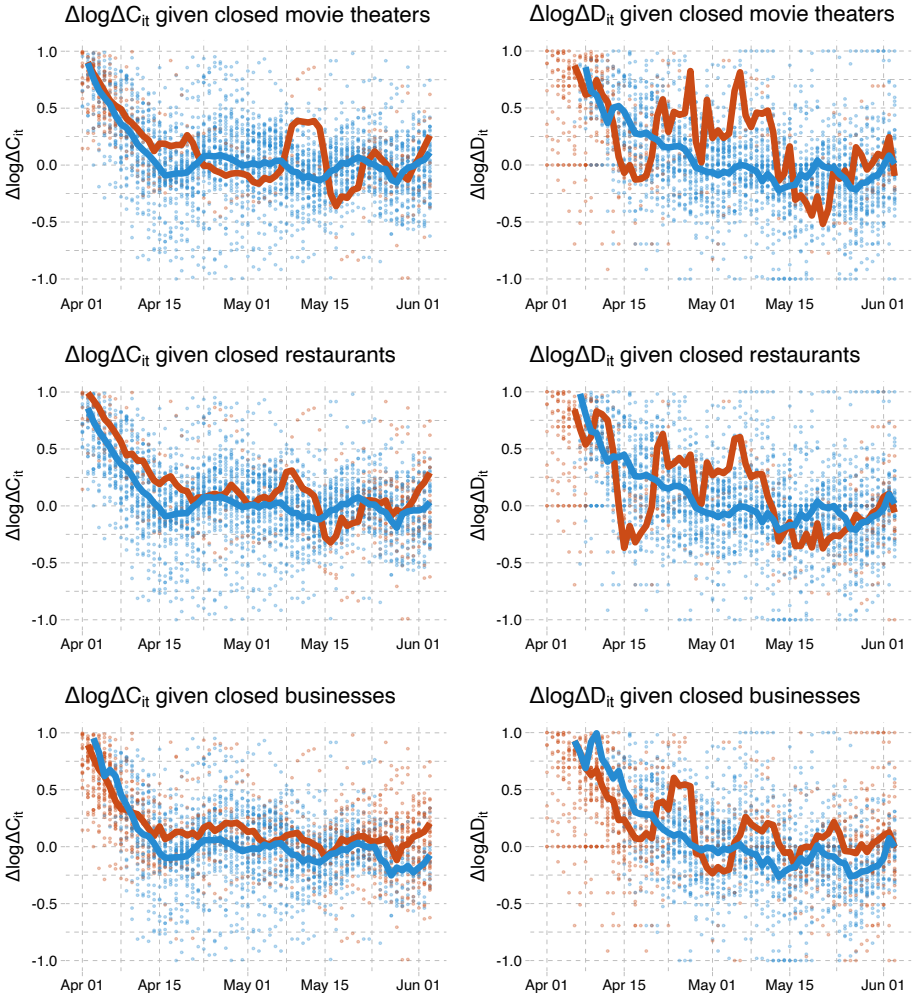

In these figures, red points are the case or death growth rate in states without each policy 14 (or 21 for deaths) days earlier. Blue points are states with each policy14 (or 21 for deaths) days earlier. The red line is the average across states without each policy. The blue line is the average across states with each policy.

Covid Economics 35, 7 July 2020: 116-175

	N	Min	Median	Max
State of emergency	51	2020-02-29	2020-03-11	2020-03-16
Date closed K 12 schools	51	2020-03-13	2020-03-17	2020-04-03
Closed day cares	15	2020-03-16	2020-03-23	2020-04-06
Date banned visitors to nursing homes	30	2020-03-09	2020-03-16	2020-04-06
Closed non essential businesses	43	2020-03-19	2020-03-25	2020-04-06
Closed restaurants except take out	48	2020-03-15	2020-03-17	2020-04-03
Closed gyms	46	2020-03-16	2020-03-20	2020-04-03
Closed movie theaters	49	2020-03-16	2020-03-21	2020-04-06
Stay at home shelter in place	42	2020-03-19	2020-03-28	2020-04-07
End relax stay at home shelter in place	33	2020-04-24	2020-05-15	2020-06-05
Began to reopen businesses statewide	49	2020-04-20	2020-05-07	2020-05-29
Reopen restaurants	41	2020-04-24	2020-05-15	2020-06-01
Reopened gyms	31	2020-04-24	2020-05-16	2020-06-01
Reopened movie theaters	19	2020-04-27	2020-05-15	2020-06-01
Resumed elective medical procedures	35	2020-04-20	2020-04-30	2020-05-29
Mandate face mask use by all individuals in public spaces	17	2020-04-08	2020-04-20	2020-05-29
Mandate face mask use by employees in public facing businesses	39	2020-04-03	2020-05-01	2020-06-01
Stop Initiation of Evictions overall or due to COVID related issues	24	2020-03-16	2020-03-24	2020-04-20
Stop enforcement of evictions overall or due to COVID related issues	26	2020-03-15	2020-03-23	2020-04-20
Renter grace period or use of security deposit to pay rent	2	2020-04-10	2020-04-17	2020-04-24
Order freezing utility shut offs	34	2020-03-12	2020-03-19	2020-04-13
Froze mortgage payments	1	2020-03-21	2020-03-21	2020-03-21
Waived one week waiting period for unemployment insurance	36	2020-03-08	2020-03-18	2020-04-06

TABLE 13. State Policies

Estimates of the time between symptom onset and case reporting or death are less common. Using Italian data, Cereda et al. (2020) estimate an average of 7.3 days between symptom onset and reporting. Zhang et al. (2020a) find an average of 7.4 days using Chinese data from December to early February, but they find this period declined from 8.9 days in January to 5.4 days in the first week of February. Both of these papers on time from symptom onset to reporting have large confidence intervals covering approximately 1 to 20 days.

Studying publicly available data on infected persons diagnosed outside of Wuhan, Linton et al. (2020) estimate an average of 15 days from onset to death. Similarly, using publicly available reports of 140 confirmed Covid-19 cases in China, mostly outside Hubei Province, Sanche et al. (2020) estimate the time from onset to death to be 16.1 days.

Based on the above, we expect a delay of roughly two weeks between changes in behavior or policies, and changes in reported cases while a corresponding delay of roughly three weeks for deaths.

A.7. **Counterfactuals for Massachusetts and Illinois.** Figures 25 and 26 present the fit of estimated cases as well as the counterfactual effect of mandating masks on April 1st in Massachusetts and Illinois, respectively. Figures 27 and 28 show the counterfactual effect of leaving non-essential business open in Massachusetts and Illinois, respectively.

Covid Economics 35, 7 July 2020: 116-175

FIGURE 25. Effect of mandating masks on April 1st in Massachusetts

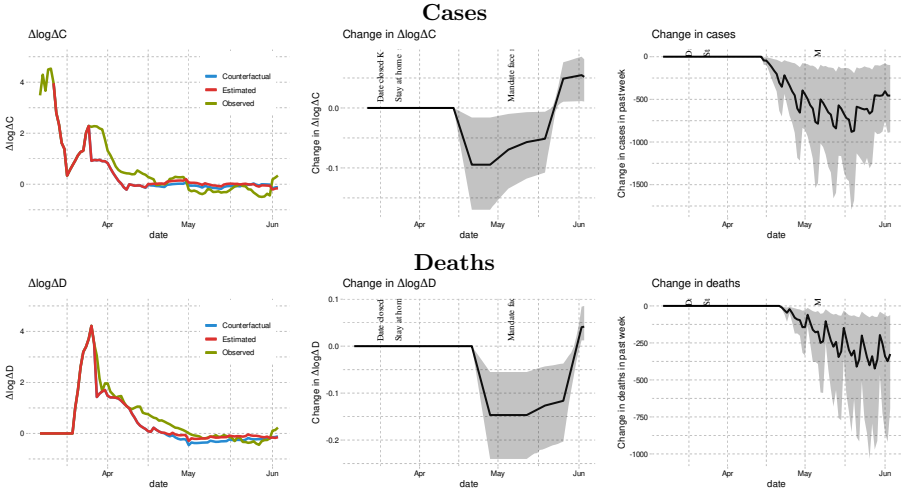

To compute the estimated and counterfactual paths we use the average of two estimated coefficients as reported in column "Average" of Table 7. We set initial $\Delta \log \Delta C$ and $\log \Delta C$ to their values first observed in the state we are simulating. We hold all other regressors at their observed values. Error terms are drawn with replacement from the residuals. We do this many times and report the average over draws of the residuals. The shaded region is a point-wise 90% confidence interval.

Covid Economics 35, 7 July 2020: 116-175

Covid Economics 35, 7 July 2020: 116-175

FIGURE 26. Effect of mandating masks on April 1st in Illinois

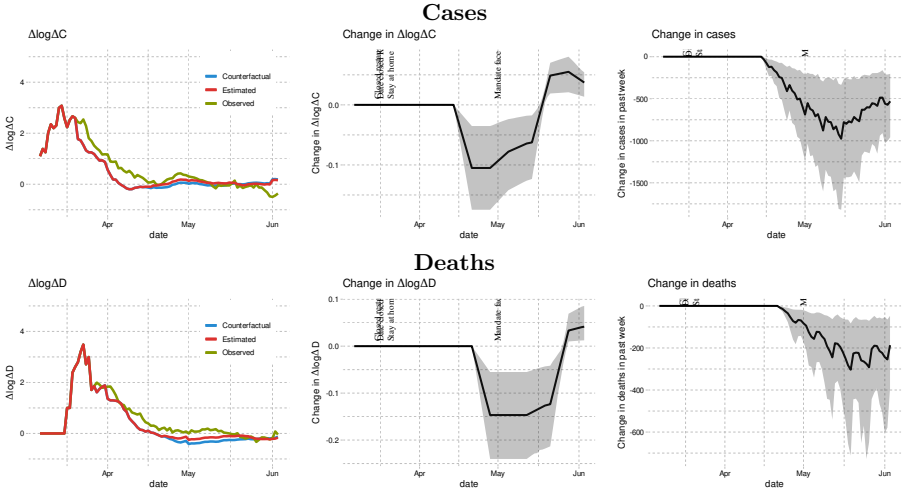

To compute the estimated and counterfactual paths we use the average of two estimated coefficients as reported in column "Average" of Table 7. We set initial $\Delta \log \Delta C$ and $\log \Delta C$ to their values first observed in the state we are simulating. We hold all other regressors at their observed values. Error terms are drawn with replacement from the residuals. We do this many times and report the average over draws of the residuals. The shaded region is a point-wise 90% confidence interval.

Covid Economics 35, 7 July 2020: 116-175

FIGURE 27. Effect of leaving businesses open in Massachusetts

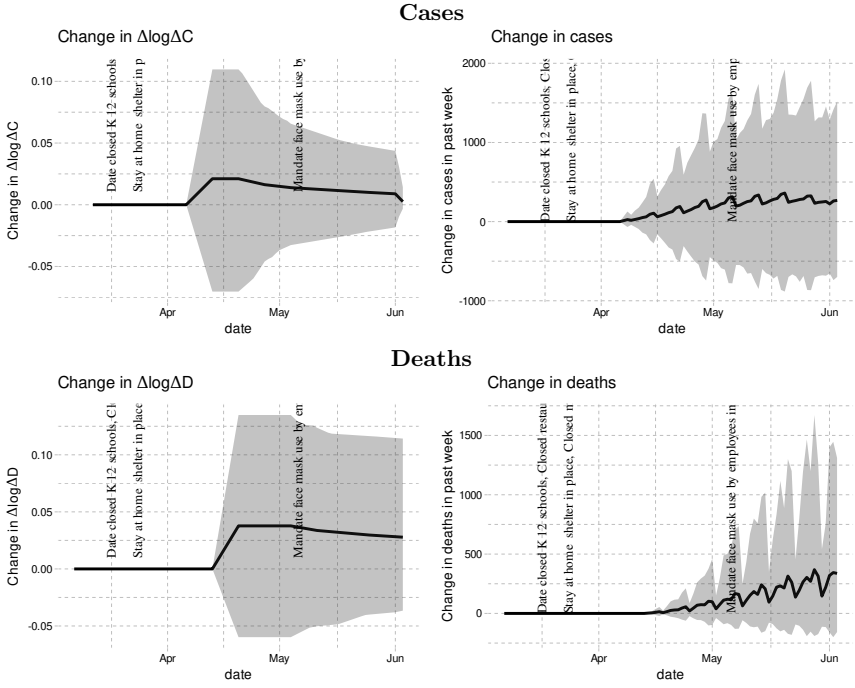

Covid Economics 35, 7 July 2020: 116-175

FIGURE 28. Effect of leaving businesses open in Illinois

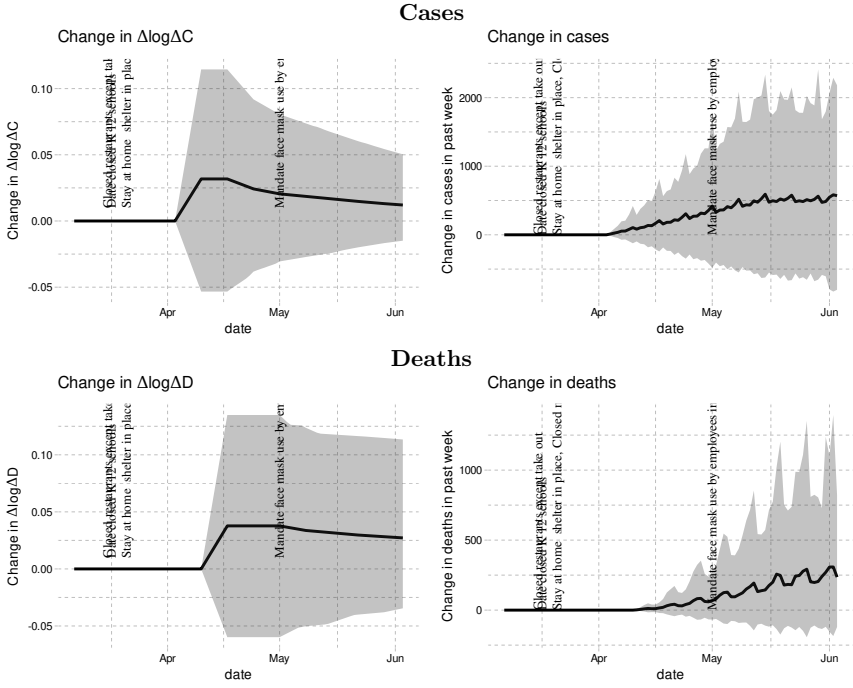

A.8. **Counterfactuals with National Cases as Information Variables.** Figures 29-31 present the results of counterfactual analyses that include the national cases/deaths as the information variables. To create this figure, we repeat the same counterfactual simulation that we did for Washington with each state. For each state, we hold national cases constant, but endogenize state specific information. Thus, these figures should be interpreted as an average of state specific counterfactuals, and not a national counterfactual.

The counterfactual results of mask policies, shelter-in-place, and closing non-essential businesses remain robust with respect to the inclusion of national case/death variables. This contrasts to the resulting counterfactual of removing all policies discussed in section 6.

Covid Economics 35, 7 July 2020: 116-175

FIGURE 29. Effect of mandating masks for employees on April 1st under a specification with both state-level cases/deaths and national-level cases/deaths as information

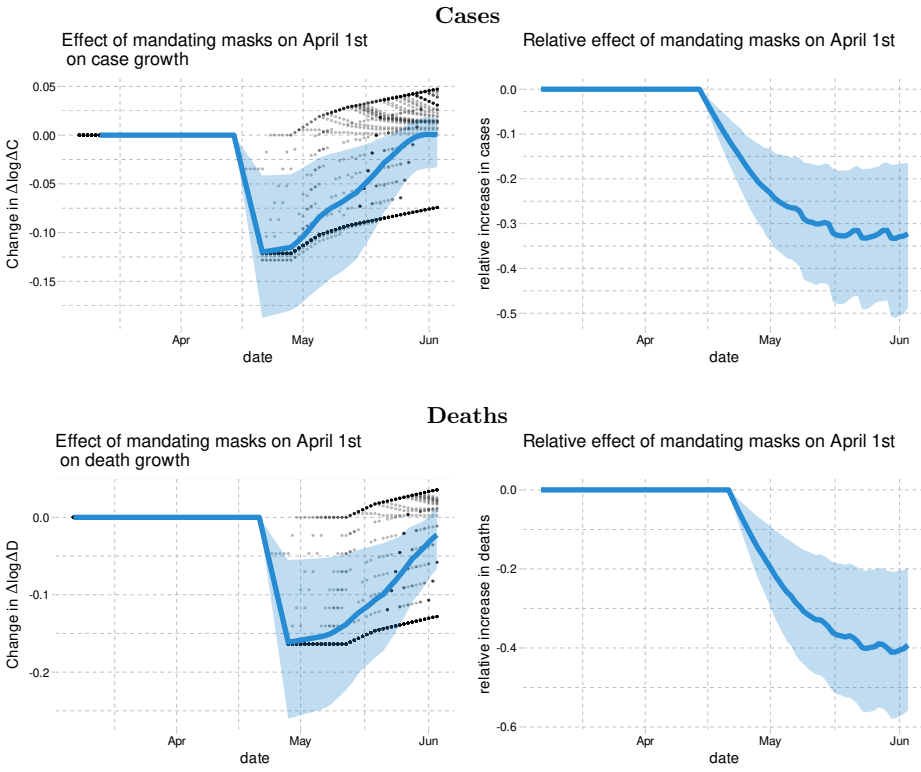

Covid Economics 35, 7 July 2020: 116-175

FIGURE 30. Effect of leaving non-essential businesses open under a specification with both state-level cases/deaths and national-level cases/deaths as information

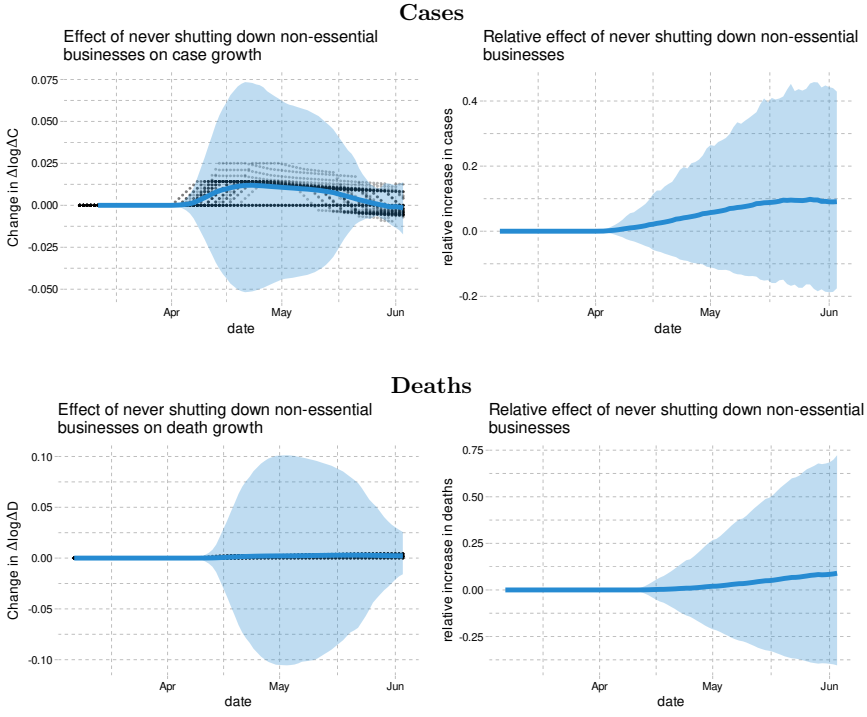

Covid Economics 35, 7 July 2020: 116-175

FIGURE 31. Effect of having no stay-at-home orders under a specification with both state-level cases/deaths and national-level cases/deaths as information

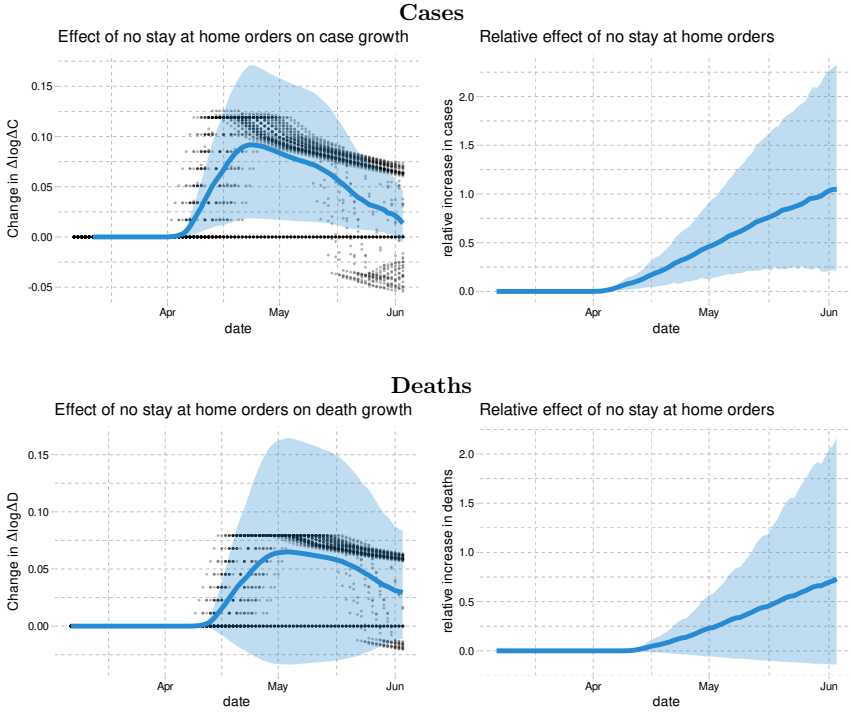

Covid Economics **Issue 35, 7 July 2020**

Transition model for coronavirus management[1]

Antoine Djogbenou,[2] Christian Gourieroux,[3] Joann Jasiak,[4] Paul Rilstone[5] and Maygol Bandehali[6]

Date submitted: 1 July 2020; Date accepted: 1 July 2020

This paper examines the individual records of patients treated for COVID-19 during the early phase of the epidemic in Ontario. We trace out daily transitions of patients through medical care of different intensity and address the right truncation in the database. We also examine the sojourn times and reveal duration dependence in the treatments for COVID-19. The transition model is used to estimate and forecast the counts of patients treated for COVID-19 in Ontario, while accommodating the right truncation and right censoring in the sample. This research is based on the Public Health Ontario (PHO) dataset from May 07, 2020.

Covid Economics 35, 7 July 2020: 176-119

1 The authors thank the Ontario Ministry of Health for information and J. Wu and V. Aguirregabiria for helpful comments. C. Gourieroux acknowledges the financial support of Agence Nationale de la Recherche (ANR-COVID) grant ANR-17-EURE-0010.

2 York University.
3 University of Toronto, Toulouse School of Economics and CREST.
4 York University.
5 York University.
6 York University.

THIS VERSION: July 6, 2020

1 Introduction

A basic epidemiological model SIRD distinguishes 4 individual states [see e.g. Vinnicky, White (2010), Yan, Chowel (2019), Toda (2020)]. State S of susceptible comprises individuals who are healthy and not yet immunized. In state I of infected, the individuals are infected and not yet recovered. An infected individual can either recover, i.e. move to state R, or decease, i.e. move to state D. This causal scheme, which can be summarized as follows:

$$S \to I \begin{array}{l} \nearrow R \\ \searrow D \end{array}$$

is a chain that involves two episodes. The first episode between states S and I concerns the propagation phase of the disease and the process of detection of infected individuals. The second episode concerns the disease monitoring that follows a diagnosed infection. In this paper, we are interested in the analysis of individual medical care histories. Therefore, we divide the state of infected into states of medical care of increasing intensity, which depend on the severity of symptoms. These additional states include Hospitalization, Intensive Care Unit (ICU), Ventilation and Intubation.

Our analysis of COVID-19 infections in Ontario is based on daily records of 18722 individuals who were diagnosed with COVID-19 over the period of 104 days between January 23 and May 05, 2020 and reported in the Public Health Ontario (PHO) database of combined records from iPHIS (integrated Public Health Information System) and CORES (Toronto Public Health Coronavirus Rapid Entry System).

The objective of our research is to introduce a modelling approach for the analysis of counts of patients under medical care that produces reasonably accurate results, given the complexity of problems encountered in the data. Some of those problems are related to the data collection method. These are, for example, the missing and misreported dates of medical treatments, the recovery dates being unavailable, or unknown outcomes of medical care. Other problems are related to the statistical analysis of data with right censoring and truncation, which is amplified by the delay in data reporting of up to 10 last days of the sampling period.

The publicly available data on COVID-19 infections in Ontario are aggregated at various levels and communicated in the PH0 reports [see, e.g. PHO (2020 a), (2020

Covid Economics 35, 7 July 2020: 176-119

b)], which contain the counts of confirmed, and deceased individuals by age or region of Ontario. The Enhanced Epidemiological Summary [PHO 2020 a)] shows additional counts of patients who are hospitalized and in the ICU. Below, we display similar series of counts computed from our sample until May 04, after a preliminary individual data adjustment for misreported data and outliers. Figures 1 and 2 show the cumulated and daily counts of Diagnosed and Deceased individuals, respectively.

Covid Economics 35, 7 July 2020: 176-119

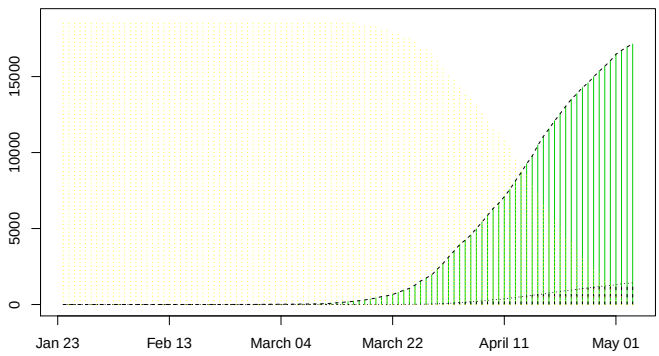

Figure 1: Counts of Undiagnosed, Diagnosed and Deceased

These figures can be compared with Figures 1 and 4 in PHO (2020 b, pages 3 and 6) reporting the confirmed cases of infections and deaths over a longer sampling period ending on May 27. The patterns prior to May 05 revealed in Figures 1 and 2 above are close to those displayed in Figures 1 and 4 of PHO (2020 b). Note that due to a reporting lag, the last few daily counts in Figures 1 and 2 need to be considered with caution.

In Figure 2, the daily increments reveal hump-shaped patterns, which resemble the curves of standard epidemiological SIRD models. The dates of peaks are, approximately,

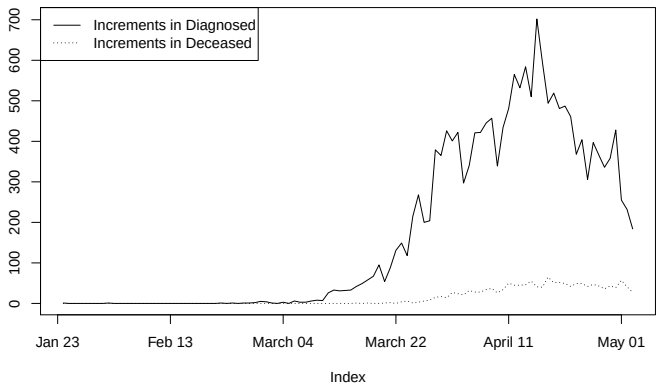

Figure 2: Daily Counts of Diagnosed and Deceased

April 15 for the diagnosed and May 01 for the deceased patients. There is a delay of about 2 weeks between the two peaks due to the length of medical treatment for COVID-19. There also is a seasonal (weekly) effect in reporting that creates a jagged pattern in Figure 2.

The curves of daily counts display a non-stationary behaviour with a phase of growth followed by a decline. The aggregated data do not allow us to disentangle the possibly multiple sources of this non-stationarity, such as the non-stationary propagation of the epidemic with an exponential increase in the early phase, the effect of disease detection that depends on the reliability and availability of tests, and the efficiency of treatment for COVID-19. The latter one can improve over time due to advancement of knowledge about the disease, or worsen due to shortages of medical staff and equipment.

Figure 3 below, which shows the daily counts of patients under medical care, can be compared with Figure 1 (ICU, Hospitalization) in PHO (2020 a, page 3) based on a shorter sampling period, ending on April 22 [or with Figure 2 in PHO (2020 c)]. We observe that the peak of the curve depends on the sampling period. The timing of that peak seems to be, approximately, April 1 in Figure 1, PHO (2020 a) and April 11 in Figure 3 given above. This is likely due to a kind of truncation bias. This bias is discussed later

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

Covid Economics 35, 7 July 2020: 176-119

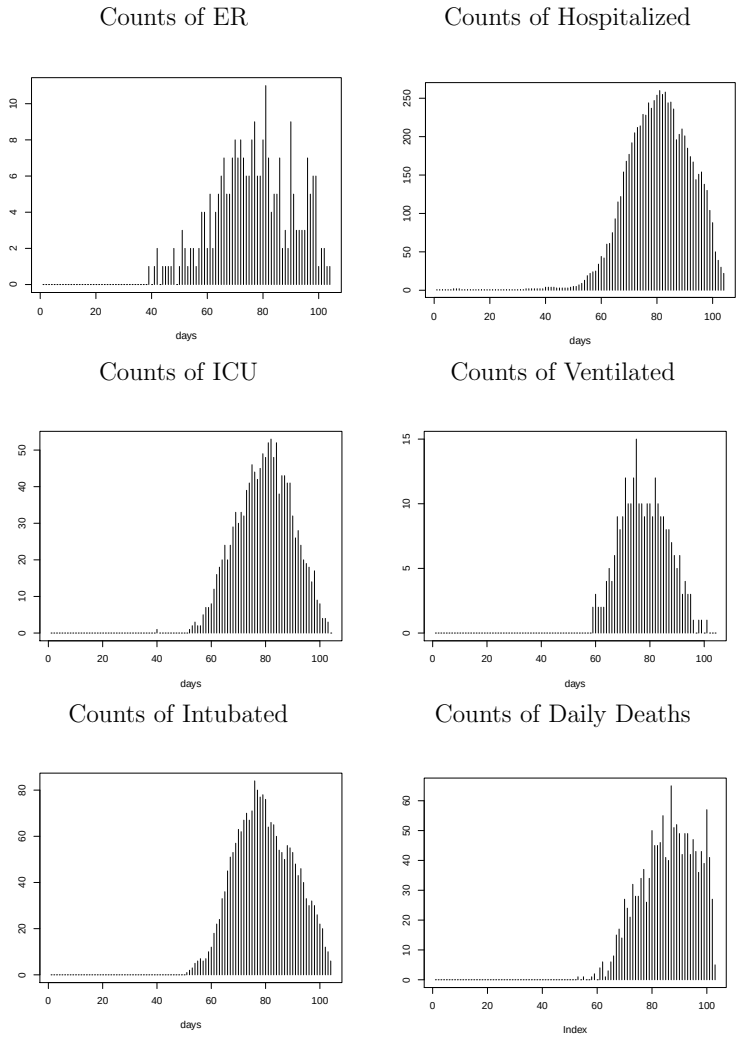

Fig 3. Counts of Medical Care

in Section 2.3 and adjusted for by considering the state transitions instead of "crude" counts.

Figure 4 presents changes in daily counts displayed in Figure 3. These patterns can be compared with Figure 3, Aguerragibiria et.al. (2020), which illustrates observations

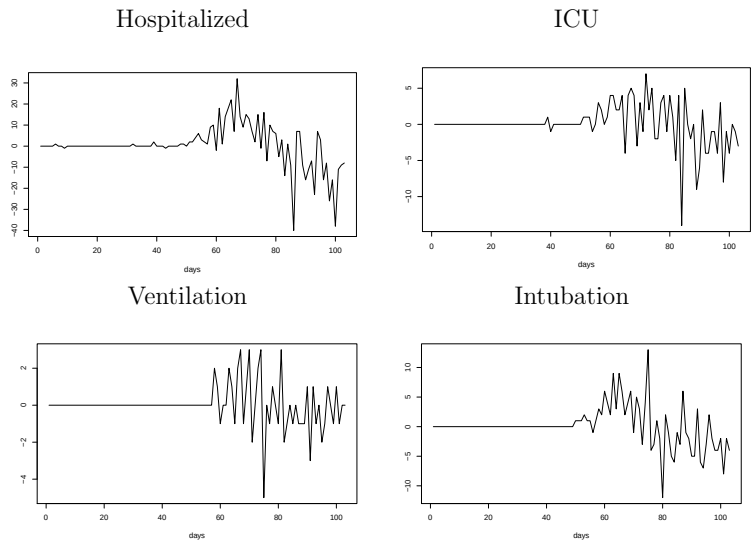

Fig 4. Changes in Daily Counts of Medical Care

Covid Economics 35, 7 July 2020: 176-119

recorded over a period ending on April 22. The difference between the patterns displayed in the two figures seemingly confirms a misleading effect of truncation. Nevertheless, both figures reveal daily fluctuations due to the fact that the counts of individuals under medical care are determined by both the entry and exit effects.

Our approach to examining the individual records relies on a transition model, where each individual is represented by a "history" variable, i.e. a sequence of states [see, e.g. Gourieroux, Jasiak (2007), Chapter 8]. It allows us to adjust for right censoring and attenuate the truncation bias. The state dynamics are defined through transition matrices, which are functions of the current and past environment.

An advantage of the transition model is that it allows for including time and individual dependent explanatory variables, such as the time already spent in the state, i.e. the duration dependence.

The analysis of the transition model with duration dependence leads to the following observations:

The longer an individual stays hospitalized, the lower the probability of moving to the ICU, or of being intubated. The probabilities of recovering and of dying of COVID-19

increase with the length of medical treatment. The probability of death seems to increase with time spent in all types of medical care, at rates depending on the intensity of care. The results concerning the probability to recover are affected by the lack of the date of recovery and need to be interpreted with caution.

The paper is organized as follows. Section 2 describes the states of medical care and illustrates the individual histories of transitions between states. The truncation problem is discussed too. Section 3 presents the empirical data analysis and summary statistics computed under a simplifying assumption of homogeneous Markov chain. Then, the assumption of stationarity of the medical care process is discussed. Section 4 presents the transition model with duration dependence, where the probabilities of transition from each state depend on the time already spent in that state. Section 5 concludes. The summary statistics of sojourn times in each state are given in Appendix A.1. The supplementary figures are given in Appendix A.2.

Covid Economics 35, 7 July 2020: 176-119

2 The Discrete Time Transition Model

This Section introduces the states and the transitions and describes the discrete time dynamic transition model based on the individual histories.

2.1 Individual Histories

Below, we provide a few examples of individual histories found in the PHO data set:

1. Case reported on April 17, recovered by May 05.
2. Case reported on March 29, in the ER on March 26, returned and hospitalized on April 03, in the ICU and under ventilation from April 6 until April 23, not recovered by May 05.
3. Case reported on April 16, hospitalized on April 24, not recovered until May 05.
4. Case reported on April 02, hospitalized on March 31, in the ICU April on 02, died on April 03.
5. Case reported on April 24, not recovered.

These examples of individual histories of patients underlie our approach and suggest the selection of states given below.

Covid Economics 35, 7 July 2020: 176-119

2.2 States

The dataset allow us to distinguish 9 latent states. The states are redefined for compatibility with the transition model and the availability of data.

State 0. Undiagnosed: from the beginning of the sampling period until the case reported date, or a transition date to the next state. On "the case reported date" the individual is diagnosed as COVID-19 infected and enters a follow-up stage that includes self-isolation and/or states of medical treatment.

State 1. D: Domiciled: from the case reported date until the medical care, death or recovery. As the population of Ontario was supposed to self-isolate during the sampling period, we include in state D individuals who are not currently hospitalized and not yet recovered: see below the definition of recovered. Therefore, we have D = isolation + self-isolation. The isolation can be prior to, or after, a medical treatment, which are denoted by D^1 and D^2, respectively. The records contain a limited amount of information on individual isolation. There are very few cases of isolation reported in the dataset and their locations are unknown.

State 2. ER: Emergency Room

State 3. H: Hospitalized due to COVID-19 (hospitalization for other reasons is disregarded)

State 4. ICU: Intensive Care Unit

State 5. V: Ventilation

State 6. T: Intubation

To help adjust breathing, a machine is used to move air in and out of the patient lungs. Ventilators (also called respirators) are machines of different types, including computerized microprocessor controlled machines, as well as CPAP (Continuous Positive Airway Pressure) and non-invasive ventilators. This state of medical care is called Ventilation. The state Intubation refers to placing a tube in the patient's throat to help move air in and out of the lungs while protecting the airway, which is a long term ventilator dependence with a tracheotomy cannula.

In many cases, the time spent in Hospitalization, Intensive Care, Intubation, or under a Ventilator overlap. Given that the states should be disjointed events, we proceed as follows: We treat an individual as being intubated if this stage overlaps with others

(because being intubated means that the individual is already in intensive care, under a ventilator and breathing through a tube). Similarly, if a patient is in intensive care and on a ventilator, we consider him/her to be using a ventilator. When a patient is in intensive care and under hospitalization, we consider him/her as in intensive care.

This distinct states do not obey the relationship $H \supset ICU \supset V \supset T$, and the differences between the states considered, marked with a star, and the traditional counts given in the PHO reports are as follows:

$$\begin{cases} T^* = T \\ V^* = V - T \\ ICU^* = ICU - V \\ H^* = H - ICU \end{cases} = \begin{cases} T = T^* \\ V = V^* + T^* \\ ICU = ICU^* + V^* + T^* \\ H = H^* + ICU^* + V^* + T^* \end{cases}$$

For ease of exposition, the non-starred symbols are used, assuming the reader is aware of the above distinction.

Accordingly, among the 18722 individuals diagnosed in our data set, 2047 individuals went through medical care for COVID-19 during the sampling period of 104 days. The total counts of individuals who at some point over the sampling period were in the distinct (starred) states of medical care given above, are: 138 in ER, 1376 in H (Hospitalization), 243 in ICU, 46 in V (Ventilation) and 244 in T (Intubation).

State 7. R: Recovered.

This state is an important component of a SIRD model. However, for COVID-19, the notion of recovery is not clearly defined and the date of recovery is unknown. The database combines two sources mentioned earlied, which are the iPHIS and CORES that are seemingly not reporting the recovery in the same way. In the most recent PHO practice, the recovery date is determined conventionally as follows. According to the PHO documents "the following cases are considered resolved" (but referred to as "recovered" in the daily Ontario provincial report for the media):

-cases that are reported as "recovered" in iPHIS based on the local public health unit assessment;

-cases that are not hospitalized and are 14 days past their symptoms onset date, or specimen collection date;

-cases that are currently hospitalized, have a case status of "closed" indicating that

Covid Economics 35, 7 July 2020: 176-119

public health follow up is complete and are 14 days past their symptoms;

There is a lack of coherency in this definition, as several individuals reported "recovered" by the iPHIS may still have COVID-19 symptoms, and the definition depends on the local health unit.

In our analysis, we have adopted a similar approach to setting conventionally the date of state R. For that reason, some results concerning the time to recover have to be considered with caution.

State 8. DE: Deceased

Individuals who are not recovered or deceased can remain in their last state on May 05, causing a right censoring problem.

Covid Economics 35, 7 July 2020: 176-119

2.3 Truncation

Let us now consider the state of all individuals on May 05. On that day, the dataset reports 13218 individuals recovered and 1429 deceased. Therefore, on May 05, there are 4075 right censored histories. Among those, 4046 individuals are in state D, 1 in state ER, 22 are hospitalized and 6 are intubated. Due to the lack of a clear definition of recovery, there are 4046 individuals in state D who have been diagnosed since less than 14 days.

Let us now discuss the right truncation of individual histories for individuals who are under medical care on May 05. On May 05, 1043 Ontarians are hospitalized [see, Global News May 05, 2020]. Among them 223 are in the ICU and 166 are under ventilation (according to the PHO definition). At the same time, 3504 critical care beds are available, including 2811 beds equipped with ventilators [see, Office of the Premier, News, April 16, 2020, 1:00 pm]. Therefore, the database suffers from a right truncation problem as very few individuals are reported under medical care on May 05.

As individuals undergoing medical treatment for COVID-19 may even stay for 3 weeks under medical care, the counts given in Figure 3 are reliable up to day 80 - 85 (April 11 - April 16). In particular, the decrease in counts observed in Figure 3 is misleading and essentially due to the truncation in the data base.

In fact, the data available on May 05 show that the curves of counts need to be projected upward to obtain correct predictions for May 05 and the following days [see, Section 4.3]. Similar biases may also arise in the analysis of crude sojourn times in each

state [see, e.g. Lapidus et al. (2020) for application to COVID-19 and Blanhaps et al. (2013), Schweer, Wickelhaus (2015) for theoretical results].

2.4 Specification of transition probabilities

The transitions of interest are conditional on being diagnosed. The following conditional state transitions are assumed to occur :
From D → D, or D → ER, or D → H or D→ ICU, or D → V, or D → T, or D → R, or D → DE .
From ER → D, or ER → ER, or ER → H, or ER → R, or ER → DE.
From H → D, or H → ER, or H → H, or H → ICU, or H → V, or H → T, or H → R or H → DE.
From ICU → D, or ICU → H, or ICU → ICU, or ICU → V or ICU → T, or ICU → R, or ICU → DE .
From V → D, or V → H, or V → ICU or or V → V, or V → T, or V → DE.
From T → D, or T → H, or T → ICU, or T → V, or T → T, or T → R, or T → DE.
From R → R : absorbing state.
From DE → DE : absorbing state.

During a medical care episode some transitions are known to have probability 0. For example, an individual cannot go back to state ICU after recovery, etc.

2.5 Joint Distribution of Individual Histories

The probabilistic model concerns the histories of diagnosed individuals over the set of states given above. The individuals are indexed by i, i=1,..,n, where n is the total size of the sample. The model is applied to the individual histories starting from a common date t_0 equal to the first detection date of January 23, 2020.

As the individuals are diagnosed at different dates, we introduce the state 0 of Undiagnosed in order to align the individual histories on a common time scale. The histories can be characterized either:

i) by the dates of jumps from one state to another, or equivalently by the ordered sequence of occupied states with given sojourn times in each state (this is how the examples of individual histories in Section 2.1 are reported),

Covid Economics 35, 7 July 2020: 176-119

ii) or, by a qualitative time series with given occupied state indicators on each day. For example, the individual history:
00011111111111111111112333333444441111111111111111111111117
represents an individual, undiagnosed for 3 days of the sampling period, isolated (state D) for 18 days, who enters the ER and then remains in hospital for 6 days. Next, the patient is moved to the ICU for 5 days, returns to state D and recovers.

These two representations are equivalent and their use depends on the purpose of a study. The database relies on representation i) which takes less space, while a dynamic model can be applied to either representation. Our estimation method is based on the qualitative time series representation b) in order to accommodate easily the right censoring.

Let us denote by $\{Y_{i,t}, t = t_0, t_0 + 1, ..\}, i = 1, ..., n$ the individual histories since the common date t_0. At the beginning of the history, the individual is still undetected and undiagnosed with $Y_{i,t} = 0$, by convention. After diagnosis date $S_{0,i}$ that depends on each individual, $Y_{i,t}$ is a qualitative variable with 8 possible states described above. The main probabilistic assumption is the following:

Assumption A.1

The individual histories $\{S_{0,i}, Y_{i,t},\ t \geq S_{0,i}\}, i = 1, ..., n$ are independent.

As we focus our analysis on the medical care history conditional on the diagnosis, only the conditional distribution of $\{Y_{i,t}, t \geq S_{0,i}\}$ given $S_{0,i}, Y_{i,S_{0,i}}$ is relevant. This conditional distribution is characterized by a sequence of transition matrices $P_{i,t}$ that provides the conditional distributions of $Y_{i,t}$ given $S_{0,i}, \underline{Y_{i,t-1}}$, where $\underline{Y_{i,t-1}}$ contains all past observations up to and including time $t - 1$.

Assumption A.1 allows for a variety of specifications for the sequence of transition matrices, some of which are examined in the empirical analysis in Sections 3 and 4.

i) **Homogeneous Markov Chain**

Assumption A.2 a)

$P_{i,t} = P$, independent of i and t, $i = 1, ..., n$, $t > S_{0,i}$.

This assumption can be equivalently written under representation i) in terms of dates of jumps and states following those jumps. This representation is denoted by:

Covid Economics 35, 7 July 2020: 176-119

$\{S_{l,i}, \tilde{Y}_{l,i} = Y_{S_{l,i},i},\ l = 0, 1, ...\}$.

The successive sojourn times are denoted by $D_{l,i} = S_{l+1,i} - S_{l,i}, ...$

Then, we obtain the following result [see, e.g. Cinlar (1969), Cox (1970)]:

Proposition 1: The homogeneous Markov Chain assumption (i.e. Assumption A.2 a)) is satisfied if and only if:

i) The series of distinct successive states is Markov with transition matrix $\tilde{P}$ that contains the probabilities of each jump conditional on the previous state that is $\tilde{p}_{j,k} = p_{j,k}/(1 - p_{jj}),\ \forall j, k,\ j \neq k$ and $\tilde{p}_{jj} = 0,\ \forall j$.

ii) Conditional on the series of $\tilde{Y}_{l,i}, l = 0, 1, ...$, the sojourn times $D_{l,i}, l = 1, ..., n$ are independent, such that the distribution of $D_{l,i}$ is a geometric distribution with parameter p_{jj}, where $j = \tilde{Y}_{l,i}$.

Such an independence property of the sojourn times conditional on the occupied state characterizes a renewal process, and facilitates simulations as well as the derivation of asymptotic properties of the estimators. The homogeneous Markov chain is the benchmark model that can be extended in various aspects.

ii) **Time Dependent Chain**

Assumption A.2 b): $P_{i,t} = P_t$, independent of i and dependent on t.

This model allows us for taking into account the calendar time effects, such as the shortages of medical staff and hospital beds, the advances of knowledge about the disease, or the fact that individuals get diagnosed earlier that helps them recover.

ii) **Chain with Duration Dependence**

Assumption A.2 c): $P_{i,t} = P_{t-S_{l_t,i}}$, where $S_{l_t,i}$ is the date of the most recent jump.

This transition depends on the time already spent in the current state.

Under Assumption A.2 c) the process maintains some properties of a renewal process, although the sojourn time distributions are no longer geometric distributions and $\tilde{Y}_{l,i}, l = 1, ..., n$ is no longer a Markov chain.

iv) **Chain with Cohort Effect**

Assumption A.2 d): $P_{i,t} = P_{S_{0,i}}$

The transition depends essentially on the diagnosis time. If we consider a given cohort (or generation) of individuals $\mathcal{P}_{s_0} = \{i : S_{0,i} = s_0\}$ composed of individuals diagnosed on the same date s_0, we have:

Covid Economics 35, 7 July 2020: 176-119

$P_{i,t} = P_{s_0}$, for $i \in \mathcal{P}_{s_0}$

Therefore, for each cohort, we have a homogeneous Markov chain, but the chain is no longer homogeneous if all cohorts are confounded due to some aggregation bias.

v) **Further extensions**

It is possible to consider more complex specifications by introducing jointly the time dependence, duration dependence, a cohort effect and/or also the effects of individual characteristics such as the gender, age, co-morbidity, etc.

Covid Economics 35, 7 July 2020: 176-119

3 Descriptive Empirical Analysis of Transitions

This section presents time independent summary statistics that can be easily computed by averaging individual histories over time and individuals. The statistics discussed in this section are transition frequencies and densities of sojourn times in a given state. These statistics are interpreted under the homogeneous Markov chain assumption. The time dependence is discussed at the end of this section.

3.1 Transition Matrix

Table 1 below presents the transition matrix P estimated from the entire sample.

Table 1. Estimated Frequency Matrix (%)

	D	ER	H	ICU	V	T	R	DE
D	96.04	0.03287	0.1237	0.02031	0.00481	0.02191	3.529	0.2271
ER	68.25	27.51	3.175	0	0	0	0.5291	0.5291
H	14.49	0.01337	81.62	0.4678	0.08019	0.4143	0.0401	2.874
ICU	9.242	0	5.0	81.21	0.303	1.591	0.07576	2.576
V	11.48	0	1.481	1.111	82.96	0.3704	0	2.593
T	4.813	0	0.4627	2.175	0.3702	88.99	0.09255	3.1
R	0	0	0	0	0	0	1	0
DE	0	0	0	0	0	0	0	1

The matrix can be interpreted as follows. Regardless of individual characteristics, the probability that an individual who is currently hospitalized is admitted to the ICU is 0.46% and that he/she dies is 2.87%, regardless of how long the patient has stayed in hospital.

The first six lines of this transition matrix represent a "production process" of the medical care system, with inputs of infected individuals, and outputs of recovered and deceased individuals. Along with the information on the time spent in each state, this can be seen as a basic tool for the budgeting of medical process, planning and scenario analysis [see, e.g. Alvarez et al. (2020), DiDomenico et al. (2020), Aguirregabiria et al. (2020), Atkeson (2020)].

3.2 Distributions of Sojourn Times

Let us now examine the durations of various states of medical treatment conditional on transitioning to another state. The analysis concerns the entire sampling period of 104 days.

We provide the sample distributions of sojourn times in the six states of interest. In order to (partly) eliminate the right censoring/truncation bias, the analysis concerns only complete sojourn times. For example, the distribution of a sojourn time in state ICU conditional on exit to state Intubation is estimated from individuals who have accomplished that transition.

Thus, the adjustment is carried out by examining the sojourn times conditional on the exit state. It simplifies the estimation at the expense of disregarding incomplete sojourn times and the decreasing numbers of individuals left at the end of the sample. This last effect is due to delays in reporting.

Figures 5-7 below show the distributions of sojourn times in states H, ICU and T, conditional on the exit states. Additional summary statistics of these distributions (mean, variance, quantiles) are provided in Appendix A.1. These empirical distributions are evaluated from samples of different sizes, such as N=215 for state H of Hospitalization before state DE of death, or N=6 for state H of hospitalization before state V of Ventilation [see, Appendix A.1.]. In Figures 5-7, we only display the distributions evaluated from a sufficiently large number of observations.

The histograms show decreasing patterns, except for state T of Intubation illustrated in three panels of Figure 7. The exit state is unknown at the entry time into state T of Intubation. However, it can be considered as a measure of severity of the disease with an effect of the length of required Intubation. Under the assumption of perfect foresight

Covid Economics 35, 7 July 2020: 176-119

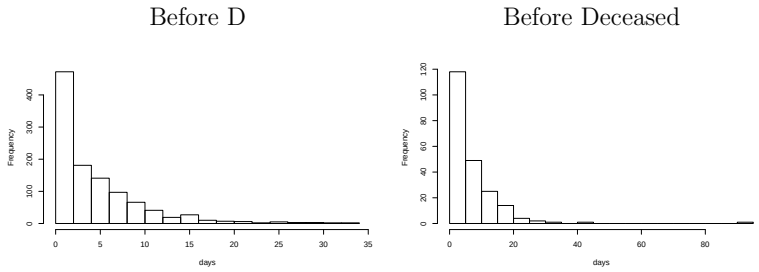

Fig 5. Durations of State Hospitalization Prior to Transition

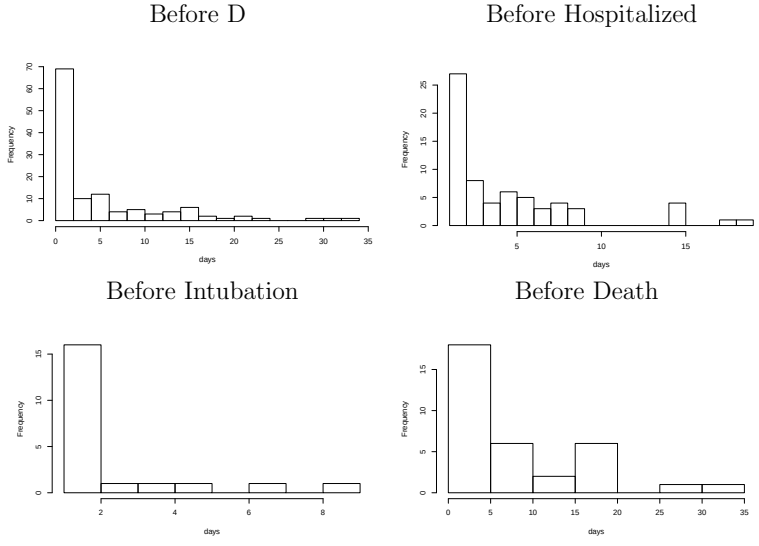

Fig 6. Durations of State ICU Prior to Transition

of the future severity, we distinguish 3 severity levels for the intubated individuals: $v_1 = D, v_2 = ICU, v_3 = DE$.

These three distributions in Figure 7 feature fat right tails with 32 days of maximum Intubation. Depending on the severity level, 25% of individuals spend more than 8.5 days in level v_1, 16.5 days in level v_2 and 14.5 days in level v_3 [see Appendix A.1, 6]. We observe a similar pattern in the mean duration with 6.6 days in level v_1, 11.8 days in level v_2 and 9.1 days in level v_3. The decrease of the number of days between levels v_2 and v_3 reveals that a higher effort in medical care increases the probability of survival.

Covid Economics 35, 7 July 2020: 176-119

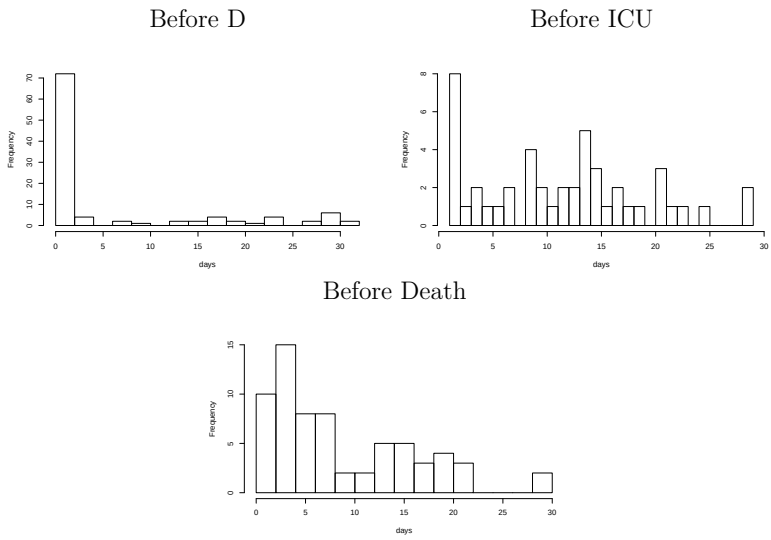

Fig 7. Durations of State Intubation Prior to Transition

These sojourn time distributions are valid state by state and have to be used with caution for deriving conclusions on joint transitions. It is unlikely that an individual transitioning through states of H, ICU and T has an expected time of medical care equal to the sum of average times in H, ICU and T. There may exist a complicated negative or positive dependence between the sojourn times in the successive states. A patient with a severe condition may stay for a short time in ICU and a long time in T. Such complicated dependencies are better captured by means of a duration dependent transitions model [see, Section 4], than by an analysis of the joint distribution of sojourn times. Nevertheless, the correlation between the duration in H before exit to ICU and the duration in ICU before exit to Intubation is -0.128, the correlation between duration in ICU before exit to Intubation and duration in Intubation before Death is -0.233. These negative correlations support the above discussion.

3.3 Homogeneous Markov Chain Model

The summary statistics given in Sections 3.1 and 3.2 can be interpreted as evidence in favour or against the homogeneous Markov chain assumption A.2 a). In particular,

Covid Economics 35, 7 July 2020: 176-119

the empirical transition matrix given in Table 1 is the maximum likelihood estimator of matrix P in a homogeneous Markov Chain. Under A.2 a), the medical care process remains stationary and all non-stationarities in the counts of deceased and recovered are induced by the non-stationarity in the counts of diagnosed and their increase during the early phase of the epidemic. This implies that the medical staff does not change their practices regarding the patient's treatment during the epidemic.

There is evidence in support of the homogeneous Markov chain assumption, such as the decreasing patterns of some duration densities, which resemble geometric densities and the fact that some of the densities do not depend on the state of exit. As well, the mean sojourn time computed directly from the duration data is often close to the mean sojourn time computed from the elements of the transition matrix in Table 1.

Let us, for example, consider the state T of Intubation with p_{jj} =0.89 (see, Table 1). Under Assumption A.2 a), the sojourn time follows a geometric distribution with mean $1/(1\text{-}p_{11}) = 9.0$ days, which can be compared with the expected values reported in Appendix A.1, 6.

On the contrary, the homogeneous Markov chain assumption seems inadequate for some other states. Moreover, the sojourn times of the same individual in various states may be correlated, as mentioned above. Therefore, the homogeneous Markov Chain assumption does not seem to hold, although it conveniently provides a simple framework for computing summary statistics that can serve as benchmarks for comparison [see, Section 4.3.1]. The alternative assumptions to replace A.2 a) are those of time dependence and duration dependence, for example.

Let us first examine if the medical care process is non-stationary. This can be done graphically by comparing the daily estimated transition matrices $\hat{P}_t$. By definition, the estimated transition probabilities $\hat{P}_{jk}$ in Table 1 are weighted averages:

$$\hat{P}_{jk} = \sum_{t=t_0}^{T-1} w_{jt} \hat{P}_{jk,t},$$

with weights that depend on the conditioning state and are proportional to he number of individuals in that state on day t.

Figure 8 shows the trajectories of daily transition probabilities $\hat{P}_{j,k,t}$ for selected states over the sampling period. The series have to be considered with caution as the number

THIS VERSION: July 6, 2020

Covid Economics 35, 7 July 2020: 176-119

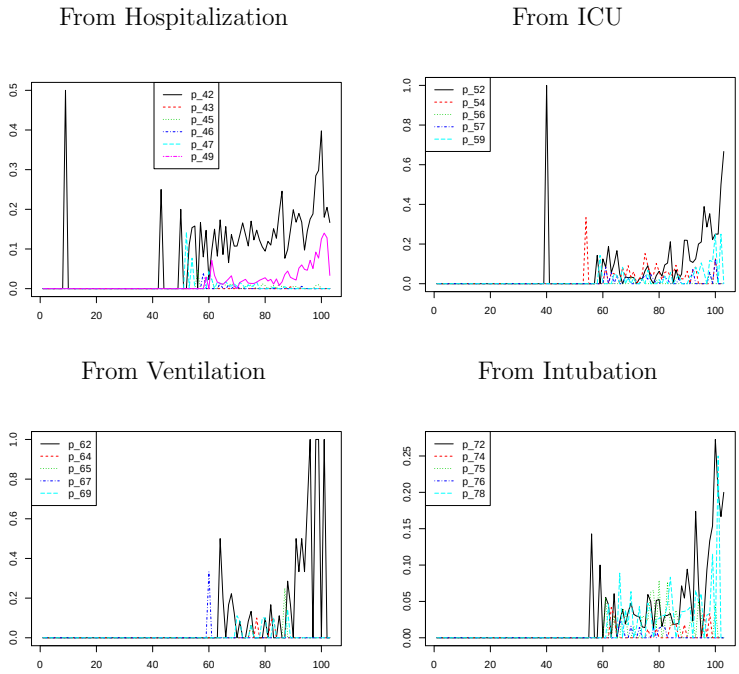

Figure 8 Transitions from States over 104 days

of individuals in a given state at a given date can be small. Therefore, we focus on the transitions from H and T to D (black lines), and from H and T to DE (pink and blue lines, respectively) computed from larger sub-samples. These series display upward trends, while the transitions from H to H and T to T (not reported) have downward trends by the unit mass restriction. Thus, the time spent in these states seems to decline. That can be due to a change in the treatment process, or to a left censoring effect. Indeed, the time of the diagnosis is not equal to the time of infection. If more tests were performed over time, infectious individuals could be detected earlier and then it would be easier to cure them of COVID-19. The distinction between these alternative explanations is out of the scope of this paper.

Another remark concerns the evolution of the number of deaths. It is increasing smoothly from state H, which suggests it is being controlled. On the contrary, the evolu-

tion of deaths from state T is more erratic, which suggests there might be a pressure on the number of beds available for intubation.

4 Model with Duration Dependence

This section extends the benchmark model by introducing duration dependence. First, we describe the specification of the transition matrix. Next, we provide the estimation results. These results are used for adjusting the count summary statistics for right truncation bias and for prediction making.

Covid Economics 35, 7 July 2020: 176-119

4.1 The Model

The (8×8) transition probability matrix P_t has components $p_{j,k,i,t}$. The rows sum up to 1. Each row represents the probabilities of exit from a state j to another state k . The transition probabilities are specified as (conditional) multinomial logit models [see, McFadden (1984)]. The transition probabilities can depend on the time already spent in a state, denoted by $Dur_{it}(D)$ for state D. They depend on the state and vary with both individuals and time.

For example, let us illustrate the transitions $p_{j,k,i,t}$ from state $j = 1$ of D for individual i at time t:

$p_{1,1,i.t} = (\text{D} \rightarrow \text{D}) \propto 1$
$p_{1,2,i.t} = (\text{D} \rightarrow \text{ER}) \propto exp(\beta_{1,1} + \beta_{1,2} Dur_{it}(D))$
$p_{1,3,i,t} = (\text{D} \rightarrow \text{H}) \propto exp(\beta_{1,3} + \beta_{1,4} Dur_{it}(D))$
$p_{1,4,i,t} = (\text{D} \rightarrow \text{ICU}) \propto exp(\beta_{1,5} + \beta_{1,6} Dur_{it}(D))$
$p_{1,5,i,t} = (\text{D} \rightarrow \text{V}) \propto exp(\beta_{1,7} + \beta_{1,8} Dur_{it}(D))$
$p_{1,6,i,t} = (\text{D} \rightarrow \text{T}) \propto exp(\beta_{1,9} + \beta_{1,10} Dur_{it}(D))$
$p_{1,7,i,t} = (\text{D} \rightarrow \text{R}) \propto exp(\beta_{1,11} + \beta_{1,12} Dur_{it}(D))$
$p_{1,8,i,t} = (\text{D} \rightarrow \text{DE}) \propto exp(\beta_{1,13})$

where $Dur_{i,t}(D)$ denotes the duration (time) already spent in the current state D and the symbol $\propto$ denotes "proportional to", with the proportionality coefficient such that the above transition probabilities sum up to 1.

We mentioned earlier that state D comprises individuals of two types: D^1 prior to a

potential medical care and D^2 returned from medical care. Therefore the transition to a medical care concerns individuals in D^1 and depends on time spent in state $D = D^1$.

4.2 Estimation

The model is estimated by the Maximum Likelihood. The results of the estimation of the model are presented in Table 2, which contains the estimated coefficients for each row of the transition matrix. In each panel, the first coefficient is the intercept and the second one is the duration sensitivity.

The state "R" of Recovery is only observed at the end of the sample, and many individuals remain reported in state D in the dataset while they could have recovered. Therefore the transition from state "D" to state "R" of Recovery, computed from the estimated parameters is necessarily underestimated. Nevertheless, the dataset includes reliable information on the transition to/from the medical states and the state "Death". Thus, the estimates can be used to estimate accurately the dynamic of the probabilities and counts in these states.

Table 2: Estimated Coefficients

	TRANSITION FROM D												
States	ER		H		ICU		V		T		R		DE
Parameters	$\beta_{1,1}$	$\beta_{1,2}$	$\beta_{1,3}$	$\beta_{1,4}$	$\beta_{1,5}$	$\beta_{1,6}$	$\beta_{1,7}$	$\beta_{1,8}$	$\beta_{1,9}$	$\beta_{1,10}$	$\beta_{1,11}$	$\beta_{1,12}$	$\beta_{1,13}$
Estimates	-8.0924	0.0040	-5.0664	-0.0665	-6.9397	-0.0635	-9.1849	-0.0262	-6.4233	-0.0861	-2.7289	-0.0217	-6.0468
	TRANSITION FROM ER												
States	D	H	R	DE									
Parameters	$\beta_{2,1}$	$\beta_{2,2}$	$\beta_{2,3}$	$\beta_{2,4}$									
Estimates	0.9086	-2.1595	-3.9512	-3.9512									
	TRANSITION FROM H												
States	D		ER		ICU		V		T		R		DE
Parameters	$\beta_{3,1}$	$\beta_{3,2}$	$\beta_{3,3}$	$\beta_{3,4}$	$\beta_{3,5}$	$\beta_{3,6}$	$\beta_{3,7}$	$\beta_{3,8}$	$\beta_{3,9}$	$\beta_{3,10}$	$\beta_{3,11}$	$\beta_{3,12}$	$\beta_{3,13}$
Estimates	0.2103	-0.2535	-1.6251	-3.7254	-3.3540	-0.2278	-3.4636	-0.7252	-3.0971	-0.3065	-8.1585	0.0299	-3.3466
	TRANSITION FROM ICU												
States	D		H		V		T		R		DE		
Parameters	$\beta_{4,1}$	$\beta_{4,2}$	$\beta_{4,3}$	$\beta_{4,4}$	$\beta_{4,5}$	$\beta_{4,6}$	$\beta_{4,7}$	$\beta_{4,8}$	$\beta_{4,9}$	$\beta_{4,10}$	$\beta_{4,11}$		
Estimates	-0.2942	-0.2116	-0.7782	-0.2372	-4.2901	-0.1215	-1.6329	-0.3035	-7.9989	0.0593	-3.4509		
	TRANSITION FROM V												
States	D		H		ICU		T		R		DE		
Parameters	$\beta_{5,1}$	$\beta_{5,2}$	$\beta_{5,3}$	$\beta_{5,4}$	$\beta_{5,5}$	$\beta_{5,6}$	$\beta_{5,7}$	$\beta_{5,8}$	$\beta_{5,9}$	$\beta_{5,10}$	$\beta_{5,11}$		
Estimates	0.7679	-0.2843	-2.2051	-0.1448	-1.8467	-0.2322	7.5903	-9.7537	-53.9343	-66.3286	-3.4657		
	TRANSITION FROM T												
States	D		H		ICU		V		R		DE		
Parameters	$\beta_{6,1}$	$\beta_{6,2}$	$\beta_{6,3}$	$\beta_{6,4}$	$\beta_{6,5}$	$\beta_{6,6}$	$\beta_{6,7}$	$\beta_{6,8}$	$\beta_{6,9}$	$\beta_{6,10}$	$\beta_{6,11}$		
Estimates	-0.3381	-0.2170	-3.3238	-0.1377	-2.1517	-0.1031	-3.8112	-0.1127	-7.5641	0.0351	-3.3569		

The duration dependence is introduced for each state except the state of Emergency Room ER. Indeed, a majority of individuals stay in ER for one day only. Table 2 does not report

Covid Economics 35, 7 July 2020: 176-119

the results concerning probabilities of remaining in a given state, which are assumed to be proportional to 1, for identification. The values of parameters of a multinomial model do not have a direct interpretation. Instead the ratios of transition probabilities are interpretable (in the 2-state case, the odd-ratios are interpretable).

Let us consider the state of intubation T. After one day spent in this state, the ratio of the transition to D (i.e. the way to recovery) and the transition to death DE are:

$$\exp(\beta_{6,1} + \beta_{6,2} - \beta_{6,11}).$$

After 10 days, this ratio is:

$$\exp(\beta_{6,1} + 10\beta_{6,2} - \beta_{6,11}).$$

Therefore, the ratio of these two ratios 10 days/1 day is:

$$\exp(9\beta_{6,2}) = \exp(-0.22 \times 9) < 1.$$

Thus, when the time spent in Intubation T increases, the chance to exit on the way to recovery diminishes, as compared to the earlier days.

We illustrate these effects conditional on the state sojourn time (duration), which we set to vary from 1 to 15 days in Figures 15-19 in Appendix A.2.

Figure 15 depicts the dynamics of the transition probabilities from state "D" in terms of its duration. This figure contains eight panels. Each panel shows the probability of transition to states: D, ER, H, ICU, V, T, R, and DE. The probability that someone in state D stays in state D remains above 0.93 and increases with the duration of state D. Indeed, isolated individuals mostly have mild symptoms and are more likely to stay isolated until their recovery. The probability of recovering varies between 0.060 after 1 day and 0.045 after 15 days. However, we expect this probability to be higher due to the lack of observability of the date of recovery. The transition probabilities to the other states are negligible.

Figure 16 displays the transitions from Hospitalization. Individuals after one day of Hospitalization likely have mild symptoms and have a probability 0.46 of moving to state D and becoming isolated. This probability decreases as the duration increases while

Covid Economics 35, 7 July 2020: 176-119

the risk of remaining hospitalized increases. Similarly, Figures 17-19 show the transition probability from the states "ICU", V of "Ventilation" and T of "Intubation".

The probability of remaining in these states increases quickly with time spent in each state. The probability of death seems to increase with time spent in all types of medical care, at rates depending on the intensity of care. The rate of increase is the highest for V, followed by that for ICU, which itself is followed by a slightly lower rate for T.

In general, the probabilities of transition to death after long durations are close to the transition probabilities in Table 1, computed under the homogeneous Markov chain assumption.

The results concerning the probability to recover are affected by the lack of the date of recovery and need to be interpreted with caution.

4.3 Fitted Values and Predictions

Let us now explain how the estimated transition model can be used to compute the fitted values of counts over the observation period and the forecasts. First, we consider the computation under the homogeneous Markov chain assumption, and next under the assumption of duration dependence. We focus on the counts of patients who are hospitalized (state H), in the ICU and state T of Intubation.

4.3.1 **Homogeneous Markov Chain**

Recall that state D includes two types of histories: D^1 are individuals who are isolated at the detection date and D^2 are individuals who are on their way to recovery after a transition through H, ICU, V and T. For subsequent computation, we need to separate these two types of observations and focus on D^1. The first row of the transition matrix in Table 1 is replaced by:

Table 3. Modified Row 1 of Frequency Matrix (%)

	D^1	ER	H	ICU	V	T	R	DE
D^1	93	0.0	1	0.2	1	0.2	3.6	1

The modified transition matrix is denoted by $\tilde{P}$. After being diagnosed, an individual is assigned to a state. The time-independent (constant) probabilities of assignments estimated from the entire sample are:

Table 4. Probabilities of Assigment (%)

	D	ER	H	ICU	V	T	R	DE
v	93	1.0	5.0	1.0	0	0	0	0

We have a causal scheme with 3 episodes:

1. the detection episode that ends at the case reported date.
2. the assignment episode, with the selection based on v given in Table 4,
3. the medical care process summarized by matrix $\tilde{P}$.

At time T=May 05, we observe for each individual either a recovery, or a death or a censored duration in the last state, such as one of the medical care states.

The fitted values, i.e. the expected counts, are computed from the sequence of inputs X_t, which are the counts of new diagnosed on day t (see Figure 2). Let N_t denote the row vector of length 8, with elements equal to the counts of individuals in each state on day t. The fitted counts are given by:

$$EN_t = \sum_{\tau=1}^{t} (X_\tau v' \tilde{P}^{t-\tau}), \quad t = 1, ..., 104. \tag{4.1}$$

Thus, we evaluate the future expected counts for each cohort (generation) τ and sum over the past cohorts. Figure 9 below provides the fitted counts of H, ICU and T. They are close to the patterns shown in Figure 3 up to day 80 (April 11), and are significantly different afterwards. Indeed, the fitted values provide an adjustment for the right truncation bias revealed in Figure 3. It is therefore important to use the model-based figures instead of data-based figures in the presence of truncation and reporting lag problems, to avoid misleading conclusions, especially concerning the peak of the epidemic. Figure 9 shows the peak of state H of Hospitalization that seems about to appear at the beginning of May for state H (from recent data, we know that there was a peak in the first wave, followed by a peak in the second wave). The peak in state T of Intubation is not visible from the fitted counts due to long sojourn times of state T.

4.3.2 **Model with Duration Dependence**

A similar analysis can be performed by using the model with duration dependence discussed in Sections 4.1-4.2. Then, equation (4.1) becomes:

$$EN_t = \sum_{\tau=1}^{t} (X_\tau v' E[\tilde{P}_{\tau+1}, \tilde{P}_{\tau+2} \cdots \tilde{P}_t)]), \quad t = 1, ..., 104. \tag{4.2}$$

Covid Economics 35, 7 July 2020: 176-119

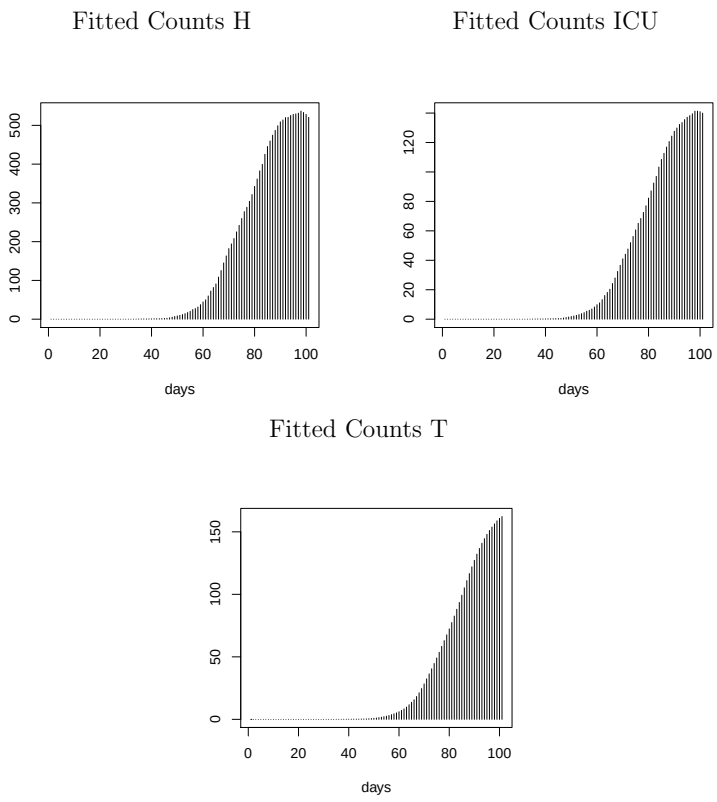

Figure 9. Fitted Counts of Medical Care over 104 days

Covid Economics 35, 7 July 2020: 176-119

It accounts for the fact that the matrices depend on time through the durations spent in each state. As the individual durations Dur_{it} are random, the product of matrices has to be integrated out, which can be easily done by simulations.

Figures 10 and 11 show the estimated and observed counts of patients under medical care and the cumulative counts of deaths, respectively. Figure 10 illustrates the good fit of the transition model to the data and the satisfactory predictive power of the model.

THIS VERSION: July 6, 2020

Covid Economics 35, 7 July 2020: 176-119

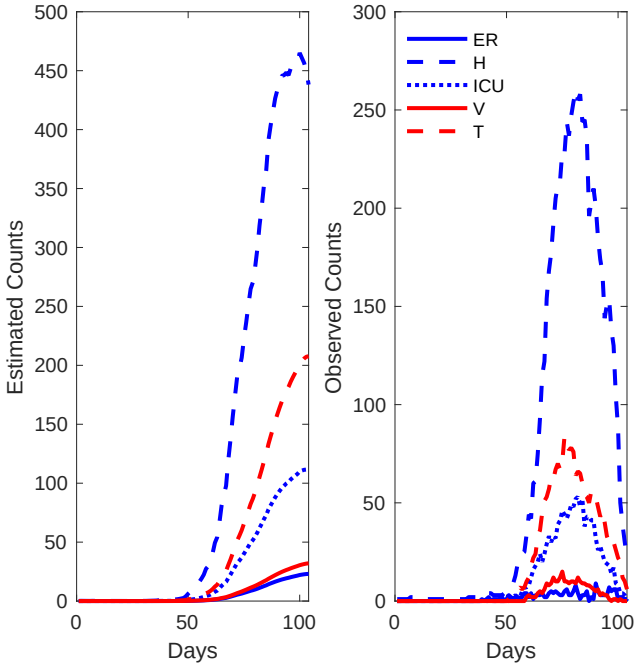

Figure 10: Estimated and Observed Counts of Medical Care

The model predicts, as expected from Figure 3, higher than recorded medical state counts over the last part of the sampling period until May 05. Figure 11 plots the estimated and the observed number of cumulative deaths. The figure shows that the total death is accurately predicted. The figures confirm that the model provides an adjustment for right truncation and reporting lag. As mentioned earlier the data suffer from a truncation bias and a lag in reporting that start around day 80 (April 11) of the observational period and generate decreasing patterns of daily counts, found ex-post incompatible with the publicly available counts provided later by the PHO. The model is shown to "adjust" for these biases.

4.3.1 **Predictions**

A similar approach can be applied to predict the future counts and the numbers of

THIS VERSION: July 6, 2020

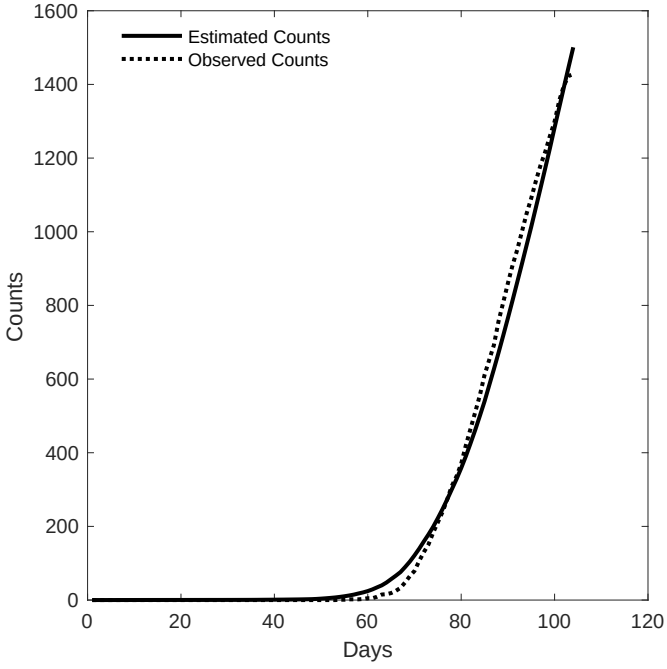

Figure 11: Estimated and Observed Counts of Deaths

Covid Economics 35, 7 July 2020: 176-119

beds required for the different types of medical care [see Grasselli et al. (2020), Murray (2020), for predictions of beds in Italy and the US, respectively].

Due to the time of medical treatment, we recommend predicting a term structure of counts, up to say 20 days rather than computing one day ahead forecasts only. This long-run prediction requires additional forecasting of future inputs $X_{T+1}, X_{T+1}, ..., X_{T+H}$. This is only possible if the model is completed by a dynamic model for the X_t's, from a SI component of the SIRD-type of model, obtained by logit adjustments for the series of cumulated counts of diagnosed individuals [Berkson (1944)]. This is clearly out of the scope of the present paper.

THIS VERSION: July 6, 2020

5 Concluding Remarks

This paper examined the process of medical treatment during the early phase of COVID-19 epidemic in Ontario. We investigated the medical care dynamics and its adjustment over time. We considered a benchmark transition model based on the homogeneous Markov chain assumption and studied its extension that accounts for time dependence and duration dependence. The duration dependence revealed in some of the transitions points out to an important limitation of standard SIR models. The advantage of the transition model is the adjustment for truncation that helps eliminate misleading results based on a naive interpretation of the series of aggregate counts.

We have disregarded the effect of the detection process of infected individuals. Indeed, the detection process may induce both left truncation and left censoring. The analysis of left truncation in the case of COVID-19, is complicated further by the large number of asymptomatic individuals who are difficult to detect [see e.g. Gourieroux, Jasiak (2020) for an approach to recovering the unobserved counts of undetected infectious individuals]. The left censoring arises, as the detection date does not coincide with the infection date, and the difference between these dates may depend on the detection efforts and the numbers of tests performed.

The transition model-based analysis could be further extended to accommodate the individual characteristics and spatial dependence between the infections in Ontario.

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

REFERENCES

Alvarez, F., Argente, D. and F. Lippi (2020): A Simple Planning Problem for COVID-19 Lockdown, Covid Economics, Issue 14, May 06.

Aguirregabiria, V., Gu, J., Luo, Y. and P. Mira (2020): A Dynamic Structural Model of Virus Diffusion and Network Production: A First Report. Working Paper 665, University of Toronto.

Atkeson, A. (2020): What Will Be The Economic Impact of Covid-19 In The US? Rough Estimates of Disease Scenarios, NBER Working Paper 26867.

Blanghaps, N. Nov, Y. and G. Weiss (2013): Sojourn Times Estimation in and M/G/∞ Queue with Partial Information, Journal of Applied Probability, 27, 314-324.

Berkson, J. (1944): Application of the Logistic Function to Bio-Assay, JASA, 339-357.

Cinlar, E. (1969): Markov Renewal Theory, Advances in Applied Probability, 1, 123-187.

Cox, D. (1970): Renewal Theory, Wiley, London.

DiDomenico, L., Pullano, G., Coletti, P., Hens, N. and V. Colizza (2020):Expected Impact of School Closure and Telework to Mitigate COVID-19 Epidemic in France, www.epicx-lab.com/covid-19.html]

Gourieroux, C., and J. Jasiak (2007): Econometrics of Individual Risks: Credit, Insurance and Marketing, Princeton University Press.

Gourieroux, C., and J. Jasiak (2020): Time Varying Markov Processes with Partially Observed Aggregate Data: An Application to Coronavirus, Arxiv 2005.04500.

Grasselli, G., Pesenti A., and M. Cecconi (2020): Critical Care Utilization for the Covid-

Covid Economics 35, 7 July 2020: 176-119

19 Outbreak in Lombardy, Italy: Early Experience and Forecast During an Emergency Response, Journal of the American Medical Association, forthcoming.

Kermack, W., and A., McKendrick (1927): A Contribution to the Mathematical Theory of Epidemics,Proceedings of the Royal Statistical Society, A, 115,700-721.

Lapidus, N., Zhou, X., Carrat, F. Riou, B. , Zhao, Y., and G. Hejblum (2020): Biased and Unbiased Estimation of the Average Length of Stay in Intensive Care Unit in the Covid-19 Pandemic, medRxiv.

McFadden, D. (1984): Econometric Analysis of Qualitative Response Models,in Handbook of Econometrics,Vol 2,1395-1457, Elsevier.

Murray, C. (2020): Forecasting COVID-19 Impact on Hospital Bed-Days, ICU-Days, Ventilator-Days by US State in the Next 4 Months, Medrxiv, March 30.

Office of the Premier, News, April 16, 2020, 1:00 pm: Ontario Significantly Expands Hospital Capacity to Prepare for Any COVID-19 Outbreak Scenario.

Public Health Ontario (PHO) 2020a: Enhanced Epidemiological Summary: COVID-19 and Severe Outcomes in Ontario, April 22 https://www.publichealthontario.ca/-/media/documents/ncov/epi/covid-19-severe-outcomes-ontario-epi-summary.pdf?la=en

Public Health Ontario (PHO) 2020b: Epidemiological Summary: COVID-19 in Ontario: January 15, 2020 to May 27, 2020 https://www.publichealthontario.ca/-/media/documents/ncov/epi/2020/covid-19-daily-epi-summary-report.pdf?la=en

Public Health Ontario (PHO) 2020c: Enhanced Epidemiological Summary: COVID-19 in Ontario: A Focus on Diversity, May 14, 2020 https://www.publichealthontario.ca/-/media/
documents/ncov/epi/2020/covid-19-daily-epi-summary-report.pdf?la=en

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

Schweer, S. and C. Wickelhaus (2015): Nonparametric Estimation of the Service Time Distribution in the Discrete Time GI/G/∞ Queue with Partial Information, Stochastic Processes and Their Applications, 125, 233-253.

Toda, A. (2020): Susceptible-Infected-Recovered (SIR) Dynamics of COVID-19 and Economic Impact, ArXiv:2003.11221v2

Vynnicky,.E., and R. White (eds)(2010): An Introduction to Infectious Disease Modelling, Oxford Univ Press,

Yan, P., and G., Chowell (2019): Quantitative Methods for Investigating Infectious Disease Outbreaks, Springer.

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

Appendix A.1

Sojourn Times - Descriptive Statistics

1. State D

Duration in D before ER.

There are N =123 patients with only 2 durations over the first 30 days.

	N	min	25%	median	75%	max	mean	var
total	123	1	2	3	8	40	5.813	37.940

Duration in D before Hospitalized

	N	min	25%	median	75%	max	mean	var
total	463	1	2	4	8	36	5.460	20.309

Duration in D before ICU

There are only 5 cases over the first 30 days. One individual spends one day to ICU and returns to state D.

	N	min	25%	median	75%	max	mean	var
total	77	1	2	3.5	6	32	4.644	19.432

Duration in D before Ventilation

There are 18 individuals who move from D to Ventilation. These transitions took place after April 02.

	N	min	25%	median	75%	max	mean	var
total	18	1.0	2.0	4.0	6.5	26.0	5.611	40.722

Duration in D before Intubation

	N	min	25%	median	75%	max	mean	var
total	82	1	2	4	6	14	4.365	8.185

Duration in D before Recovered.

These individuals may have undergone medical treatments, returned to D and were reported as Recovered.

THIS VERSION: July 6, 2020

	N	min	25%	median	75%	max	mean	var
total	13207	1	17	24	33	103	25.147	109.524

Duration in D before Deceased

There are no transitions from D to Deceased over the first 62 days. One individual dies after 49 days in isolation without any medical treatment.

	N	min	25%	median	75%	max	mean	var
total	851	1	4	7	11	49	7.994	35.323

The durations prior to transitions from state 3 of De are given in Figure 5.

2. State ER

Duration in ER before return to D

	N	min	25%	median	75%	max	mean	var
total	137	1	1	1	1	25	1.781	12.157

Duration in ER before Hospitalization

The results are based on N=6 patients.

	N	min	25%	median	75%	max	mean	var
total	6	1	1	3	5	7	3.333	7.066

One individual has recovered on the last day of sampling period on day 104 after staying in ER. Another individual died after 2 days in ER.

3. State Hospitalization

Duration in Hospitalization before return to D

	N	min	25%	median	75%	max	mean	var
total	1084	1	1	3	7	34	4.959	27.008

Duration in Hospitalization before ER

One individual made that transition after 1 day in hospital.

Duration in Hospitalization before ICU

These durations are recorded after March 20.

	N	min	25%	median	75%	max	mean	var
total	35	1.0	1.0	2.0	3.5	15.0	3.085	10.668

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

Duration in Hospitalization before Ventilation.

These transitions start on day 59 of the sampling period, about March 22.

	N	min	25%	median	75%	max	mean	var
total	6	1.00	1.25	2.50	3.00	3.00	2.166	0.966

Duration in Hospitalization before Intubation

These transitions start on day 53 , i.e. about March 16

	N	min	25%	median	75%	max	mean	var
total	31	1	1	2	3	12	2.774	6.780

Duration in Hospitalization before Recovered

There are only 3 transitions after hospitalization of 28, 9 and 39 days.

Durations of Hospitalization before Death

The are no transitions to state 9 from Hospital over the first 30 days.

	N	min	25%	median	75%	max	mean	var
total	215	1	3	5	10	92	7.525	71.568

4. State ICU

Duration in ICU before D

There are only 2 transitions before day 60 of March 23.

	N	min	25%	median	75%	max	mean	var
total	122	1.00	1.00	1.00	6.75	33.00	5.327	47.627

Duration in ICU before Hospitalization

There is only one transition before day 60.

	N	min	25%	median	75%	max	mean	var
total	66	1	2	3	6	19	4.863	18.673

Duration in ICU before Ventilation

There are only 4 durations between days 71 (April 3) and 99 (April 30) of 1,3 ,6 and 7 days.

Duration in ICU before Intubation

There are no transitions before March 23 and only 4 durations between April 22 and the end of sample

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

	N	min	25%	median	75%	max	mean	var
total	21	1	1	1	2	9	2.285	4.914

Duration in ICU before Recovered

One individual makes a transition into R on day 104 (May 04) after 20 days in the ICU.

Duration in ICU before Death

There are no transitions before day 60 (March 23).

	N	min	25%	median	75%	max	mean	var
total	34	1	2	5	14	35	4.863	18.673

5. State Ventilation

Duration in Ventilation before return to D

There are no durations before day March 23 (day 60).

	N	min	25%	median	75%	max	mean	var
total	31	1	1	1	2	26	4.774	57.580

Duration in Ventilation before return to Hospitalized

There are only 4 durations between days 76 and 84 of length 2,8,1, and 25 days, respectively.

Duration in Ventilation before return to ICU

There are 3 durations between days 76 and 88 of length 1, 8 and 9.

Duration in Ventilation before Intubation

There is only one 1-day duration on day 61 (March 24).

Duration in Ventilation before Death

	N	min	25%	median	75%	max	mean	var
total	7	5.0	7.5	9.0	11.0	16.0	9.571	12.952

6. State Intubation

Duration in Intubation before return to D

There is only one transition before day 60 (March 23).

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

	N	min	25%	median	75%	max	mean	var
total	104	1.0	1.0	1.0	8.5	32.0	6.586	96.380

Duration in Intubation before return to Hospitalized

There are no transitions before day 60 (March 23).

	N	min	25%	median	75%	max	mean	var
total	10	1.00	7.50	10.50	11.75	17.00	9.9	22.988

Duration in Intubation before return to ICU

There are no transitions before day 60 (March 23).

	N	min	25%	median	75%	max	mean	var
total	47	1.0	5.5	12.0	16.5	29.0	11.80851	59.679

Duration in Intubation before Ventilation

All durations occurred between days 67 and 83.

	N	min	25%	median	75%	max	mean	var
total	8	1.0	1.0	10.0	18.5	29.0	11.25	.

Duration in Intubation before Recovered

There are 2 durations with transition on day 104 of 35 and 7 days.

Duration in Intubation before Death

	N	min	25%	median	75%	max	mean	var
total	67	1.0	3.0	7.0	14.5	30.0	9.149	53.886

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

Appendix A.2

Additional Figures

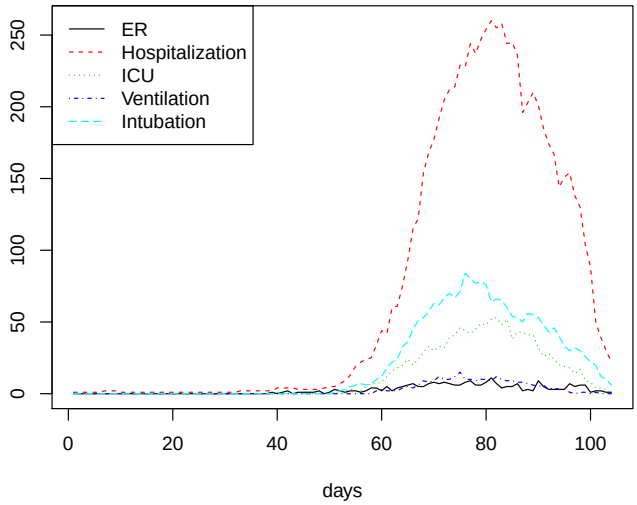

Figure 12: Counts of Medical Care States 3 to 7

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

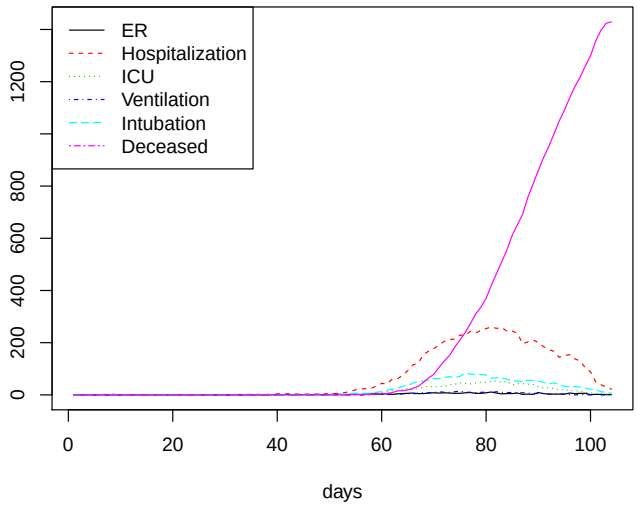

Figure 13: Counts of Medical Care States 3 to 7 and Cumulated Deceased

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

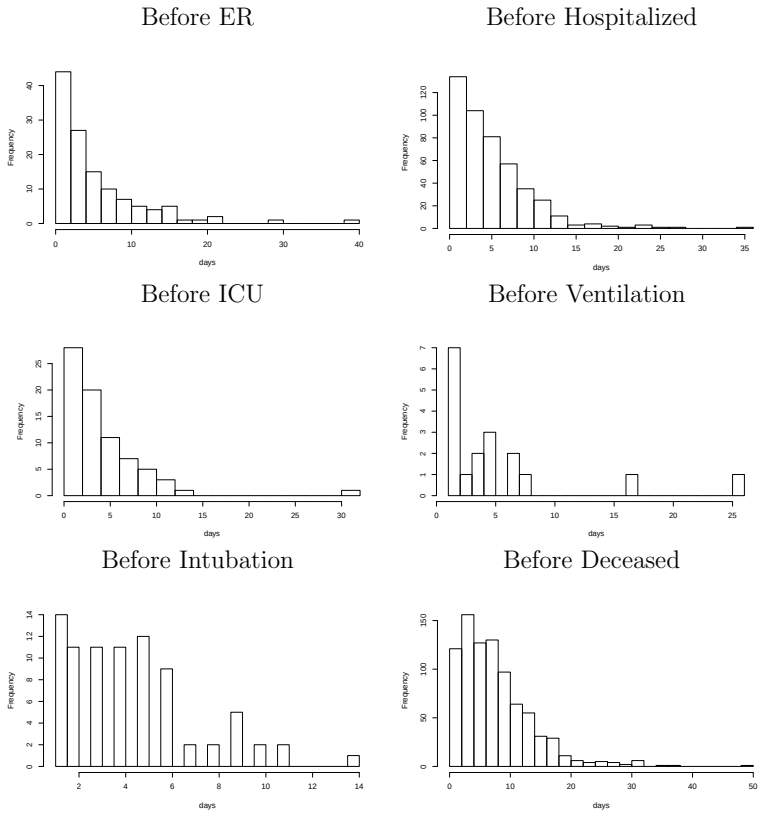

Figure 13: Durations of State D Prior to Transition

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

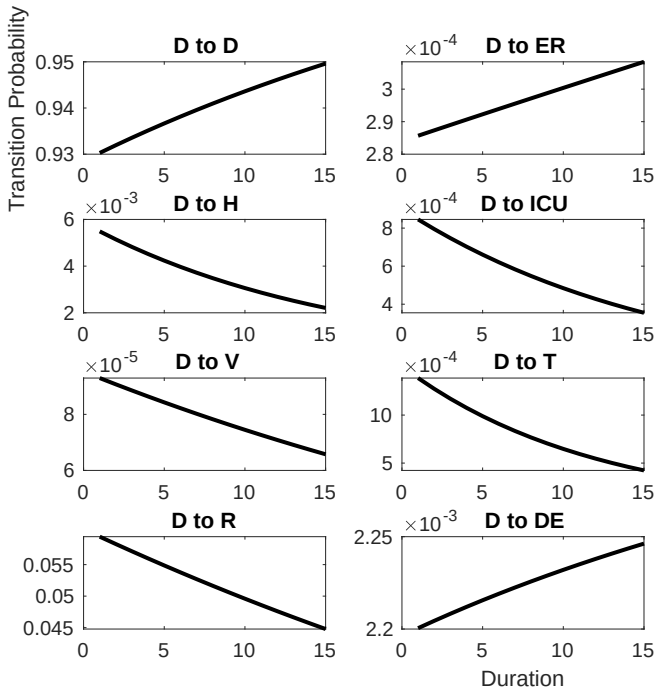

Figure 15: Transition Probabilities from "D" as Functions of Duration

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

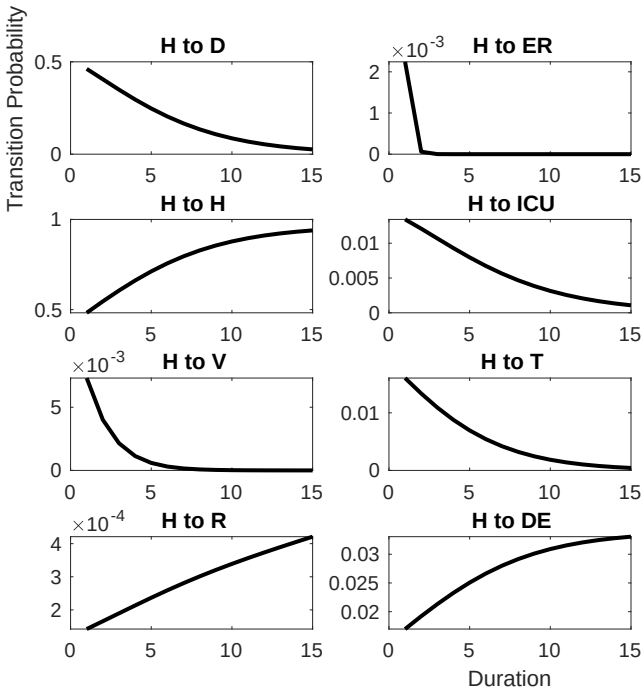

Figure 16: Transition Probabilities from "Hospitalized" as Functions of Duration

Covid Economics 35, 7 July 2020: 176-119

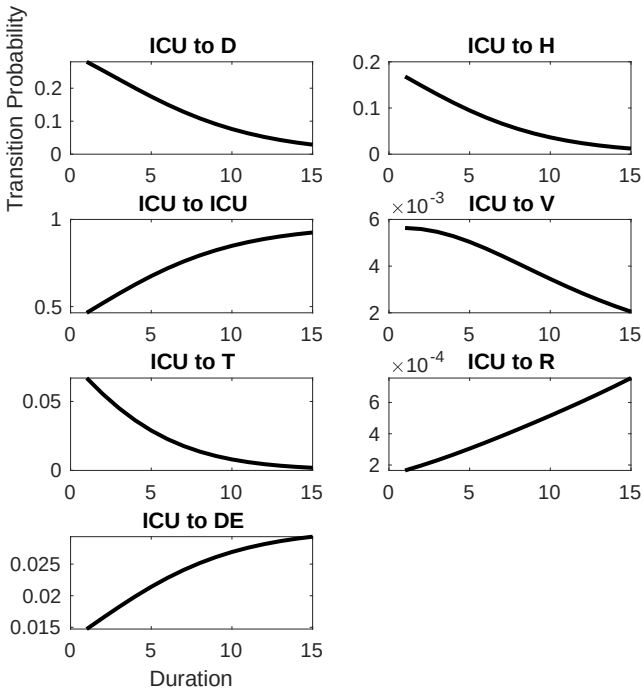

Figure 17: Transition Probabilities from "Intensive Care Unit" as Functions of Duration

Covid Economics 35, 7 July 2020: 176-119

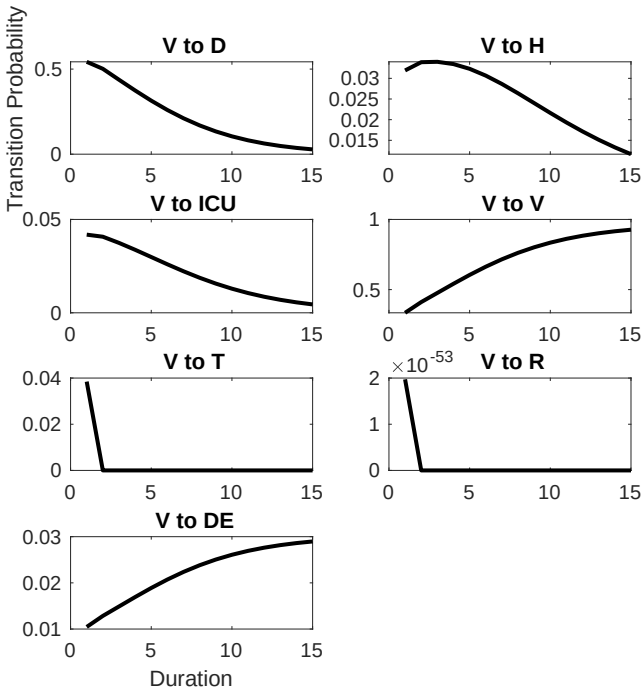

Figure 18: Transition Probabilities from "Ventilation" as Functions of Duration

Covid Economics 35, 7 July 2020: 176-119

THIS VERSION: July 6, 2020

Covid Economics 35, 7 July 2020: 176-119

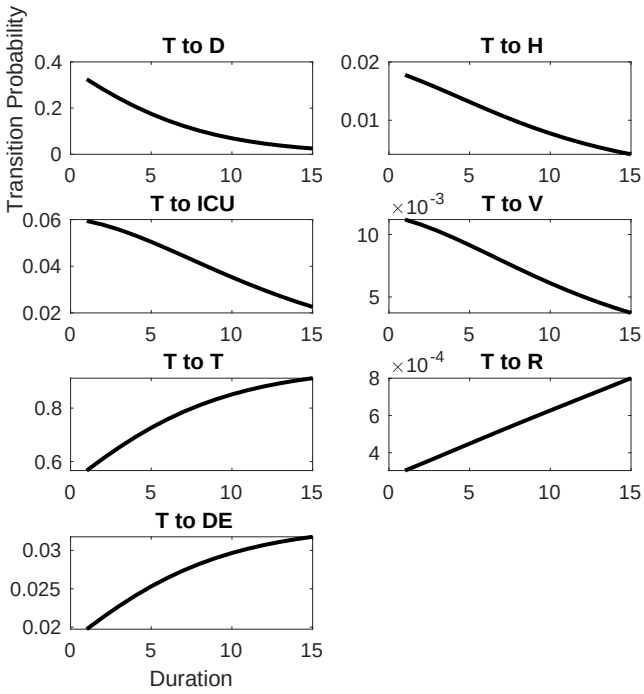

Figure 19: Transition Probabilities from "Intubation" as Functions of Duration

Covid Economics **Issue 35, 7 July 2020**

Task content and job losses in the Great Lockdown[1]

Filippos Petroulakis[2]

Date submitted: 1 July 2020; Date accepted: 4 July 2020

I examine the short-term labor market effects of the Great Lockdown in the United States. I analyze job losses by task content (Acemoglu & Autor 2011), and show that they follow underlying trends; jobs with a high non-routine content are especially well-protected, even if they are not teleworkable. The importance of the task content, particularly for non-routine cognitive analytical tasks, is strong even after controlling for age, gender, race, education, sector and location (and hence for differential demand shocks). Jobs subject to higher structural turnover rates are much more likely to be terminated, suggesting that easier-to-replace employees were at a particular disadvantage, even within sectors; at the same time, there is evidence of labor hoarding for more valuable matches. Individuals in low-skilled jobs fared comparatively better in industries with a high share of high-skilled workers.

Covid Economics 35, 7 July 2020: 220-256

1 The views expressed in this paper are those of the author alone and may not necessarily represent those of the Bank of Greece or the Eurosystem. I would like to thank Theodora Kosma, Katerina Nikalexi, Lara Vivian, and Agostino Consolo for comments and discussions. All errors are mine.

2 Economic Analysis and Research Department, Bank of Greece.

Covid Economics 35, 7 July 2020: 220-256

1 Introduction

The COVID-19 pandemic and the associated non-pharmaceutical interventions (NPIs) to contain it have rendered a massive shock on the global economy. In the United States, an unprecedented 22 million people applied for unemployment benefits over the space of 4 weeks, an almost 30-fold increase over the previous 4-week period.

Substantial policy measures have been enacted in virtually all advanced economies, and the United States presents an interesting case, in that policy measures were designed to support workers, through a large expansion of unemployment benefits, rather than jobs, which was the case in a number of countries in Europe which implemented employment support measures, such as wage subsidies. At this early stage, employment losses at the level of the United States have not been reported elsewhere.

At first glance, the reaction may appear excessive to what is likely to be a temporary shock. On the one hand, firms may fire workers to reduce costs in the face of dwindling revenue, particularly if they have other liabilities close to maturity, or because the uncertainty has increased their risk aversion. On the other hand, search costs for good job-worker matches are high, and labor shedding is inefficient; one would then expect that firms shed their hardest-to-replace workers last.

In this paper, I investigate the drivers of short-term job losses in the United States as a result of COVID-19 and the "Great Lockdown" recession. I focus on the United States both because it represents the more extreme case of labor market shock, partially because it is well-known to represent the most dynamic labor market in the world (Elsby et al. 2013), but also because it provides high-frequency public use micro-data.

I use CPS micro-data, and consider all exits from employment; that is, I consider individuals who were employed in February or March 2020, and look at monthly transitions (and hence only use the continuing sample of the CPS). I group together all transitions out of employment, including out of the labor force, as it is likely that a number of exits from employment are erroneously recorded as exits from the labor force precisely because of the force from the shock.[1]

[1]The standard definition of unemployment as individuals not working but looking for work is misleading in a context where individuals are forced to stay at home. Overall, the number of individuals outside the labor force rose by 8.7 million in April relative to February.

Covid Economics 35, 7 July 2020: 220-256

I focus on the role of structural drivers of employment losses, and in particular the task characteristics of jobs, using the task content framework of Acemoglu & Autor (2011). While other papers (e.g. Cajner et al. 2020, Chetty et al. 2020, Cowan 2020) have also studied worker transitions in response to the COVID-19 shock, the task framework is useful in examining whether and why certain workers were more affected than others. The task framework looks at jobs as a collection of tasks, which require certain skills to be completed, hence breaking down the skill content of each job. It can provide a comprehensive framework for analyzing employment losses, by considering job characteristics, and shed light on the types of jobs that were most affected, going beyond demographics, education or sectoral effects. The task framework is particularly suited to studying the pandemic shock, due to the pronounced differences it is expected to have had on workers depending on the relationship of the tasks they perform to technology, location, and other actors. For instance, the most salient effect of COVID-19 on the labor markets is the rise of teleworkability, as individuals switch to working from home.[2] While we expect individuals in teleworkable occupations to be less affected, other issues may matter, such as the relative supply of particular skills. The task framework can help in analyzing these issues.

I further examine whether firms were less likely to layoff hard-to-replace employees. Labor hoarding (Guerrieri et al. 2020) may be optimal if the shock is expected to be temporary and labor search is subject to high costs. Put another way, firms may be less likely to destroy high-value matches if they can afford to keep them until the recovery, even at a short-term cost. By contrast, the uncertainty and possible credit tightening caused by the shock may lead employers to layoff easy-to-replace employees (or sever low quality-matches in general) in order to reduce costs. I use a simple measure based on average monthly frequency of turnover, which has the advantage of being easily and transparently estimated with a large sample, for both sectors and occupations.

I find strong evidence that job loss patterns resemble underlying trends. Task content is of central importance; in particular, individuals in jobs with a high non-routine cognitive analytical (NRCA) and personal (NRCP) content are especially well-protected, even if their jobs are not teleworkable. The importance of the task content, particularly for non-routine cognitive analytical tasks, is strong even after controlling for age,

[2] Brynjolfsson et al. (2020) report that half of previously employed workers work from home, 70% of which used to commute.

Covid Economics 35, 7 July 2020: 220-256

gender, race, education, sector and location. Results are robust to controlling for the extent of local lockdowns (at the CBSA level) using high-frequency data on people movements and credit card payments. These results are consistent with the findings of Cajner et al. (2020), who use administrative payroll data to find much larger employment declines for lower paid workers. Furthermore, I find patterns consistent with labor hoarding for workers in NRC-intensive jobs; controlling for the average turnover rate at the sector or occupation level wipes out any role for task content for all job types, except for those intensive in non-routine cognitive tasks. This also possibly implies a role for preemptive layoff for easy-to-replace workers. I also find an important role for complementarities for NRCA-intensive jobs; in industries with a large concentration of such jobs, lower skilled workers are less likely to lose employment than individuals in similar occupations in other industries, even after controlling for sectoral characteristics.

The results are consistent with the work of Jaimovich & Siu (2020) and Hershbein & Kahn (2018), who show that job polarization, the switch away from routine jobs towards non routine cognitive and manual jobs that has characterized labor markets in advanced economies for the past three decades (Autor 2015), occurs largely in recessions, when routine jobs are disproportionately destroyed; Jaimovich & Siu (2020) also show that polarization implies jobless recoveries, as routine jobs are permanently lower. Foote & Ryan (2015) further show that middle-skill workers are predominantly employed in volatile sectors.[3] While it is too early to tell what the implications of COVID-19 are, let alone whether the short-term effects of the shock will be persistent, a corollary of that research could be that structural changes in employment also take place primarily in recessions (Burger & Schwartz 2018). For instance, it is less costly (in the sense of foregone sales) to implement organizational changes in downturns and new firms are likely to operate with newer technology vintages than older firms.[4]

[3]It should be noted that structural change, in particular polarization, has not been found to be associated with jobless recoveries in advanced economies outside the United States (Graetz & Michaels 2017). For the Great Recession, in particular, a number of commentators (e.g. Rothstein 2017) have argued for a protracted weak demand explanation for the slow recovery.

[4]There are already news reports of sharp organizational changes brought forward as a result of the crisis. Carphone Warehouse, a large UK phone retailer announced that it would cut 2900 jobs and close all 531 standalone stores, citing changes to the industry unrelated to COVID-19 (Warrington 2020). At the same time, other organizations have halted large reorganizations, which could have a smoothing effect on the shock. HSBC announced postponement of a massive global overhaul in order to be able to function smoothly during the crisis (Crow 2020).

I finally briefly consider whether a major reallocation shock is under way. Indeed, certain sectors, related to home consumption and leisure, are booming. Barrero et al. (2020) provide anecdotal evidence of substantial hiring sprees in booming sectors, and even exchanges by firms in affected sectors to reallocate workers to booming firms. They then provide direct evidence of an important reallocation effect of the shock; using high frequency firm-survey data, they report three jobs created for every ten destroyed, a large number given the unprecedented overall contraction. They also construct a forward-looking reallocation measure, using firm employment expectations, and document a sharp rise in expected reallocation. I examine whether such a phenomenon can already be detected by publicly available data. There is no meaningful change in the share of workers switching occupations or industries within the month, suggesting that any major reallocation would be taking place within occupations or industries. Furthermore, while I confirm the relative creation and destruction magnitudes in aggregate data, by comparing the hire and separation rates in the Job Openings and Labor Turnover Survey (JOLTS), I note that the hire rate is similar to historical averages, suggesting that evidence of a major reallocation shock may be premature until new data is released.

Covid Economics 35, 7 July 2020: 220-256

2 Data

The primary data source is the Basic Monthly Sample (BMS) from the Current Population Survey (CPS), the primary source of information for the US labor market. I use primarily the 2012-2020 sample, to focus on recent trends. I also rely on data going back to 2005 for some exercises.

I base my analysis of task content on the job skills measures created by Acemoglu & Autor (2011). They use data from the O*NET (Occupational Information Network) study, which provides survey-based measures of work abilities (e.g. manual dexterity), activities (e.g. thinking creatively), work context (e.g. face-to-face discussion) and skills (e.g. social perceptiveness) for each occupation. They classify each occupation according to six tasks, collected in three broad groupings: non-routine cognitive (NRC), routine (R), and non-routine manual (RM), approximating the top, middle and lower ends of the skills distribution (conventionally defined by a one-on-one mapping

of wages to skills). Each of these is further broken down into two subgroups: non-routine cognitive analytical (NRCA) and personal (NRCP), routine cognitive (RC) and manual (RM), and non-routine manual physical (NRMPH) and personal (NRMPE).[5] It has been long recognized that a simple high-, middle-, and low-skill categorization may be insufficient to capture the intricacies of the labor market. Figure A1 in the appendix shows a mapping between skills, tasks and occupations.[6]

The proxy for teleworkability is the index created by Dingel & Neiman (2020), at the occupation level. They also use responses to the O*NET survey, and based on questions relating to work context and activities of an occupation, such as working outdoors or handling objects or machinery, they classify the feasibility of working from home, for each occupation, creating a binary indicator equal to 1 for teleworkable occupations. Teleworkable jobs include all or most jobs in computer and mathematical, education, legal, business and management occupations, and opposite for building, food preparation, construction and production occupations. The mapping from O*NET to CPS is not one-to-one, so for a few CPS code with multiple values I take the average (across 4-digit codes) and recode the indicator to 1 if the average value is above 0.5, and 0 otherwise.

To measure occupation- and sector-specific match quality, I take the CPS basic monthly micro-data files from 2005-2015 and calculate, for each occupation and industry, the share of workers that is laid-off every month. I take the median monthly value for each year to account for seasonality and then average them over the entire 15-year period in order to sweep away cyclical forces.[7] This measure, which does not include job-to-job transitions, is a proxy for the ease with which firms can replace their employees, and hence for the average job-specific match quality (for each sector and occupation). The downside of this measure is that it may be higher for declining industries. As an auxiliary indicator, I also use the gross worker flow rate, given by the

[5]Measures for each task are standardized, then summed and the sum is standardized again.

[6]Acemoglu & Autor (2011) do not explicitly use the NRMPE category in their paper, only define it in their code. I further use the refinement of Dias da Silva et al. (2019), who exclude management and professional occupations (Census codes 10 to 3599), as medical professions were ranked highly in NRMPE, as well as NRCP, making distinction between the two groups difficult. The NRMPE category is designed to capture in-person services, which are thought to be a growing part of the new economy.

[7]Median estimates are preferable to seasonally adjusted estimates because seasonal industries will have a higher turnover on an annual basis. I also compute instead seasonally adjusted estimates and the results are qualitatively similar.

Covid Economics 35, 7 July 2020: 220-256

sum of transitions out of employment, transitions into employment, and job-to-job transitions, as a share of total employment, for the industry or occupation. Job-to-job transitions are calculated as the share of workers who remain employed in consecutive months but report that they are not working for the same employer, following Fallick & Fleischman (2004). [8]

The sample is composed of individuals aged 15 or over. As the interest of this paper is on employment losses, I focus on transitions from employment to non-employment, and hence use the panelized version of CPS, using the approach of Nekarda (2009).[9] The main dependent variable (ENE) is equal to zero for individuals who work at time t-1 and remain employed at time t, and 1 for those who work at t-1 but do not at t. This implies that people who drop out of the labor force are also included, to get a more complete measure of transitions outside of employment. I only exclude individuals who voluntarily leave their jobs.[10] Focusing on those formally unemployed (individuals without work but seeking to work) may be substantially misleading in this case. Indeed, employment in April 2020 fell by 24.7 million relative to February (in seasonally unadjusted terms), but unemployment rose by 16.3 million, meaning that roughly a third of job losers exited the labor force.[11] Invariably, this choice comes with some issues, seasonality being a clear one, as workers in seasonal industries may drop out from the labor force in the low seasons; however, this should be controlled for using industry dummies and month dummies.

I work with two-level NAICS sectors, and group some similar sectors together, as some have too few observations in each month to record transitions out of employment. I end up with 35 different sectors (down from 51 used in the detailed group two-digit classification in CPS). I cluster standard errors at the occupation level, though results change little if I cluster at the sector-level instead.

[8] There is a large number of respondents who are employed in consecutive months yet as reported as "not in universe" for this question, but it is impossible to know why.

[9] I panelize using the code of Kevin Rinz, available at `kevinrinz.github.io/data.html`.

[10] Individuals who were let go and found a new job are not considered as having transitioned. This is related to the well-known time aggregation bias in CPS data (Shimer 2012), but is not a concern in this case, as I am interested in how the shock affected transitions out of employment. If individuals manage to find another job within the month, this would imply that the industry/occupation where this worker is employed is resilient enough to the shock. Exceptions are cases where individuals change occupations/industry, or individuals transitioning voluntarily to other jobs.

[11] The BLS noted that correcting for the excess transitions out of the labor force would increase the unemployment rate in April by almost 5 percentage points. See here.

Covid Economics 35, 7 July 2020: 220-256

CPS is a rotating panel, where households are recorded for four consecutive months, and where leakages from one month to other are around 6%; as such, while for two adjacent months the matching is around 70%, it is less than 50% for non-consecutive months (Rivera Drew & Warren 2014). As such, I prefer to study transitions from February to March and March to April. With such a large increase in unemployment, I expect effects in April to be mostly cumulative, justifying this approach.

Figure 1: Evolution of task content of the average worker

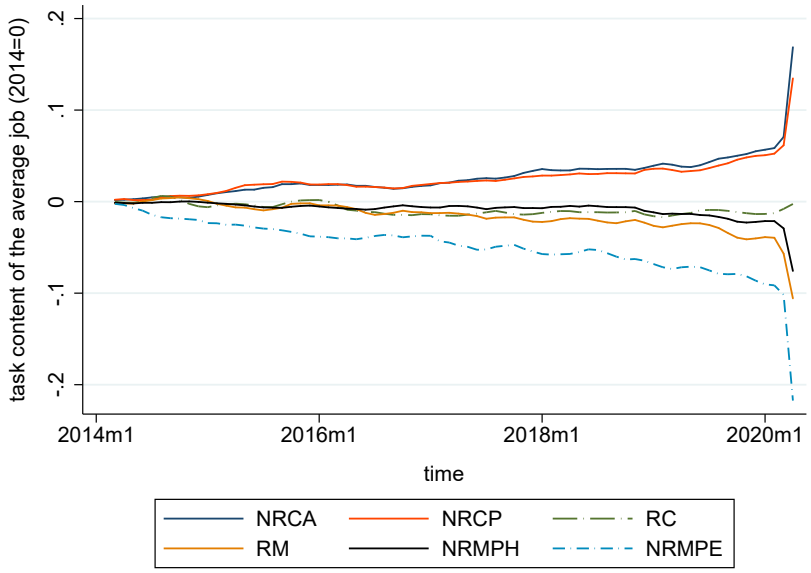

Notes: Each line shows the content of the average job for the specific task at each point in time. Task indices are standardized, so the scale is in standard deviation form Sample is from January 2015 to April 2020. Seasonally adjusted using X13-Seats. All series set to 0 in January 2014, and shown in 3-month moving averages until February 2020, in order to smooth the series. The adjusted and unadjusted trends are indistinguishable before March 2020.

The main result is previewed with graphical evidence in Figure 1, where I plot, for the 2014-2020 period, the content of the average job for each of the six tasks, across all jobs in the US economy. Each index is standardized, so the scale is in standard deviation form. Jobs intensive in non-routine cognitive tasks have been on a long-run upward trend for the past few years. For these the trend accelerated in 2019, and rose

Covid Economics 35, 7 July 2020: 220-256

substantially in March and April 2020. The opposite is true for all types of manual tasks, while the average RC task content has been mostly flat for the past five years.

These patterns indicate an increase in the share of jobs intensive in tasks typically associated with high skills, a flat profile for middle skills, and a reduction for low skills. As such, this is not indicative of job polarization. Figure 2 replicates and extends to recent years results from Autor (2015), and shows (smoothed) changes in employment shares and wages at different intervals for each occupational skill percentile (in 1980, a conventional starting point for such analyses). Unlike the 2007-2012 period, where job losses due to the financial crisis occurred in the middle of the distribution, the recovery period of 2012-2018 exhibits a positive linear relationships between skills and relative employment growth. At the same time, wage gains were relatively broadly shared. This could be indicative of skill bias, together with an increase the relative supply of high skills, preventing the skill premium from rising, as was the case in the 1970s (Katz & Murphy 1992). This is purely a conjecture though, and beyond the scope of the paper to study. Clearly though the more recent period is not associated with an increase in job polarization, but rather an evolution in employment shares increasing in skill.

3 Estimation

3.1 Task content of occupations

In this section, I examine how the probability of leaving employment depends on the task characteristics of the job, using the task indicators of Acemoglu & Autor (2011). For ease of interpretation, I convert the continuous task index to a dummy variable, equal to 1 for occupations whose values of the index is above the 80^{th} percentile of the index, and zero otherwise.[12] Recall that the dependent variable is ENE, which is equal to 1 for individuals transitioning out of employment, and 0 if they remain employed. I run several models where I regress ENE on the binary task indicator, dummies for March and April 2020, and the interaction of the task indicator with the pandemic month dummies. I run models for each of the task indicators separately.

[12] Results are qualitatively very similar with the continuous indicator as well.

Covid Economics 35, 7 July 2020: 220-256

Covid Economics 35, 7 July 2020: 220-256

Figure 2: Smoothed employment and wage changes by occupational skill percentile

(a) Employment

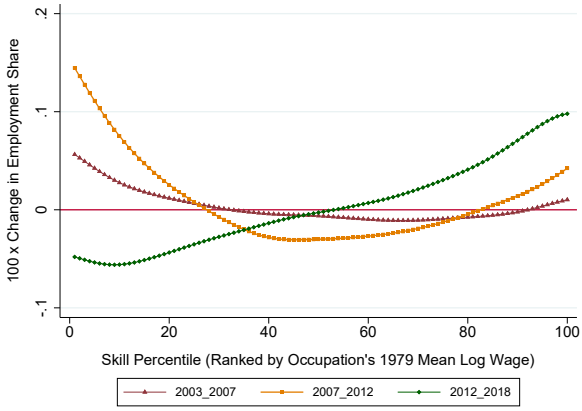

(b) Wages

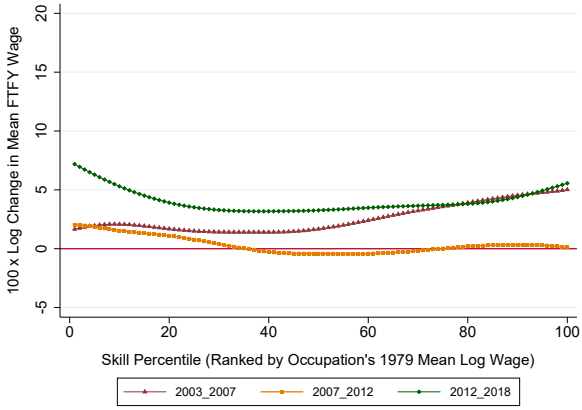

Notes: Calculated using 2003, 2006-2008, 2012, and 2014-2018 American Community Survey Integrated Public Use Microdata Series (IPUMS) files. The top figure plots changes in employment shares by 1980 occupational skill percentile rank using a locally weighted smoothing regression (bandwidth 0.75). The bottom figures plots similarly defined log wage changes for full-time, full-year workers. Skill percentiles are measured as the employment (annual hours) weighted percentile rank of an occupation's mean log wage in the Census IPUMS 1980 5 percent extract. Consistent occupation codes for Census years 1980, 1990, and 2000, and 2008 are from Autor (2015).

More formally, the model I study is:

$$ENE_{ijo,t} = \alpha_0 + \alpha_1 task_{ijo,t-1} + \alpha_2 task_{ijo,t-1} \times Mar20_t + \alpha_3 task_{ijo,t-1} \times Apr20_t + \alpha_4 Mar20_t + \alpha_5 Apr20_t + \alpha_6 \Theta_{ijo,t-1} + \gamma_j + \delta_t + \epsilon_{ijo,t}, \quad (1)$$

where $ENE_{ij,t}$ denotes the outcome for individual i in sector j, in occupation o, at time t, $task_{ijo,t-1}$ is the binary indicator for whether the individual is in a job with a high intensity of the given task at time t-1, $Mar20$ and $Apr20$ are indicators for March and April 2020, γ_j are sector dummies, δ_t are time dummies, $\Theta_{ijo,t-1}$ are lagged individual controls, and $\epsilon_{ij,t}$ is the error term. Recall that the sample is limited to individuals employed in the previous month.

As such, I examine whether task intensity matters in general for transitions out of employment, and whether the COVID-19 shock altered any prevailing pattern. Of course, losing one's job at any given month could have little consequences for one's employment trajectory, but it is important to control for time-invariant patterns, to avoid conflating them with pandemic effects. All regressions include monthly dummies to control for seasonality, and some specifications also include sector-time dummies. Individual controls include age and its square, gender, and an indicator for whether the person is Caucasian. The education variable is categorical for less than high-school, high-school/GED graduate and college graduate. This model is similar to the one used by Cajner et al. (2020).

Results are given in Table A1 in the Appendix; I report here the main coefficients in graphical form for convenience, using the odds ratio coefficient from a logistic regression, for the specification with all controls and sector-time fixed effects, for the March and April samples from 2012-2020. As the probability of losing employment rose so much in April 2020, the odds ratio, which gives the relative odds of losing employment for each value of the binary variable, has the desirable feature of scaling responses, which is convenient for a graphical representation.[13]

Each set of lines in Tables A1 responds to regression results from (1) for each of the

[13]The downside of logistic regressions is that convergence is problematic when including a large set of fixed effects (e.g. with sector-time dummies), which is why linear probability models are preferred as a baseline.

Covid Economics 35, 7 July 2020: 220-256

six task indices. It reports the α_1, α_2 and α_3 coefficients, and each column lists results from alternative specifications. For ease of interpretation, the task-COVID interaction coefficients are given in additive form to the pre-COVID coefficient, hence capturing the total effect of the pandemic shock.

The first set of results consider the role of the non-routine cognitive analytical (NRCA) content of the job. NRCA content is an important predictor of employment losses - individuals with high NRCA occupations have a 3pp lower probability of losing employment in March 2020. What is more, the effect remains strong even when demographics, education and job intensity are taken into account. Column 4 further shows results after controlling for industry-month fixed effects, which is the easiest way through which I can control for differential demand shocks, most notably those related to exposure to social distancing. The coefficient changes little and remains highly significant. This is a key finding: individuals with NRCA-intensive jobs were more protected from the pandemic shock irrespective of the industry they worked in.

NRCA content is more important for the non-teleworkable occupations (column 5). A large part of this is driven by sector-specific effects, as the inclusion of industry fixed-effects reduces substantially the magnitude of the NRCA coefficient, though it is still statistically significant. This suggests that, for the subset of teleworkable occupations, sectoral differences account for almost half of the effect, possibly indicating again that industries abundant in NRCA jobs were hit comparatively less. Of course, industry fixed effects remove time-invariant effects as well, which could be particularly important regarding job-sector-specific match value. High-NRCA individuals may be more dear to their firms, possibly due to high job-specific match surplus, or relatively low supply of the particular skills they have. For teleworkable occupations, by contrast, while NRCA-intensity by itself is less important than for non-teleworkable ones, the coefficient does not change after controlling for sector fixed-effects. Even within the subset of teleworkable jobs and within industries, individuals in NRCA-intensive occupations are less likely to be hit by an employment loss, which indicates that such a supply story is likely to hold.

The pre-COVID coefficient shows that individuals in NRCA-intensive jobs are always less likely to transition out of employment in any given month. At the same time, this effect is magnified in March, and the coefficient on the task-COVID interaction term is

Covid Economics 35, 7 July 2020: 220-256

Figure 3: Odds ratio coefficient of ENE on task content

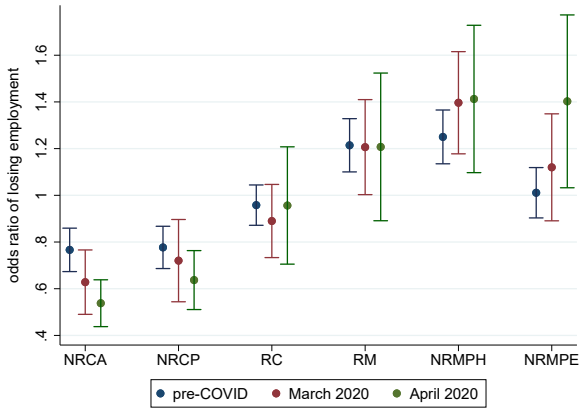

Notes: Odds ratio coefficients with 95% confidence intervals from logistic regressions of the probability of losing employment to a binary of high-task content, its interaction with a COVID dummy, relevant covariates and time-sector fixed effects. Sample is composed of February to March and March to April transitions from 2012-2020.

Covid Economics 35, 7 July 2020: 220-256

statistically significantly different from the pre-COVID coefficient. In the pre-COVID sample, individuals in jobs with high-NRCA content were 2pp less likely to transition out of employment, relative to a baseline probability of 4% for other individuals. In March, those in high-NRCA jobs were 3pp less likely to lose their jobs, while 5.5% of those in low-NRCA jobs lost their jobs. This is strong evidence that the pandemic shock served to exacerbate existing patterns, in particular providing for employment protection for jobs with high NRCA content. This pattern survives across all specifications for NRCA.

For April, the results are very similar, only now the magnitude of the coefficients is substantially higher, because the loss of employment in the aggregate is three times higher than March. Figure 3 shows graphically odds ratio coefficients of a logistic regression of the specification with fixed effects (column 4).[14] The blue dots give estimates for α_1, the red dots for α_2, and the green dots for α_3. As these are scaled (relative to the probability of losing employment for the baseline group), they help with inter-

[14]The coefficients come from regressions only including March and April for 2012-2020, as the inclusion of sector-time dummies renders the estimation unstable.

preting the magnitudes. I see that the role of NRCA task content in protecting from employment losses is even higher in April, where individuals in these jobs had 40% lower odds of losing their jobs than others.

Results are roughly similar for non-routine cognitive personal (NRCP) jobs, with the coefficient being negative and statistically significant across most specifications, with the exception of the subsample of teleworkable jobs, and for non-teleworkable jobs once controlling for sector fixed-effects. Intuitively, for jobs that can be executed remotely, NRCP elements are less relevant. The estimated coefficients are also larger in magnitude for the COVID-19 period, but by a smaller margin than NRCA, however the difference becomes statistically significant in April for all specifications.

Individuals in jobs with a high content of routine cognitive (RC) element fared better only once controlling for demographics, as they have substantially lower educational attainment than individuals in NRC-intensive jobs, and in particular for the non-teleworkable jobs. In this case as well, the magnitude is substantially higher than for the pre-COVID sample, but similar qualitatively. This may be somewhat surprising as these jobs are typically administrative and clerical (Acemoglu & Autor 2011), and are thought to be at-risk for automation. At the same time, such occupations that have still survived previous automation waves may be substantially more difficult to automate at the margin. The focus is on monthly transitions out of employment, making it difficult to discern long-term patterns. Moreover, sector-specific effects seem to drive all of the variation, and, in any case, any effect is wiped out in April.

On the other hand, routine manual (RM) and, especially, non-routine manual physical (NRMPH) jobs were at a clear disadvantage. This is likely to reflect the fact that these jobs are typically executed with physical presence, while requiring little training and have low match value; as such, they are hit both with supply and demand shocks. Again though, such patterns were clearly present for the pre-COVID sample as well. These effects are not driven by sector-specific shocks, which could be indicative of supply forces. Indeed, the coefficient for NRMPH in April becomes much larger once controlling for sector-specific effects. This is indeed consistent with employers cutting NRMPH-intensive jobs to reduce operating expenses in the face of uncertainty shocks, as such positions are relatively easy to fill-in once uncertainty recedes.[15]

[15]The very high coefficient of the NRMPH index for teleworkable occupations is due to the fact that

Covid Economics 35, 7 July 2020: 220-256

Finally, individuals employed in occupations with a high non-routine manual personal (NRMPE) content did not differ from average, irrespective of the specification. Note again that I have removed management and professional jobs from this group; otherwise the coefficient would indicate they had fared relatively better, especially in April, and in particular for non-teleworkable jobs, as the majority of them would have been in medical occupations, which were in particularly high demand due to the pandemic shock.

To recap, I find that non-routine cognitive jobs were substantially less affected by the pandemic shock. High NRCA jobs have enjoyed a premium in the form of additional job security. Individuals in routine cognitive jobs were also somewhat protected, at least given their relatively lower educational attainment, but in this case sectoral-specific effects drive the results. On the other hand, manual jobs, particularly non-routine manual physical jobs, were especially affected. In all these cases, the pandemic shock exacerbated preexisting patterns. It should be noted that the COVID-19 shock does not resemble an automation shock as such, but rather a skill-biased shock, in that lower skills seems to have been disproportionately affected. At the same time, social distancing measures necessitate labor substitution technologies by firms, and it is unclear whether this will remain once the shock passes. Finally, the results are consistent with the analysis of Cajner et al. (2020), who focus on separations by wage groups, and find that sectoral effects (which account for the differential shocks of the pandemic) played a relatively minor role.

An interesting parallel to this analysis is how the pattern of separations relates to that in the Great Recession. To investigate this, I run the baseline model (1) on a sample covering only the COVID-19 Shock (February-April 2020) and the Great Recession (October 2008-December 2009). The Great Recession started in December 2007, but it was relatively mild at its early stages, and non-seasonally adjusted employment was rising until late Summer 2008. As such, I start the after the Lehman shock, when the recession became especially deep and unemployment started to climb fast.

Results are shown in table 1. Starting with NRCA, in the Great Recession individuals employed in high-NRCA jobs had 29% lower odds of losing their job relative to others. This effect was magnified substantially during the pandemic, with indi-

the sample is very small, as this category only includes three occupations.

Covid Economics 35, 7 July 2020: 220-256

viduals in these jobs having 32% lower odds of losing their job relative to the Great Recession. As such, the protection individuals in NRCA jobs received was then substantially higher in the current recession relative to the Great Recession. There could be a number of reasons behind this result, most prominently the fact that the Great Recession was a large protracted demand shock, with employment falling continuously for over a year, while the COVID-19 shock was sharper but much shorter. Firms are hence perhaps less likely (in relative terms) to let go of their higher skilled workers in the current recession, expecting a faster increase in demand.

Workers in NRCP-intensive jobs were similarly less likely to lose their jobs in the Great Recession, but this effect remained the same in the pandemic. More interestingly, workers in RC-intensive occupations were not more likely than the rest to separate in the Great Recession, and this did not change in the current recession. Given that this group never recovered from its losses in the Great Recession, it then seems that such occupations were not disproportionately hit, but rather new jobs were not created once the recovery started. By contrast, while individuals in RM jobs were substantially more likely to be separated in the Great Recession, the difference is much weaker, and not statistically significant, for the COVID-19 episode. At the same time, this group did recover from the Great Recession, but slowly.

For NRMPH-intensive jobs, results are similar with RM jobs, as those groups overlap to some extent. Finally, by far the largest difference is in the NRMPE group, which was hit the hardest in the COVID-19 shock, as it mostly comprises of high-contact personal service jobs. In the Great Recession, this was group was not affected differently than average (and also recovered faster).

The recovery from the Great Recession was also slow (even though the expansion was the longest on record). It took almost 8 years from its peak (October 2009 to March 2017) for unemployment to reach its pre-crisis trough of 4.4%, and even though it fell as low as 3.6% in February 2020, the employment rate never recovered, due to a persistent decline in the participation rate. (Jaimovich & Siu 2020) have shown that the employment losses from the Great Recession were substantially heavier for routine occupations; they argue that the loss of routine jobs and the delayed adjustment of workers to new jobs can explain the jobless recovery. Routine cognitive occupations (using the definition of Jaimovich & Siu 2020), in particular, never recovered, staying

Covid Economics 35, 7 July 2020: 220-256

Table 1: The Great Lockdown versus the Great Recession

	(1) NRCA	(2) NRCP	(3) RC	(4) RM	(5) NRMPH	(6) NRMPE
index $\times GR$	-0.007***	-0.008***	-0.004	0.012***	0.012***	-0.000
	(-3.75)	(-4.33)	(-1.63)	(3.91)	(3.88)	(-0.15)
index $\times Mar20$	-0.013***	-0.011**	-0.005	0.007*	0.012***	0.006
	(-4.05)	(-2.49)	(-1.30)	(1.82)	(3.00)	(1.13)
index $\times Apr20$	-0.073***	-0.063***	0.008	0.037**	0.043***	0.053**
	(-8.30)	(-5.69)	(0.48)	(2.10)	(2.69)	(2.44)
N	678522	678522	678522	678522	678522	678522
Dem/phics	X	X	X	X	X	X
Education	X	X	X	X	X	X
Sector FE	X	X	X	X	X	X

The dependent variable is a dummy equal to 1 if an individual lost their job, 0 otherwise, for those employed in the previous month. Index is a dummy equal to 1 for occupations above the 80^{th} percentile of each task index. Each column shows results from regressing the dependent variable on the respective task indicator (index), period dummies, their interactions, and controls (demographics, education controls, sector fixed effects. The task indicators are denoted as defined in section 2. Errors clustered at the occupation level. The sample is October 2008-December 2009 and March-April 2020.

Covid Economics 35, 7 July 2020: 220-256

flat close to their trough of around 33 million throughout the recovery, while routine manual occupations reached their pre-Recession levels by 2018, despite a deeper fall. More generally, they argue that all recessions since 1991, when polarization was already under way, have been characterized by such a pattern.

The nature of the two shocks and the speed of transmission are completely different, but to the extent that jobless recoveries are a structural feature of polarized labor markets, it could be instructive to compare the two episodes in that regard once the dust settles and, in particular, recalls have run their course.

3.2 The role of turnover

In the previous section I showed evidence that occupational task content is an important predictor of job losses. While individuals with higher skills clearly fared better, it is still unclear whether this is because they are in relatively short supply, or whether there is also a role for firm-specific human capital in driving these results. Search and

matching frictions are large in the labor market, and employers may be willing to hoard labor in the face of a temporary shock in order to avoid losing hard-to-replace workers; at the same time, they may respond to uncertainty shocks by laying off workers who skills are in abundance and have little job-specific match capital.

I test this hypothesis by augmenting the baseline model with a turnover indicator:

$$ENE_{ijo,t} = \alpha_0 + \alpha_1 index_o + \alpha_2 turnover_{j/o} + \alpha_3 \Theta_{ijo,t-1} + \gamma_j + \epsilon_{ijo,t}, \quad (2)$$

As before, I run separate regressions for each of the task indices, in binary form. The first indicator is given by the share of workers laid off in the median month each year, averaged from 2005 to 2015, at either the occupation or the sector level. While this more directly measures the ease of replacement of employees, it would also be higher for declining occupations and industries. As such, I use an additional indicator in the form of total turnover rate, given by gross worker flows in and out of an industry or occupation over employment.[16]

Results are shown in Table 2, for each of the six task categories; columns 1-3 show results for the layoff rate, and columns 4-6 for total turnover. All models include demographics and education controls.[17] I focus on April transitions only, as over 90% of employment losses are recorded in April data.

The first column includes sectoral turnover; for all six of the task indices, the coefficient on the sectoral turnover indicator is large and statistically significant, and is essentially unchanged. This shows that sectors that tend to layoff a larger fraction of their workers at a given point in time also laid off a larger fraction of their workers due to the COVID-19 shocks; the coefficient of around 5 indicates that an additional 1pp in median layoff intensity led to 5% additional layoffs in this industry.

There are several possible implications of this finding. On the one hand, this is strong confirmation of the idea that the COVID-19 shock exacerbated pre-existing patterns. Moreover, it may indicate a role for uncertainty in exacerbating the shock; as firms

[16] An improvement would be to also have information on job tenure; unfortunately, while CPS does have a biennial Job Tenure supplement, it is not, to my knowledge, available for 2020.

[17] As the measures of turnover are based on past data, I cannot include previous years in the estimation sample. However, as the previous section made clear, the COVID-19 shock was distinctly larger to allow for a separate analysis.

Covid Economics 35, 7 July 2020: 220-256

are uncertain about the effects of the COVID-19 shock, they may have hedged by laying off workers that they deem easier to replace once the shock recedes. A similar interpretation could hold for demand concerns. At the same time, one concern is that sectoral turnover is correlated with COVID-19 demand or lockdown shocks, and so I am only capturing these effects.

This match quality interpretation is strengthened by results in column 2, which shows results for regressions where I instead include occupational turnover, and the coefficient is largely the same. Finally, column 3 includes a sector fixed effect together with the occupation turnover variable, hence controlling for sector-specific COVID-19 shocks. The turnover coefficient changes little, suggesting that even within sectors, occupations with a higher turnover suffer more job losses.[18] The specifications with total turnover in columns 4-6 show much the same picture. In this case, adding a sectoral fixed effect to the specification with the occupation turnover leads the coefficient to fall by about a quarter, but it remains large and statistically significant.

As regards the task indices, and focusing on the specifications with occupational turnover and sector fixed effects, which provide the cleanest estimates, only the coefficient for NRCA and NRCP are statistically significant; for all other tasks, the coefficients lose much of their magnitude and are not significant. The ease with which firms can replace employees working in these tasks seems to go a long way into explaining job losses during the pandemic shock. But individuals in jobs with high intensity of non-routine cognitive tasks are less likely to lose their jobs, even controlling for occupation turnover, and even controlling for sectoral characteristics and demand shocks. This could be evidence of labor hoarding behavior for hard-to-replace employees.

Covid Economics 35, 7 July 2020: 220-256

3.3 The role of sectoral structure

The above analysis has indicated a central role for the NRCA-content of jobs as being an important driver of relative employment losses due to the COVID-19 shock.

[18] An alternative way of capturing local shocks is by using Google's Mobility Reports, which track people movement across several dimensions and can provide a measure of the extent of the lockdowns. Controlling for such measures at the county level in fact strengthens the main results, but reduces the sample by over 60%, as most CPS observations do not report counties, only states or metro areas. If I instead fill in missing county values with state values and using this imputed measure as a control, results change little, while the coefficient on the mobility measure has the expected sign.

Covid Economics 35, 7 July 2020: 220-256

Table 2: Turnover

	(1)	(2)	(3)	(4)	(5)	(6)
	Terminations			Total turnover		
nr cog anal						
index	-0.064***	-0.047***	-0.034***	-0.054***	-0.019*	-0.024***
	(-5.40)	(-4.03)	(-3.86)	(-5.15)	(-1.73)	(-2.59)
turnover	4.908***	5.951***	6.354***	1.760***	2.368***	1.682***
	(4.58)	(4.45)	(6.35)	(6.47)	(6.75)	(7.50)
nr cog pers						
index	-0.043***	-0.029**	-0.033***	-0.042***	-0.013	-0.026**
	(-3.35)	(-2.30)	(-3.10)	(-3.48)	(-1.08)	(-2.54)
turnover	5.248***	6.542***	6.566***	1.846***	2.430***	1.709***
	(5.16)	(5.02)	(7.20)	(6.88)	(7.27)	(8.55)
r cog						
index	-0.016	-0.008	0.005	-0.007	0.001	0.005
	(-0.88)	(-0.48)	(0.33)	(-0.43)	(0.06)	(0.37)
turnover	5.296***	6.892***	7.205***	1.857***	2.480***	1.831***
	(5.17)	(5.26)	(7.54)	(7.05)	(7.76)	(9.15)
r man						
index	0.023	-0.007	-0.004	0.024	-0.013	0.000
	(1.17)	(-0.30)	(-0.23)	(1.29)	(-0.62)	(0.02)
turnover	5.243***	7.150***	7.283***	1.857***	2.530***	1.820***
	(4.96)	(4.71)	(6.47)	(6.73)	(7.14)	(7.60)
nr man phys						
index	-0.001	-0.049**	-0.006	0.012	-0.032	0.008
	(-0.03)	(-2.00)	(-0.31)	(0.69)	(-1.55)	(0.49)
turnover	5.390***	8.700***	7.351***	1.865***	2.606***	1.787***
	(4.60)	(5.00)	(6.46)	(6.79)	(7.60)	(8.19)
nr man pers						
index	-0.006	-0.005	-0.017	-0.006	-0.004	-0.020
	(-0.31)	(-0.24)	(-1.37)	(-0.35)	(-0.22)	(-1.48)
turnover	5.287***	6.917***	7.035***	1.859***	2.472***	1.803***
	(5.42)	(5.26)	(7.43)	(6.84)	(7.58)	(8.87)
N	33245	33245	33245	33245	33245	33245
Dem/phics	X	X	X	X	X	X
Education	X	X	X	X	X	X
Industry Turnover	X			X		
Occupation Turnover		X	X		X	X
Sector FE			X			X

The dependent variable is a binary indicator equal to 1 if an individual lost their job, 0 otherwise, for those employed in the previous month. Turnover is the median layoff probability (cols 1-3) or total turnover rate (cols 4-6) in each year, averaged over 2005-2015, either at the industry or the occupation level. Index is a binary indicator equal to 1 for individuals whose jobs are in the top quantile of the task index in *t*-1, for each task type. The sample includes transitions from March to April 2020. Errors clustered at the occupation level.

In this section, I dig deeper into this issue. Somewhat surprisingly, differential demand shocks do not seem to matter much for the magnitude of the NRCA-coefficient. At the same time, there may still be a sectoral element. In particular, it is possible that because such industries are knowledge-intensive, job-specific human capital is higher, and firms would be unwilling to sever these relationships even in bad times. A corollary would be to check whether this spills over to low-NRCA jobs as well.

I do so with the following model:

$$ENE_{ijo,t} = \beta_0 + \alpha_1 lowNRCA_{io,t-1} + \alpha_2 lowNRCA_{io,t-1} \times sectorNRCA_{j,t-1} + \alpha_3 sectorNRCA_{j,t-1} + \alpha_4 \Theta_{ij,t-1} + \gamma_j + \epsilon_{ij,t}, \quad (3)$$

where $lowNRCA$ is the negative of the NRCA index dummy, while $sectorNRCA$ is the raw sectoral average of the NRCA index. I expect a positive α_1 and negative α_3. A negative α_2 would indicate that low-skill individuals experience fewer unemployment transitions in high-NRCA sectors. For brevity, I focus on April transitions here as well, and consider March in a final specification where I pool the pre-pandemic period as well.

Results are given in table 3. Column 1 shows results from a baseline with no controls: the α_1 and α_3 coefficients have the expected signs, and the interaction coefficient α_2 is negative and significant, confirming the above intuition. One concern is that this may simply reflect differential demand shocks, as NRCA-intensive sectors experienced lower employment losses. While sectoral NRCA intensity should control for a substantial part of this, it is important to try and purge the model of demand variation. Column 2 augments the model with sectoral fixed effects. The interaction coefficient falls by almost half in magnitude, but remains negative and significant. Column 3 further augments the model with demographics and education controls, including an indicator for urban workers, and its interaction with $sectorNRCA$. The effect may be driven by differences between urban and rural areas; rural/urban location could bias results in either direction. While high NRCA sectors are typically located in cities, rural areas were presumably less hit by the COVID-19 shock. Nevertheless, the main result is little affected by the inclusion of controls.

Finally, column 4 shows results from a regression of a pooled model for the entire

Covid Economics 35, 7 July 2020: 220-256

Table 3: The role of sectoral NRCA intensity

	(1)	(2)	(3)	(4)
lowNRCA× Apr20	0.1142***	0.1064***	0.0771***	0.0943***
	(7.53)	(8.56)	(6.34)	(7.85)
lowNRCA× sectorNRCA	-0.0867***	-0.0452***	-0.0401***	-0.0412***
× Apr20	(-3.91)	(-2.91)	(-2.69)	(-2.74)
sectorNRCA× Apr20	-0.0391***			
	(-3.82)			
lowNRCA × pre-COVID				0.0089***
				(6.61)
lowNRCA × Mar20				0.0198***
				(4.82)
lowNRCA × sectorNRCA				-0.0077***
× pre-COVID				(-3.41)
lowNRCA × sectorNRCA				-0.0121*
× Mar20				(-1.66)
N	33245	33245	32930	4204153
Dem/phics			X	X
Education			X	X
Sector FE		X	X	
Sector-Time FE				X

The dependent variable is a binary indicator equal to 1 if an individual lost their job, 0 otherwise, for those employed in the previous month. Errors clustered at the occupation level. $lowNRCA$ is a dummy equal to 1 for occupations below the 80^{th} percentile of the NRCA index, and $sectorNRCA$ is the sectoral average of the NRCA index. The sample in columns 1-3 is April 2020; in 4 it is comprised of all months in 2012-2020.

Covid Economics 35, 7 July 2020: 220-256

2012-2020 sample, with appropriate interaction terms for the variables of interest. As before, coefficients for March and April are given in additive form relative to pre-COVID values. In this case as well, these patterns held in the pre-COVID sample, but at a substantially smaller degree. In terms of its relative magnitude, it appears that employment losses from the pandemic shock are scaled up versions of usual patterns. The triple interaction coefficient of interest retains its size and significance even for March, and is substantially larger in April.

These results suggest a complementarity between high- and low-skilled workers in sectors with a relatively high share of NRCA-intensive jobs. Aghion et al. (2017) present evidence of such a channel for innovation; the wage premium earned by employees of innovative firms is higher for low- than high-skilled employees. They build a model to rationalize their finding, and one prediction of the model is that in

fact worker quality is more firm-specific for low- than high-skill workers, implying longer job tenures for low-skilled employees in innovative firms, as firms will invest more resources to train them.

3.4 Controlling for economic activity with high-frequency data

In the previous exercises, I controlled for differential shocks using different combinations of fixed effects, which provide a transparent means of controlling for such issues, at the cost of losing variation. An alternative way of accounting for local economic activity is by using some of the newly released high-frequency data. In this section, I revisit the analysis of the previous sections using two such datasets. The first is Google's Mobility Reports, which track people's using GPS data. They track people's movement across several dimensions relative to pre-COVID values and reports percentage changes in visits, providing a proxy for the extent of the lockdowns. The best coverage and most relevant variables for my purposes are given by the retail and recreation (RER) and workplaces (WORK) indices; for instance, March 22 saw a 44% reduction in visit to retail and recreation places relative to the January 3-February 6 average. The second source is the spending data of the Opportunity Insights Economic Tracker project of Chetty et al. (2020), which track credit/debit card payments at the zip code level, as collected by a private provider, capturing around 10% of total transactions, with a higher representation of expenditure on accommodation, food and retail. The authors show that this data tracks aggregate spending quite well. The measure I use the percentage change in total spending relative to January (SPEND).

I aggregate county-level data to the core-based statistical area (CBSA) level (using population weights), because county of residence is reported for only around a third of respondents. This likely implies a substantial loss of variation; Chetty et al. (2020) report very large variations in spending changes even across adjacent zip codes. Both datasets come in daily frequencies; I aggregate daily data in lower frequencies following the timing of CPS interviews, and so April values reflect averages from 12 March to 11 April. I use binary values of the RER, WORK and SPEND variables, set to 1 for counties with above median values of each respective variable.

I estimate a variant of (1) for April only, with sectoral fixed effects and all controls

Covid Economics 35, 7 July 2020: 220-256

Table 4: High-frequency controls

	(1)	(2)	(3)	(4)	(5)	(6)
Non-routine cognitive		analytical			personal	
index	-0.063***	-0.068***	-0.074***	-0.052***	-0.057***	-0.058***
	(-5.64)	(-6.70)	(-7.01)	(-4.08)	(-5.01)	(-4.76)
high	-0.028***	-0.032***	-0.044***	-0.025***	-0.028***	-0.037***
	(-4.47)	(-4.34)	(-5.75)	(-4.25)	(-3.73)	(-4.98)
index× high	0.015	0.027***	0.039***	0.003	0.014	0.015
	(1.33)	(2.66)	(3.85)	(0.26)	(1.34)	(1.41)
Routine		cognitive			manual	
index	-0.003	-0.008	0.000	0.034	0.035*	0.037*
	(-0.17)	(-0.53)	(0.02)	(1.52)	(1.72)	(1.73)
high	-0.025***	-0.026***	-0.033***	-0.023***	-0.023***	-0.032***
	(-4.18)	(-3.78)	(-4.73)	(-4.02)	(-3.23)	(-4.61)
index×high	-0.000	0.010	-0.008	-0.011	-0.013	-0.015
	(-0.02)	(0.61)	(-0.49)	(-0.74)	(-0.93)	(-1.02)
Non-routine manual		physical			personal	
index	0.058***	0.059***	0.060***	0.047*	0.060***	0.064**
	(2.85)	(2.91)	(2.79)	(1.89)	(2.63)	(2.36)
high	-0.018***	-0.019***	-0.028***	-0.025***	-0.019***	-0.027***
	(-3.18)	(-3.04)	(-4.38)	(-4.51)	(-3.20)	(-4.57)
index×high	-0.038***	-0.034*	-0.034*	0.002	-0.028	-0.037*
	(-2.61)	(-1.73)	(-1.82)	(0.12)	(-1.50)	(-1.89)
N	23997	25103	25103	23997	25103	25103
SPEND	X			X		
WORK		X			X	
RER			X			X

The dependent variable is a dummy equal to 1 if an individual lost their job, 0 otherwise, for those employed in the previous month. Index is a dummy equal to 1 for occupations above the 80^{th} percentile of each task index, and high is a dummy equal to one for above median values of SPEND, WORK and RER, as denoted. Each panel shows results from regressing the dependent variable on the respective task indicator (index), the respective activity indicator (high), their interactions, and controls (demographics, education controls, sector fixed effects. Errors clustered at the occupation level. The sample is April 2020.

included. Results are given in table 4. Each panel reports results from each pair of broad task categories: NRC analytical and personal in the top panel, R cognitive and manual and the middle panel, NR physical and personal in the bottom panel. The task indicator is denoted as $index$ and the activity indicator denoted as $high$ (standing for high levels of activity, even though it fell in almost all areas).

Covid Economics 35, 7 July 2020: 220-256

The coefficient for $high$ is negative and significant for all activity indicators. In areas with above median levels of (relative) activity, transitions out of employment were around 3pp smaller. The effect is largest for the RER indicator, which is likely to cover mostly employees in related industries. This is consistent with the results of Chetty et al. (2020), who report that more affluent localities exhibited a larger reduction in both spending and employment; their data only cover low income employment.

This view is reinforced by results for non-routine manual jobs, where the index-high interaction coefficient is positive and significant for most specifications. This implies that the likelihood of transitioning out of employment is further mitigated for individuals in these occupations in high activity areas, relative to others. Especially for the non-routine manual occupations, the separation rates is equal to other occupations, on average, in high activity areas, as the three coefficients sum to zero. On the other hand, the index-high interaction is positive and significant for the WORK and RER indicators for NRCA-intensive jobs, meaning that the relative protection such jobs offer is smaller in high activity areas. The interactions are small and not significant for the other task indicators, suggesting that individuals in these occupations fared similarly (relative to other occupations) in high and low activity areas.

Overall, controlling directly for changes in economic activity does not materially change our results, but does provide more subtle insights for the professions that were affected the most (non-routine manual) and the least (non-routine cognitive analytical) from the shock, showing that these effects were driven primarily by areas where activity slowed down by more. These controls cannot answer whether demand or supply effects are driving the results, as either could contribute to lower activity.

3.5 Is there a reallocation wave under way?

The scale of job losses in the United States as a result of the COVID-19 have given rise to a debate regarding the extent to which jobs lost can be recovered. A combination of search costs, defaults, uncertainty and structural change may imply the permanent loss of many jobs, as well as a long delay in creating new jobs. The results of Barrero et al. (2020) certainly point towards such a story. Using survey data, they show an increase in *expected* excess reallocation (sum of creation and destruction minus net

Covid Economics 35, 7 July 2020: 220-256

employment change) of 2.4 times relative to pre-pandemic averages, at the firm level. They document that for every 10 jobs lost, 3 have already been created. Combing their results with historical data on recalls, they estimate that 42% of staffing reduction will lead to permanent job losses.

Even though the vast majority of layoffs in April are listed in CPS as temporary, other sources also indicate that a substantial part of these matches may be destroyed. For instance, although 91% of UI claimants in California expected to be recalled to their jobs in the late March, this figure had fallen to 69% by early May.[19]

The large reallocation patterns reported by Barrero et al. (2020) occur at relatively high frequencies, substantially increasing the time-aggregation bias of CPS. What the CPS can shed light on, however, is whether there is a noticeable shift of workers across occupations or sectors. This could include employed workers who switch employers or unemployed workers who take up a new job in a different occupation or sector relative to their previous employment. An increase in the fraction of workers who switch occupation or industry would point to increased reallocation. I calculate occupation and industry switchers as the fraction of workers who switch in adjacent months, both for continuing workers and for entrants from unemployment (as CPS codes the last known industry and occupation for unemployed workers). [20] I extend the sample to May, when employment rose substantially, and could give an additional data point for reallocation. I exclude individuals recalled from a temporary layoff; while Fujita (2018), in a similar exercise, does include them, the spike in recalls from temporary layoffs in May would bias the switching rate downwards. I calculate job-to-job transitions as the share of workers who are employed in consecutive months but switch employers from one month to the other.

I plot these measures in figure 4. The blue and red lines show the (non-seasonally unadjusted) shares of workers employed in a given month who switch occupations (4-digit level) and sectors (2-digit level) relative to the previous month. There is no clearly discernible increase for either series. Following an upward trend coinciding

[19]Source: California Policy Lab.

[20]Focusing only on workers employed in consecutive months yields a qualitatively very similar picture. Comparing the same month across years is complicated, as occupation codes changed in 2020, resulting in a discrete jump in January. As such, an analysis of consecutive months, removing January 2020, is more informative. I also remove June and July 2015, which exhibit a very large spike in the share of individuals who switch occupations and industries.

Covid Economics 35, 7 July 2020: 220-256

with the recovery, the share of workers employed in consecutive months that switch occupations has hovered around 7% since 2016, and 4.5% for industries. The green line shows the occupation switching rates using the coarse 22 CPS major occupation classification, to correct for the erroneous switches problem notes by Fujita (2018). The coarse switching rate is naturally lower, but the pattern is very similar, and also shows little movement during the COVID-19 shock. Finally, the yellow line shows the share of job-to-job transitions. Again, the measure is around 2% across the horizon considered, and does not move perceptibly in April or May.[21] Note that purpose here is only to see whether there is a spike in switching in April, not to analyze cyclical properties of these series, which is why I only plot non-seasonally adjusted rates.

There are a number of caveats to this analysis. First, this is not a properly defined measure of reallocation, as it focuses exclusively on the sub-sample of those currently employed, hence ignoring job losers who did not find a new job within the month. Second, it is well-known that vast majority (85-90%) of reallocation occurs within industries and localities (Davis & Haltiwanger 1992), and Barrero et al. (2020) give an excellent illustration of this phenomenon. At the same time, given the size of the shock, I expect that such a major reallocation shift should have involved at least some movement in the series plotted in figure 4. A final major caveat to this analysis compared to Barrero et al. (2020), is that it is based on data for an on-going shock. By contrast, Barrero et al. (2020) rely on a forward-looking measure, and argue that such a measure is preferable in this context, because creation lags destruction in major reallocation shocks for at least a year.[22] What this exercise can say though is that, to the extent that a large realignment of the workforce is already underway, it is occurring largely within industries and occupations.

An alternative source of information on reallocation in high frequencies is JOLTS, a survey of 16,000 establishments designed to capture worker flows in the US labor market. Table 5 shows seasonally adjusted sectoral hire and separation rates for the

[21] This measure can also include workers who are laid off and find a job prior to the CPS interview. The denominator consists only of individuals with a valid response to the relevant question, as a large number of eligible interviewees have missing values for this question, for unknown reasons. However, the stable value around 2% is consistent with estimates by Bosler & Petrosky-Nadeau (2016) with different data.

[22] Note also that the higher sampling bias (due to an increase in non-response rates) in CPS may be especially acute in this case.

Covid Economics 35, 7 July 2020: 220-256

Figure 4: Share of workers switching occupations and sectors

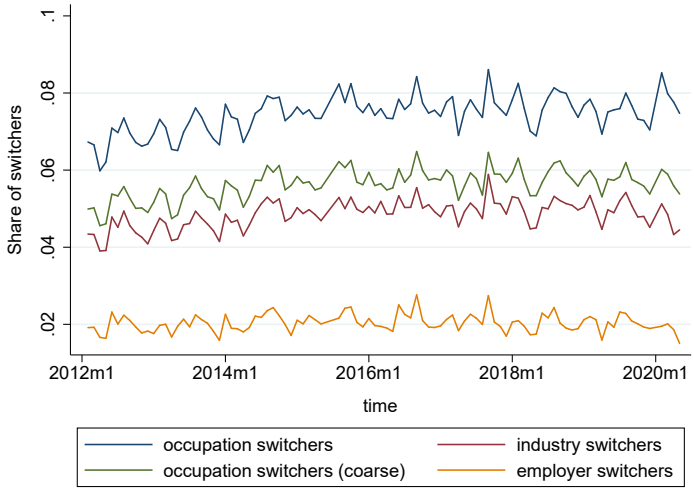

Notes: The chart shows the share of workers employed in consecutive months who switch occupations (4-digit, blue line) or sectors (2-digit, red line) from one month to another. The green line uses the CPS 22-occupation classification. The yellow line shows the share of workers who switch employers. The sample includes all CPS Basic Monthly Sample files from 2012-2020.

Covid Economics 35, 7 July 2020: 220-256

non-farm private economy. A number of issues stand out.

First, hires kept their pace relatively unabated in March, but separations increased by a factor of 2.5 relative to February. Non-farm private sector hire rate fell modestly, to 3.7%, from 4.2% in February. Seasonal adjustment may be misleading given the size of the shock; the fall in the non-seasonally adjusted hire rate was ever smaller, at 0.1pp. The separation rate, by contrast, rose substantially, from 4% to 11.1%. The layoff rate (which excludes quits) rose from 1.4% to 8.8%. Both rates fell in April, but relative to February, the relative change in the separation rate is still larger (in log points).

Second, sectoral differences in the change in hiring behavior is relatively muted. Hiring in fact increased in non-durable manufacturing and fell by only 0.2pp in retail; on the other hand, it fell substantially in entertainment and in accommodation and food services. The dispersion of the change in hiring rates across sectors was relative compressed overall; the upshot is that, while hiring was curbed, it was still taking

Table 5: Hires and Separation Rates by sector, February-April 2020

	Hire rate				Separation rate			
	(1) Feb	(2) Mar	(3) Apr	(4) Δ	(5) Feb	(6) Mar	(7) Apr	(8) Δ
Mining	3.4	2.7	1.7	-1.2	2.8	6.0	10.4	5.4
Construction	5.1	5.1	2.9	-1.1	4.7	9.9	11.9	6.2
Durable manuf	2.4	2.0	2.0	-0.4	2.3	5.9	6.7	4.0
Nondurable goods manuf	2.9	2.9	3.7	0.4	2.8	6.9	5.3	3.3
Wholesale trade	2.6	2.3	2.2	-0.4	2.6	4.2	4.7	1.8
Retail trade	5.2	4.9	5.1	-0.2	5.2	10.6	8.8	4.5
Transportation/utilities	4.1	3.8	3.2	-0.6	4.1	8.1	6.8	3.4
Information	3.3	2.6	1.3	-1.4	3.0	3.7	5.6	1.7
Finance and insurance	2.5	2.3	1.6	-0.6	2.3	2.6	1.7	-0.2
Real estate	3.4	2.9	1.5	-1.2	2.5	7.7	10.7	6.7
Professional and business services	5.1	5.1	3.5	-0.8	5.0	8.0	6.4	2.2
Education	2.5	2.4	2.0	-0.3	2.4	8.6	8.7	6.3
Health care	3.2	2.8	2.4	-0.6	2.9	6.9	6.1	3.6
Arts and entertainment	6.8	5.2	1.9	-3.3	6.3	24.5	25.6	18.8
Accommodation and food services	6.4	3.9	3.9	-2.5	6.1	34.1	23.0	22.5
Other services	3.8	2.8	2.4	-1.2	3.7	16.4	19.0	14.0
Total private non-farm	4.2	3.7	3.0	-0.9	4.0	11.1	8.7	5.9

Source: JOLTS. Sample includes all monthly data from February to April 2020, in non-seasonally adjusted form, except for the bottom line. The hire and separation rates are defined as hires or separation over employment, multiplied by 100. Columns 4 and 8 show the difference between the average value for March and April relative to February.

Covid Economics 35, 7 July 2020: 220-256

place, and in fact remained higher than its trough of 3.1% in June 2009, in March and slightly below in April.

Third, there are enormous differences across sectors in separations. The dispersion of the change in the rate is an order of magnitude higher relative to hires. Entertainment and accommodation and food services, in particular, experienced separations of over half of their February employment, given their particular exposure to the pandemic shock; on the other hand, for the finance and insurance and insurance sector, separations in fact fell in April, and total employment shrank by only 0.7% relative to February.

Fourth, the ratio of hires to separations is 0.34 for both months, almost identical to the survey findings of Barrero et al. (2020). At the same time, the hire rate is relatively close to its long-term sample average of 4%.[23] This is an interesting finding in

[23] Davis et al. (2010) noted that earlier renditions of JOLTS underestimated gross workers flows in

its own right; while the ratio of separations to hires was typically around 1, reaching a maximum of 1.25 in April 2009, it jumped to 3 in March and 2.9 in April, but this was mostly driven by the separation rate: relative to February, the log change in the separation rate was 8 times larger than the log change in the hire rate in March, and 2.3 times in April. As such, while the surprisingly strong pace of hiring in the face of such a large shock may indeed indicate the emergence of a large reallocation wave, the fact that the hire rate is so close to its usual level implies this conclusion could be premature. On the other hand, a disproportionate fraction of hires comes from young firms; given the large reduction in business formation and the tighter borrowing constraints small firms face, it is possible that the aggregate hire rate may in fact mask an increasing hire rate for large firms, which could indeed imply higher reallocation. Future data releases will likely shed light to this puzzle.

Covid Economics 35, 7 July 2020: 220-256

4 Conclusion

I examined the pattern of short-term job losses from the COVID-19 shock in the United States, using the task approach to labor markets. This framework can shed light on what types of jobs were most affected from the shock, and which ones were most resilient, offering a complementary analysis of the labor market effects of the shock to work using granular data to more precisely capture the demography, geography and scale of the shock (Cajner et al. 2020, Chetty et al. 2020). I find that job task content is an important predictor of job losses, even controlling for demographics, education, industry, teleworkability and local economic activity. The coefficients on the task content indicators become insignificant once I control for occupational turnover (which proxies for the ease of replacing workers), except for indicators for jobs intensive in non-routine cognitive (NRC) tasks; this suggests that layoffs followed usual patterns, but firms hoarded workers with NRC skills.

I also take a stab at gauging whether reallocation patterns can already be seen in publicly available data, CPS and JOLTS. In CPS I show that there is discernible uptick in the share of workers switching occupations, while, from JOLTS, I find that the hire

public use files. Following their recommendations, BLS revised the methodology in 2009 and moderated the discrepancy, though differences remain.

rate moved much less than the separation rate, and remains relatively close to its historical average, at only 0.1pp below its trough in the Great Recession, even though unemployment is currently much higher. Taken together with the results of Barrero et al. (2020), these findings imply that reallocation may have yet to take off, but if it has, it is taking place within occupation and industries. At the same time, it is important to note that at least some of the reallocation away from contact-intensive sectors (hospitality, personal services) is likely to be temporary, until vaccines become available and relevant investments take place (e.g. separating panels, germ-resistant surfaces). This makes it more likely that there will be a protracted underutilization of resources, relative to even a gradual reallocation shock, and hence a larger scope for demand-supporting policies. More broadly, the surprisingly small change in the hire rate in the face of a massive increase in the separation and layoff rates is an important area for future work.

Overall, it remains to be seen whether the patterns identified in this paper will persist. A new burgeoning literature, an offspring of the slow recovery from the financial crisis, argues for structural change as taking place in recessions and amplifying them (Chodorow-Reich & Wieland 2020, Jaimovich & Siu 2020). While COVID-19 is a singular shock, and past experience may not be especially useful, it has arguably already ushered in a wave of substantial technological diffusion and paradigm-shifting change in work patterns. Baldwin (2020) lays out the argument that the jobs that will survive are those relying on social cognition and interpersonal skills, which artificial intelligence or offshoring cannot handle, precisely the type of transition expected before the COVID-19 shock.

More speculatively, there is some evidence that such transition may have been foreshadowed by financial markets: Pagano et al. (2020) show that not only has portfolio reallocation taken place away from stock firms in sectors heavily exposed to the pandemic, towards resilient stocks, but that highly resilient firms substantially outperformed low-resilience firms for the six years prior to the pandemic. They speculate that investors may have been pricing such a risk, but another possibility is technological; given that the currently dominant technologies are also more resilient to the pandemic, firms on the forefront of production or intensive usage of these technologies were both possibly enjoying higher stock returns before the COVID-19 shock and

Covid Economics 35, 7 July 2020: 220-256

were in a better position to withstand the shock, for instance by transitioning to a new work paradigm or having sufficient cash buffers.

Covid Economics 35, 7 July 2020: 220-256

References

Acemoglu, D. & Autor, D. (2011), Skills, Tasks and Technologies: Implications for Employment and Earnings, *in* O. Ashenfelter & D. Card, eds, 'Handbook of Labor Economics', Vol. 4 of *Handbook of Labor Economics*, Elsevier, chapter 12, pp. 1043–1171.

Aghion, P., Bergeaud, A., Blundell, R. & Griffith, R. (2017), 'The Innovation Premium to Low Skill Jobs', Mimeo.

Autor, D. (2015), 'Why Are There Still So Many Jobs? The History and Future of Workplace Automation', *Journal of Economic Perspectives* **29**(3), 3–30.

Baldwin, R. (2020), 'Covid, hysteresis, and the future of work'.

Barrero, J. M., Bloom, N. & Davis, S. J. (2020), 'COVID-19 Is Also a Reallocation Shock', BFI Working Paper.

Bosler, C. & Petrosky-Nadeau, N. (2016), 'Job-to-Job Transitions in an Evolving Labor Market', *FRBSF Economic Letter* .

Brynjolfsson, E., Horton, J. J., Ozimek, A., Rock, D., Sharma, G. & TuYe, H.-Y. (2020), 'COVID-19 and Remote Work: An Early Look at US Data', NBER Working Paper No. 27344.

Burger, J. D. & Schwartz, J. S. (2018), 'Jobless Recoveries: Stagnation Or Structural Change?', *Economic Inquiry* **56**(2), 709–723.

Cajner, T., Crane, L. D., Decker, R. A., Grigsby, J., Hamins-Puertolas, A., Hurst, E., Kurz, C. & Yildirmaz, A. (2020), 'The U.S. Labor Market during the Beginning of the Pandemic Recession', NBER Working Paper No. 27159.

Chetty, R., Friedman, J., Hendren, N. & Stepner, M. (2020), 'How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data'.

Chodorow-Reich, G. & Wieland, J. (2020), 'Secular Labor Reallocation and Business Cycles', *Journal of Political Economy* pp. 2245–2287.

Covid Economics 35, 7 July 2020: 220-256

Cowan, B. W. (2020), 'Short-run Effects of COVID-19 on U.S. Worker Transitions', NBER Working Paper No. 27315.

Crow, D. (2020), 'HSBC will be forced to delay radical restructuring if crisis deepens', https://www.ft.com/content/e97f4722-693c-11ea-800d-da70cff6e4d3.

Davis, S. J., Faberman, R. J., Haltiwanger, J. C. & Rucker, I. (2010), Adjusted Estimates of Worker Flows and Job Openings in JOLTS, *in* 'Labor in the New Economy', NBER Chapters, National Bureau of Economic Research, Inc, pp. 187–216.

Davis, S. J. & Haltiwanger, J. (1992), 'Gross Job Creation, Gross Job Destruction, and Employment Reallocation', *The Quarterly Journal of Economics* **107**(3), 819–863.

Dias da Silva, A., Laws, A. & Petroulakis, F. (2019), 'Hours of work polarisation?', ECB Working Paper No. 2324.

Dingel, J. & Neiman, B. (2020), 'How Many Jobs Can be Done at Home?', *Covid Economics: Vetted and Real-Time Papers* .

Elsby, M. W. L., Hobijn, B. & Şahin, A. (2013), 'Unemployment Dynamics in the OECD', *The Review of Economics and Statistics* **95**(2), 530–548.

Fallick, B. & Fleischman, C. (2004), 'Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows', Federal Reserve Board, Finance and Economics Discussion Series Working Paper 2004-34.

Foote, C. L. & Ryan, R. W. (2015), 'Labor-Market Polarization over the Business Cycle', *NBER Macroeconomics Annual* **29**(1), 371–413.

Fujita, S. (2018), 'Declining labor turnover and turbulence', *Journal of Monetary Economics* **99**(C), 1–19.

Graetz, G. & Michaels, G. (2017), 'Is modern technology responsible for jobless recoveries?', *American Economic Review* **107**(5), 168–73.

Guerrieri, V., Lorenzoni, G., Straub, L. & Werning, I. (2020), 'Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?', NBER Working Paper No. 26918.

Covid Economics 35, 7 July 2020: 220-256

Hershbein, B. & Kahn, L. B. (2018), 'Do recessions accelerate routine-biased technological change? evidence from vacancy postings', *American Economic Review* **108**(7), 1737–72.

Jaimovich, N. & Siu, H. E. (2020), 'Job Polarization and Jobless Recoveries', *The Review of Economics and Statistics* .

Katz, L. & Murphy, K. M. (1992), 'Changes in relative wages, 1963–1987: Supply and demand factors', *The Quarterly Journal of Economics* **107**(1), 35–78.

Nekarda, C. (2009), 'A Longitudinal Analysis of the Current Population Survey', Federal Reserve Board of Governors Working Paper.

Pagano, M., Wagner, C. & Zechner, J. (2020), 'Disaster Resilience and Asset Prices'.

Rivera Drew, Julia A., F. S. & Warren, J. R. (2014), 'Making Full Use of the Longitudinal Design of the Current Population Survey: Methods for Linking Records Across 16 Months'.

Rothstein, J. (2017), 'The great recession and its aftermath: What role for structural changes?', *RSF: The Russell Sage Foundation Journal of the Social Sciences* **3**(3), 22–49.

Shimer, R. (2012), 'Reassessing the Ins and Outs of Unemployment', *Review of Economic Dynamics* **15**(2), 127–148.

Warrington, J. (2020), 'Carphone warehouse to axe 3,000 jobs as it shuts all standalone stores', https://www.cityam.com/carphone-warehouse-to-axe-3000-jobs-as-it-shuts-all-standalone-stores/.

Covid Economics 35, 7 July 2020: 220-256

A Appendix

Figure A1: Mapping of skills, tasks and occupations

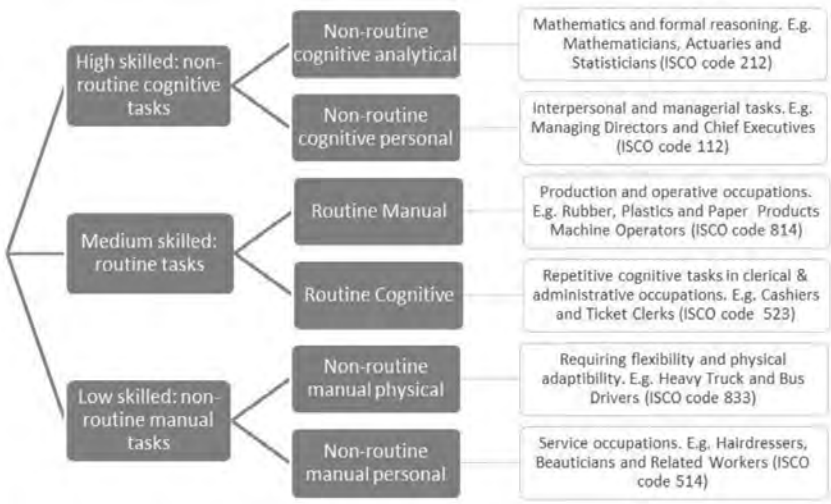

Source: Dias da Silva et al. (2019).

Covid Economics 35, 7 July 2020: 220-256

Covid Economics 35, 7 July 2020: 220-256

Table A1: Regressions for task content

	(1)	(2)	(3)	(4)	(5) Non-TW	(6) Non-TW	(7) TW	(8) TW
nr cog anal								
pre-covid	-0.021***	-0.014***	-0.005**	-0.005***	-0.024***	-0.008***	-0.004	-0.004**
	(-6.86)	(-5.72)	(-2.22)	(-3.71)	(-7.64)	(-4.54)	(-1.29)	(-2.18)
March	-0.030***	-0.023***	-0.015***	-0.014***	-0.043***	-0.019***	-0.013**	-0.013***
	(-6.53)	(-5.57)	(-3.73)	(-4.15)	(-6.33)	(-3.65)	(-2.25)	(-2.71)
April	-0.120***	-0.113***	-0.105***	-0.073***	-0.100***	-0.036	-0.068***	-0.054***
	(-8.38)	(-8.21)	(-7.70)	(-8.35)	(-3.57)	(-1.33)	(-4.44)	(-4.67)
nr cog pers								
pre-covid	-0.018***	-0.011***	-0.005**	-0.006***	-0.026***	-0.014***	-0.001	-0.003*
	(-6.09)	(-4.73)	(-2.22)	(-4.28)	(-8.99)	(-6.16)	(-0.21)	(-1.80)
March	-0.021***	-0.014***	-0.008*	-0.012***	-0.028***	-0.018*	-0.002	-0.009*
	(-4.22)	(-3.33)	(-1.92)	(-2.69)	(-3.90)	(-1.73)	(-0.43)	(-1.74)
April	-0.086***	-0.079***	-0.074***	-0.063***	-0.094***	-0.070***	-0.028*	-0.037***
	(-5.51)	(-5.30)	(-5.01)	(-5.77)	(-2.69)	(-2.62)	(-1.90)	(-2.94)
r cog								
pre-covid	-0.000	-0.005*	-0.005***	-0.002	-0.004	-0.003	-0.003	0.001
	(-0.06)	(-1.81)	(-2.67)	(-1.17)	(-0.48)	(-1.12)	(-0.78)	(0.59)
March	-0.006	-0.011**	-0.011***	-0.005	-0.015	-0.009**	-0.004	0.002
	(-0.94)	(-2.38)	(-2.73)	(-1.30)	(-1.50)	(-1.99)	(-0.57)	(0.29)
April	-0.000	-0.005	-0.005	0.008	-0.032	-0.014	0.037*	0.043**
	(-0.01)	(-0.21)	(-0.24)	(0.47)	(-0.97)	(-0.60)	(1.65)	(2.33)
r man								
pre-covid	0.017***	0.016***	0.006**	0.008***	0.010*	0.011***	0.021*	0.021**
	(3.42)	(4.56)	(2.14)	(3.94)	(1.96)	(4.72)	(1.87)	(2.05)
March	0.017***	0.016***	0.007	0.008**	0.008	0.010*	0.017	0.020*
	(2.72)	(2.94)	(1.39)	(1.97)	(1.16)	(1.96)	(1.44)	(1.79)
April	0.061***	0.060***	0.051**	0.038**	0.010	0.009	0.134**	0.116**
	(2.82)	(2.97)	(2.56)	(2.16)	(0.35)	(0.43)	(2.36)	(2.25)
nr man phys								
pre-covid	0.009**	0.016***	0.008***	0.008***	0.001	0.010***	0.014***	0.017***
	(2.08)	(4.73)	(2.78)	(3.96)	(0.25)	(3.81)	(4.13)	(4.29)
March	0.011*	0.018***	0.010*	0.013***	-0.000	0.016***	0.045***	0.050***
	(1.75)	(3.34)	(1.94)	(3.36)	(-0.03)	(3.62)	(5.67)	(5.07)
April	0.024	0.031*	0.023	0.044***	-0.033	0.011	0.027	0.051***
	(1.25)	(1.74)	(1.30)	(2.75)	(-1.30)	(0.60)	(1.61)	(3.44)
nr man pers								
pre-covid	0.006	0.000	-0.001	0.001	-0.002	-0.003	0.006	0.005
	(1.13)	(0.10)	(-0.30)	(0.27)	(-0.38)	(-1.20)	(0.64)	(0.97)
March	0.012	0.007	0.005	0.006	0.006	0.005	0.004	0.004
	(1.45)	(0.94)	(0.77)	(1.15)	(0.63)	(0.76)	(0.29)	(0.50)
April	0.093***	0.087***	0.086***	0.053**	0.076*	0.054**	0.035	0.007
	(2.67)	(2.62)	(2.59)	(2.45)	(1.82)	(2.07)	(0.83)	(0.37)
N	4204153	4204153	4204153	4204153	2505587	2505587	1698566	1698566
Dem/phics		X	X	X	X	X	X	X
Education			X	X	X	X	X	X
Sector-Time FE				X		X		X

The dependent variable is an indicator equal to 1 if an individual lost their job, 0 otherwise, for those employed in the previous month. Each panel shows results from regressing the dependent variable on the respective task indicator, interacted with time dummies. The task indicator is a binary equal to 1 for individuals whose jobs are in the top quantile of the task index in the previous month, for each task type. Regressions include monthly dummies. Errors clustered at the occupation level. TW columns include the subsample of teleworkable occupations, and opposite for Non-TW. Sample is 2012-2020.

Covid Economics **Issue 35, 7 July 2020**

Weather, social distancing, and the spread of COVID-19[1]

Daniel J. Wilson[2]

Date submitted: 1 July 2020; Date accepted: 3 July 2020

Using high-frequency panel data for U.S. counties, I estimate the full dynamic response of COVID-19 cases and deaths to exogenous movements in mobility and weather. I find several important results. First, weather and mobility are highly correlated and thus omitting either factor when studying the COVID-19 effects of the other is likely to result in substantial omitted variable bias. Second, temperature is found to have a negative and significant effect on future COVID-19 cases and deaths, though the estimated effect is sensitive to which measure of mobility is included in the regression. Third, controlling for weather, overall mobility is found to have a large positive effect on subsequent growth in COVID-19 cases and deaths. The effects become significant around 2 weeks ahead and persist through around 8 weeks ahead for cases and around 9 weeks ahead for deaths. The peak impact occurs 4 to 6 weeks ahead for cases and around 8 to 9 weeks ahead for deaths. The effects are largest for mobility measured by time spent away from home and time spent at work, though significant effects also are found for time spent at retail and recreation establishments, at transit stations, at grocery stores and pharmacies, and at parks. Fourth, I find that public health non-pharmaceutical interventions affect future COVID-19 cases and deaths, but that their effects work entirely through, and not independent of, individuals' mobility behavior. Lastly, the dynamic effects of mobility on COVID-19 outcomes are found to be generally similar across counties, though there is evidence of larger effects in counties with high cases per capita and that reduced mobility relatively late.

Covid Economics 35, 7 July 2020: 257-299

1 I thank Regis Barnichon, Karel Mertens, Enrico Moretti, Adam Shapiro and numerous FRBSF colleagues for advice and suggestions. The views expressed in this paper are solely those of the authors and do not necessarily reect the views of the Federal Reserve Bank of San Francisco, or the Board of Governors of the Federal Reserve System.

2 Federal Reserve Bank of San Francisco.

Covid Economics 35, 7 July 2020: 257-299

1 Introduction

Understanding the causal effects of social distancing behavior on subsequent growth in COVID-19 cases and deaths is clearly of first-order public health and economic importance. Yet, to date, there has been surprisingly little research estimating the full dynamic response of COVID-19 outcomes to exogenous movements in mobility.

There are a couple formidable empirical challenges that help explain this paucity of research. First, observed movements in mobility are likely to be endogenous due both to correlation with other factors that could themselves affect COVID-19 spread and to reverse causality, with publicity of current growth in cases and deaths affecting individuals' mobility (voluntarily and/or via mandatory restrictions). Weather is one such factor. There have been numerous studies in recent months on the impact of temperature and other weather variables on COVID-19 spread. The results have been mixed, with some studies finding significant effects and some not. As this paper will show, weather and mobility are very strongly correlated and hence, given the possibility that weather has a direct effect on COVID-19 outcomes, it is important to study the impacts of weather and mobility *jointly*. Second, the potential lags between mobility and COVID-19 outcomes may be quite long, requiring at least several months of post-outbreak data before one can begin to trace out the full impacts of mobility changes. The availability of geographically granular, high-frequency, real-time data along, of course, with the passage of time now open the doors to such research.

This paper is a first attempt at estimating the full dynamic response of COVID-19 outcomes to exogenous movements in mobility. I estimate this impulse response function (IRF) for COVID-19 cases and deaths up to 10 weeks ahead using a panel Local Projections estimator with county-level data. To identify plausibly exogenous movements ("shocks") in mobility, I use standard regression control techniques in a dynamic panel data framework. In particular, when regressing future COVID-19 outcomes on current mobility, I control for lagged mobility, current and lagged COVID-19 outcomes, COVID-19 testing, weather (temperature, precipitation, and snowfall), and high-dimensional fixed effects for counties and for time. Controlling for current and lagged cases and/or deaths, as well as testing, accounts for the likelihood that news of current local COVID-19 spread, which itself would predict future cases and deaths, induces people to increase or decrease their current social distancing (mobility) behavior. Controlling for lagged mobility helps ensure that current movements in mobility are not driven simply by persistence from past mobility shocks. Including county fixed effects effectively controls for many important known and unknown characteristics of local communities that can increase COVID-19 transmission and/or lethality, such as demographics, socioeconomic status, density, etc. Time (week) fixed effects absorb seasonal

Covid Economics 35, 7 July 2020: 257-299

factors, common time trends, and any policies or other factors at the national level. In an extension, I also show the results are robust to controlling for public health non-pharmaceutical interventions (NPIs), such as shelter-in-place orders and school closures.

The analysis reveals a number of important findings. First, weather and mobility are highly correlated and thus omitting either factor when studying the COVID-19 effects of the other is likely to result in substantial omitted variable bias. Second, holding fixed time spent away from home (of one many available mobility measures), temperature is found to have a negative and significant effect on COVID-19 cases for at least 6 weeks ahead and on deaths for at least 10 weeks ahead, while precipitation and snowfall have no consistent significant effects. However, while the temperate effect is robust for subsequent COVID-19 deaths, it is not robust across all measures of mobility for cases and thus the temperature result should be viewed with appropriate caution. Third, controlling for weather, overall mobility, measured by time spent away from home, is found to have a large positive effect on subsequent growth in COVID-19 cases and deaths. The effects become significant around 2 weeks ahead and persist through around 7 weeks for cases and around 8 weeks for deaths. The peak effect occurs around 5 weeks ahead for cases and around 7 weeks for deaths. Looking across subcategories of mobility, the effects are clearest for time spent at workplaces and retail & recreation, though there is also evidence of an adverse effect on deaths growth for mobility at transit stations, at grocery & pharmacy, and, to a lesser extent, at parks. Further, the effects of mobility are quantitatively large. For example, a 1% increase in time spent away from home is found to increase COVID-19 case growth over the following 8 weeks by nearly 6%. Fourth, I find that the impact of public health NPIs works entirely through affecting individual mobility behavior.

Fifth and lastly, I find that the dynamic effects of mobility are generally similar across counties, with a couple of interesting exceptions. In particular, I estimate heterogeneous treatment models allowing for mobility to affect cases or deaths growth differently across counties depending on the timing of their local onset of COVID-19 spread (first case), the timing of their initial plunge in mobility, their total cases per capita, the share of their population in nursing homes, and for mobility increases versus decreases. I find little to no evidence of heterogeneity across these dimensions except for two cases. First, there is some evidence, especially for cases growth, that mobility effects are larger when and where cases per capita are higher. Second, the impact of mobility on deaths from 2 to 5 weeks ahead is larger for counties which experienced their largest weekly drop in mobility at a relatively late date. This finding may suggest that for a given magnitude of mobility change, earlier action is more effective than later action.

Prior studies have tended to focus on the effect of non-pharmaceutical policy interven-

Covid Economics 35, 7 July 2020: 257-299

tions, especially shelter-in-place orders ("SIPOs"), on COVID-19 outcomes, rather than the direct effect of mobility.[1] For instance, Courtemanche, Garuccio, Le, Pinkston, and Yelowitz (2020) finds a sizable impact of SIPOs on subsequent COVID-19 case growth in the U.S.. Similarly, Hsiang, Allen, Annan-Phan, Bell, Bolliger, Chong, Druckenmiller, Huang, Hultgren, Krasovich, et al. (2020) exploit subnational panel data, including state-level data for the U.S., during the early months of the pandemic to estimate the impact of various NPIs on COVID-19 case growth, finding substantial beneficial impacts. These studies do not directly study the role of mobility, but rather assume that the effect of NPIs on COVID-19 cases is mediated through the channel of NPIs, reducing mobility which in turn reduces infections. Similarly, Askitas, Tatsiramos, and Verheyden (2020), which uses cross-country panel data in an event-study framework, finds that NPIs decrease the daily incidence of COVID-19. They also find that NPIs decrease mobility, though they do not assess whether NPIs affect COVID-19 cases independent of mobility or whether mobility directly affects cases.

There are at least four other notable studies of the direct effects of mobility behavior on COVID-19 outcomes. Soucy, Sturrock, Berry, Westwood, Daneman, MacFadden, and Brown (2020) uses variation across 40 global cities for the late March to mid-April time period and find a strong correlation between mobility and COVID-19 case growth 14 days ahead. Badr, Du, Marshall, Dong, Squire, and Gardner (2020) also investigate the correlation between mobility and subsequent case growth, but using panel data from the 25 U.S. counties with the most cases as of mid-April. Estimating this correlation for varying lag lengths, they find it peaks at 11 days. These two studies are correlational. As noted above, the correlation between mobility and COVID-19 outcomes can be driven by omitted variables and/or reverse causality, such that correlations may be a misleading guide to the causal impacts of changes in mobility and thus not useful for policy prescriptions. Kapoor, Rho, Sangha, Sharma, Shenoy, and Xu (2020) is one of the first studies aiming to estimate the causal effect of mobility changes on COVID-19 cases. They investigate the early declines in mobility, prior to local NPIs, using a cross-sectional regression design with U.S. county-level data. They employ an instrumental variables approach, using precipitation as an instrument for the early mobility declines. They find that these declines were associated with fewer cases and deaths up to at least 18 days later, the farthest horizon they investigate. Their results are consistent with those in this paper, though I find effects persisting as far out as 9 weeks ahead. Most recently, Unwin, Mishra, Bradley, Gandy, Vollmer, Mellan, Coupland, Ainslie, Whittaker, Ish-Horowicz, et al. (2020) estimate the impact of mobility on the transmission rate, R_t, of SARS-CoV-2, the virus that causes COVID-19. They estimate this impact with state

[1] See Nguyen, Gupta, Andersen, Bento, Simon, and Wing (2020) for a discussion of this literature as well as original evidence.

Covid Economics 35, 7 July 2020: 257-299

panel data using a logit model with state-specific random effects. Similar to this paper, they measure mobility using the Google Mobility Reports. Because the true transmission rate is unknown, they infer it from observed COVID-19 deaths using a Bayesian semi-structural model. An important element of the model is an assumed lag structure between infection and death. The authors use estimates of that lag structure from early studies of the disease in China.

As mentioned above, the evidence on weather's impact of COVID-19 is mixed. Xu, Rahmandad, Gupta, DiGennaro, Ghaffarzadegan, Amini, and Jalali (2020) finds a "modest" negative effect of temperature on covid19 case growth globally. They assume a 10-day lag, which the evidence in this paper suggests may be too short to see the full impact. Carleton, Cornetet, Huybers, Meng, and Proctor (2020) finds a negative relationship between UV light and COVID-19 case growth, while they find "weak or inconsistent lagged effects of local temperature, specific humidity, and precipitation." Similarly, Jamil, Alam, Gojobori, and Duarte (2020) find no evidence of a link between temperature and COVID-19 case growth across countries and Chinese provinces. Yet, these papers do not generally account for the fact (documented below) that weather has very strong effects on mobility behavior, which itself affects COVID-19 outcomes, making it hard to know how much their results reflect a serious omitted variable bias.

Most of the studies above and others in the nascent literature on COVID-19 impose a specific response lag between mobility or NPI/SIPOs and COVID-19 cases. Typically, they assume lags of around 14 days, which is based on the logic of a roughly 7-day incubation period (from exposure to symptoms) and a 7-day "confirmation" period from symptom onset to a positive test result. Such an assumption may miss an important third phase, the transmission propagation phase. After that initial exposed individual becomes infectious, they may spread the infection to one or more additional individuals, who may in turn spread it to others, and so on. Each of these "rounds" of infection transmission involves its own incubation and confirmation delays, potentially spreading out over time the effect of any initial shock such as mobility, weather, NPIs, etc.. This highlights the usefulness of applying an estimation method, such as local projections, that does not impose any assumed response lag and instead allows for estimation of the full impulse response function.

Covid Economics 35, 7 July 2020: 257-299

2 Data and Stylized Facts

2.1 COVID-19 Data

Daily county-level on COVID-19 cases and deaths were obtained from usafacts.org, which compiled the data from state public health agencies.[2] I also obtained alternative data on cases and deaths from the New York Times database (via tracktherecovery.org). The correlation between these two data sets is extremely high (0.9995) and the results in the paper are nearly identical using the New York Times data. Data on daily testing, which are only available at the state level, come from The COVID Tracking Project and were downloaded from tracktherecovery.org.

The national time series for COVID-19 cases and deaths are shown in Figure 1. The earliest cases occurred in late January, but nationally cases began accelerating rapidly in mid-March, with deaths accelerating around the beginning of April.[3] Yet the upsurge in cases and deaths varied substantially across counties, as shown in the histogram in Figure 2. Of the 3,016 counties with available data on cases, about one-third had their first case in or before the third week of March (15th - 21st). The week with the most first cases, in 822 counties, was the following one, March 22-28. There is a long right tail thereafter. In fact, there remain a few dozen sparsely populated counties that have no cases as of June 28, the latest date of data as of this writing.

2.2 Mobility Data

Google Mobility Reports provide percent changes in mobility relative to the Jan. 3 – Feb. 6 average. More specifically, Google describes the data as follows: "These datasets show how visits and length of stay at different places change compared to a baseline. We calculate these changes using the same kind of aggregated and anonymized data used to show popular times for places in Google Maps. Changes for each day are compared to a baseline value for that day of the week: The baseline is the median value, for the corresponding day of the week, during the 5-week period Jan. 3 – Feb. 6, 2020."[4]

The Dallas Fed Mobility and Engagement Index (MEI) is the first principal component of seven cell-phone-based geolocation variables from the data provider Safegraph. The seven

[2] `https://usafactsstatic.blob.core.windows.net/public/data/covid-19/covid_confirmed_usafacts.csv`

[3] There are two larges spikes in daily death counts in Figure 1. These are due to the addition by New York City and New Jersey of probable COVID-19 deaths based on retrospective reviews of deaths certificates. New York City added around 3,800 deaths on April 14 and New Jersey added 1,854 on June 27. I remove these added deaths from the data for the regression analyses in the paper. Their inclusion would yield large outliers in death growth rates which could unduly bias the results.

[4] The data were accessed at `https://www.google.com/covid19/mobility/`.

Covid Economics 35, 7 July 2020: 257-299

Figure 1: Daily COVID-19 Cases and Deaths in the United States

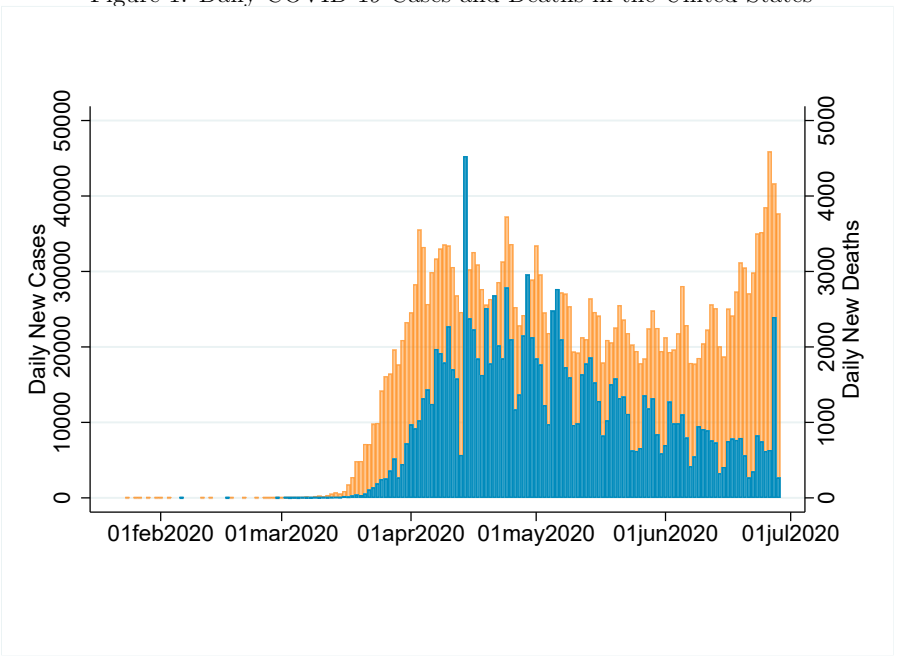

Figure 2: Timing of Onset of Local COVID-19 Outbreaks

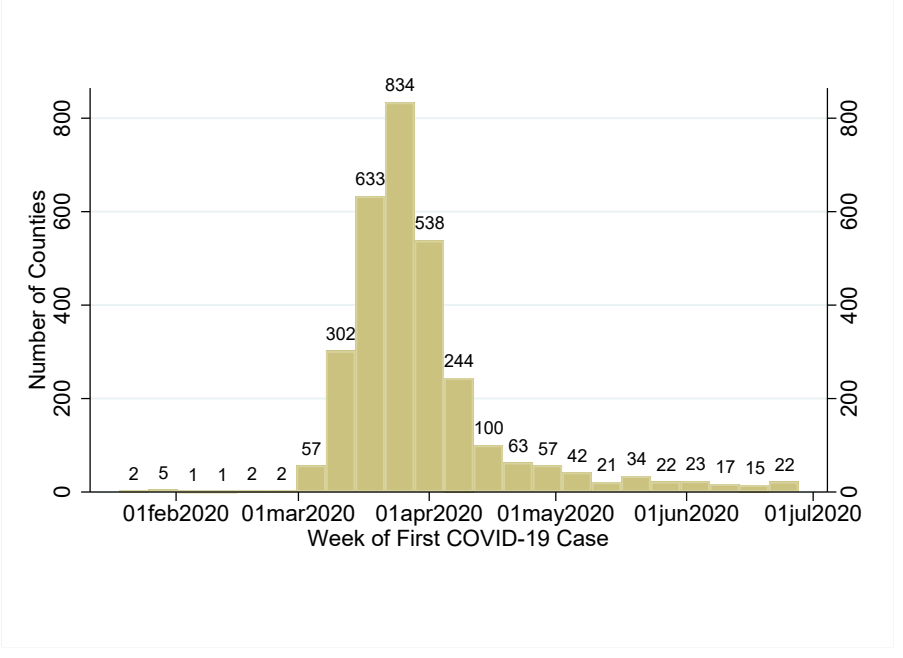

variables are as follows: the fraction of devices leaving home in a day, the fraction of devices away from home for 3-6 hours at a fixed location, the fraction of devices away from home longer than 6 hours at a fixed location, an adjusted average of daytime hours spent at home, the fraction of devices taking trips longer than 16 kilometers, the fraction of devices taking trips less than 2 kilometers, and the average time spent at locations far from home. Each variable is scaled by the weekday-specific average over January and February prior to the principal component analysis. See Atkinson, Dolmas, Koch, Koenig, Mertens, Murphy, and Yi (2020) for details.[5]

The aggregate (population-weighted sums over counties) time series movements of these mobility variables are shown in Figures 3 and 4. The graphs document a steep plunge in mobility nationally over the second half of March. Most measures bottomed out in early April and have shown gradual recovery since then. As of late June, most are at or near their January – February baseline. Notice that mobility to parks is well above that baseline.

The high level of mobility at parks illustrates an important point to keep in mind when interpreting movements in these measures of mobility: They are not seasonally adjusted. In particular, note that the Google Mobility Reports are % changes in mobility relative to a Jan. 3 – Feb. 6 baseline. January is generally the lowest month of employment in non-seasonally adjusted data on payroll employment and retail spending, due to typical layoffs and reduced spending after the December holidays. July is the second lowest month, due largely to school employees taking off the summer. Thus, one might expect some upward pressure on mobility, especially employment- and retail-related mobility measures, in the months after January and some downward pressure in July and August, though it is difficult to know how large these seasonal factors are given that the mobility data are unavailable for previous years.

In contrast to the heterogeneity across counties in timing of first COVID-19 cases (Figure 2), the plunge in mobility was remarkably synchronous across the country. Figure 5 provides a histogram showing, for each week from January to the latest week available, how many counties saw their largest weekly decline in mobility in that week. More than two-thirds of the roughly 3,000 counties with mobility data (using the MEI) saw their largest drop in mobility in the second half of March.

2.3 Weather Data

Following Wilson (2019), I construct measures of daily weather at the county level from the Global Historical Climatology Network Daily (GHCN-Daily) data set. The GHCN-

[5]The data were accessed at `https://www.dallasfed.org/~/media/documents/research/mei/MEI_counties_scaled.csv`.

Covid Economics 35, 7 July 2020: 257-299

Covid Economics 35, 7 July 2020: 257-299

Figure 3: Google Mobility Over Time

(a) At Residence (Inverted)

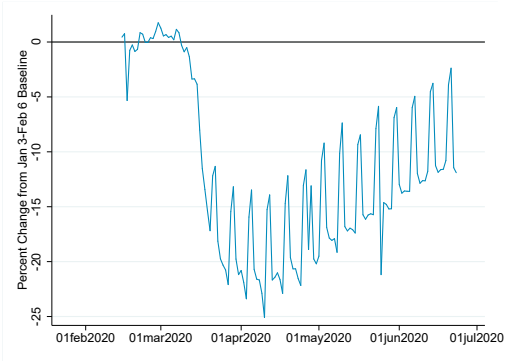

(b) At Workplace

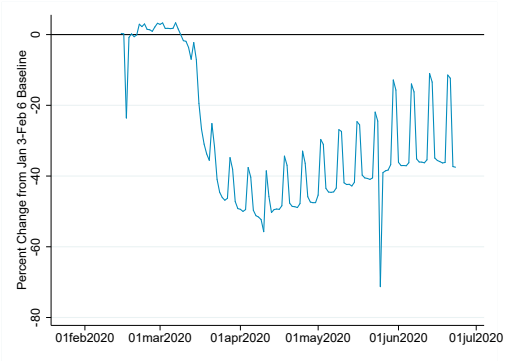

(c) At Transit Stations

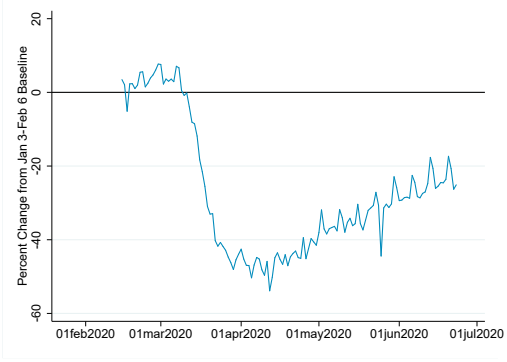

(d) At Parks

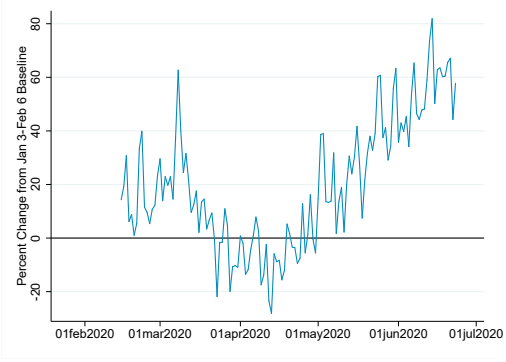

(e) At Retail and Recreation

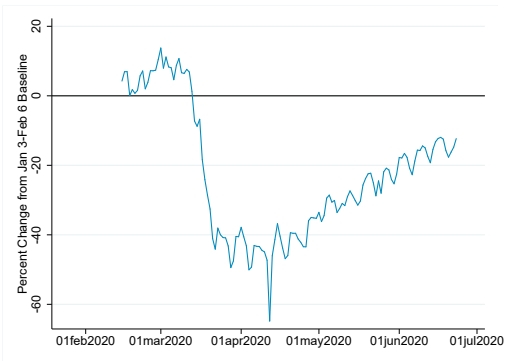

(f) At Grocery and Pharmacy

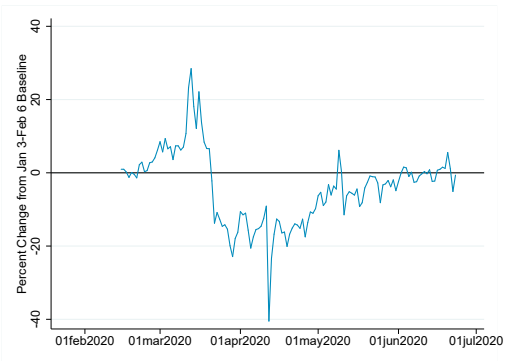

Covid Economics 35, 7 July 2020: 257-299

Figure 4: Dallas Fed Mobility & Engagement Index (MEI) Over Time

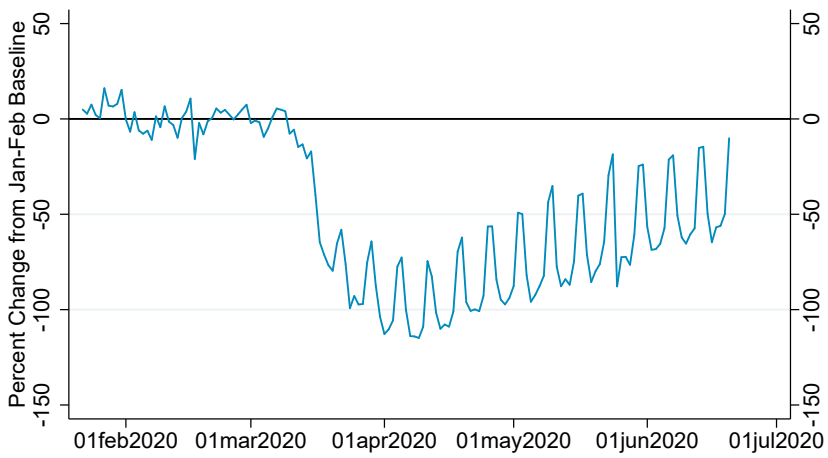

Figure 5: Timing of Local Mobility Plunge

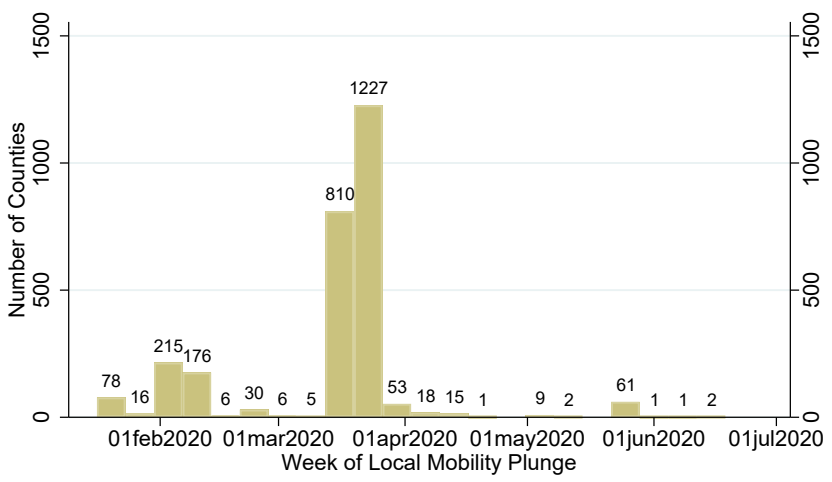

Source: Author's calculations based on Dallas Fed Mobility & Engagement Index

Covid Economics 35, 7 July 2020: 257-299

Daily is provided by the U.S. National Climatic Data Center (part of the National Oceanic and Atmospheric Administration (NOAA)) and contains daily weather measurements from a little over 4,700 weather stations throughout the United States, though not all stations provide readings every day. All stations with data on a given date are used for measuring county weather on that date. The spatial distribution of weather stations is highly correlated with the spatial distribution of population.

The readings from individual weather stations are used to estimate county-level weather using an inverse-distance weighting procedure. First, the surface of the conterminous United States is divided into a 5-mile by 5-mile grid. Second, weather values for each grid point are estimated using inverse-distance-weighted averages of the weather values from weather stations within 50 miles of the grid point. See Wilson (2019) for further details of this procedure.

This procedure yields the following daily county-level weather variables: maximum temperature (degrees Fahrenheit), minimum temperature, precipitation (mm), and snowfall (cm).

2.4 Other Data

I obtained data on start- and end-dates of various local non-pharmaceutical interventions (NPIs) from the Keystone *Coronavirus City And County Non-Pharmaceutical Intervention Rollout Date Dataset*.[6] Keystone has compiled data for all states and about 600 counties, including all counties that had at least 100 cases as of April 6, 2020. I use the county level data for those ≈600 counties and state level data for other counties. The data include start- and end-dates for the following 10 NPIs: social distancing ("social distancing mandate of at least 6 feet between people"), shelter in place ("an order indicating that people should shelter in their homes except for essential reasons"), prohibition of gatherings above 100, prohibition of gatherings above 25, prohibition of gatherings above 10, prohibitions of gatherings of any size, closing of public venues ("a government order closing gathering venues for in-person service, such as restaurants, bars, and theaters"), closure of schools and universities, closure of non-essential services and shops, closure of religious gatherings, and full lockdown. I condense these data into a single variable, the number of NPIs in place in the county on each date.

Data on population (average daily residents) in nursing homes by county as of June 4, 2020 were obtained from the Centers for Medicaid and Medicare at: `https://data.medicare.gov/Nursing-Home-Compare/Provider-Info/4pq5-n9py`.

[6] `https://raw.githubusercontent.com/Keystone-Strategy/covid19-intervention-data/master/complete_npis_inherited_policies.csv`

3 Methodology

My primary objective is to estimate the causal effects of mobility, which can be influenced by government policy choices and public opinion, on COVID-19 spread. Simultaneously estimating the causal effects of weather, which is exogenous, on COVID-19 is both a secondary objective and necessary for obtaining an unbiased estimate of mobility's effect given that weather and mobility are likely correlated. Below I describe the econometric methodology used to achieve these objectives.

3.1 Contemporaneous Effect of Weather on Mobility

I first estimate the contemporaneous effect of weather ($\mathbf{w}$) on mobility (m) using a daily county panel data model:

$$m_{it} = \mathbf{w_{it}}\gamma + \mathbf{x_{i,t}}\phi + \alpha_i + \alpha_t + \varepsilon_{it} \tag{1}$$

where i indexes counties and t indexes dates. $\mathbf{w_{it}}$ is a vector of weather variables, consisting of daily maximum temperature, precipitation, and snowfall. $\mathbf{x_{i,t}}$ is a vector of control variables, consisting of COVID-19 case growth (daily new cases divided by total cases) and testing growth. α_i and α_t are fixed effects for county and date. The rationale for the inclusion of these control variables and fixed effects is discussed in the *Identification* subsection (3.3) below.

3.2 Dynamic Effects of Mobility and Weather on COVID-19

Next, I jointly estimate the effects of mobility (m) and weather ($\mathbf{w}$) on subsequent COVID-19 growth (g) using a panel Local Projections estimator:

$$g_{i,t+h} = \sum_{\tau=0}^{1} \psi^{h,\tau} g_{i,t-\tau} + \sum_{\tau=0}^{1} \beta^{h,\tau} m_{i,t-\tau} + \mathbf{w_{it}}\delta^h + \mathbf{x_{i,t}}\phi^h + \alpha_i^h + \alpha_t^h + \epsilon_{i,t,t+h} \tag{2}$$

where i indexes counties, t indexes time (days for daily regressions and weeks for weekly regressions), and $\mathbf{x_{i,t}}$ is a vector of control variables.

I consider two different COVID-19 outcomes (g): growth in cases and growth in deaths. Growth during a given time period ($t+h$) is defined as the number of new cases (deaths) recorded during that period divided by the total number of cases (deaths) as of time t. For instance, case growth during period $t+h$ is calculated as $g_{i,t+h} = (cases_{i,t+h} - cases_{i,t+h-1})/cases_{i,t}$. Note that this variable is the *flow* of new cases (deaths) during $t+h$.

Covid Economics 35, 7 July 2020: 257-299

Covid Economics 35, 7 July 2020: 257-299

To additionally estimate effects on *cumulative* growth of cases or deaths over a given horizon, I estimate versions of equation 2 in which the dependent variable is cumulative growth (G), defined as the cumulative sum of new cases from t to $t+h$ divided by the total number of cases (deaths) as of time t: $G_{i,t+h} = (cases_{i,t+h} - cases_{i,t})/cases_{i,t}$, and similarly for deaths.

The local projections method (Jordà (2005)) traces out an impulse response function (IRF) by estimating equation 2 sequentially over horizons from $h = 1$ to some maximum horizon, H. I estimate IRFs out to $H = 10$ weeks ahead. The IRF for mobility is traced out by the sequence of $\beta^{h,0}$, while the IRF for any element of the weather vector $\mathbf{w_{it}}$ is traced out by the sequence of its element of the coefficient vector δ^h.

I estimate IRFs at a weekly frequency, using weekly aggregated data, for three reasons. First, aggregating the daily data to weekly frequency removes the sizable variation between weekdays and weekends in the time series of mobility as well as COVID-19 cases and deaths (apparent in Figures 1, 3, and 4)). Second, there likely is considerable measurement error at the daily frequency in COVID-19 cases, deaths, and testing due to reporting lags. That measurement error should largely cancel out with aggregation to the weekly level. Third, estimating equation 2 is computationally intensive and hence estimating it for daily horizons from $h = 7$ to 70 would be extremely time intensive. Nonetheless, for robustness, I also produce IRF results at the daily frequency for horizons 7, 14,...,70. The results are very similar to the weekly results, though less precisely estimated.

3.3 Identification

The causal effect of mobility, $\beta^{h,0}$, or weather, δ^h, on COVID-19 cases and deaths is unlikely to be identified by any simple cross-sectional correlations due to a variety of omitted variable and reverse causality concerns. I address these identification concerns through dynamics, control variables, and fixed effects. In terms of dynamics, equation 2 uses leads of the COVID-19 outcomes as dependent variables to mitigate the potential contemporaneous reverse causality due to local news about current cases or deaths inducing people to increase or decrease their social distancing (mobility) behavior. For the same reason, I include both contemporaneous and a one-week lag of the dependent variable as well as current growth in testing. Controlling for current testing helps mitigate concerns that public fears (or lack thereof) about community spread may affect mobility and also be correlated with future cases and deaths growth because such fears should be reflected in greater demand for tests.

I also include current cases per capita in the set of controls. Cases per capita should be roughly proportional to the share of the population no longer (or less likely to be) susceptible

to infection, which is expected to reduce the capacity for future growth in cases (and in turn deaths).[7] When the dependent variable is growth in deaths, I add current case growth as a regressor. I additionally include a one-week lag of mobility so that $\beta^{h,0}$ can be interpreted as the effect of a current "shock" or change in mobility that was not driven simply by its autoregressive properties.

As the results below will demonstrate, when studying the effect of mobility on COVID-19, it is important to control for weather. Likewise, when studying the effect of weather on COVID-19, it is important to control for mobility. As discussed in the introduction, most COVID-19 studies to date of either weather or mobility have not controlled for the other factor and could be subject to serious omitted variable bias. Thus, I include both mobility and weather variables – maximum daily temperature, precipitation, and snowfall – in all regressions unless otherwise indicated.

The county fixed effects absorb many important known and unknown characteristics of local communities that can increase COVID-19 transmission and/or lethality. These time-invariant characteristics include demographics (such as age, gender, and race), socioeconomic status, access to healthcare, population density, the presence of nursing homes or meat-packing plants, and openness to international travelers. Desmet and Wacziarg (2020) document the importance of many such time-invariant factors for COVID-19 cases and deaths among U.S. counties. In Section 5.2, I investigate whether the average "treatment" effect of mobility on COVID-19 varies across some of these county characteristics.

The time fixed effects are also crucially important. In particular, they will absorb seasonal factors and any common (i.e., national) time trends. This is particularly important given that weather, especially temperature and snowfall, obviously has strong trends over the January to June sample period, and mobility has also trended higher from late March onward.

Covid Economics 35, 7 July 2020: 257-299

3.4 Inference

The standard errors and confidence intervals reported in the paper are robust to heteroskedasticity and two-way clustering by county and state-by-time (where time is date for daily regressions and week for weekly regressions). The clustering by county allows for the possibility that errors, $\varepsilon_{i,t}$ in equation 1 and $\epsilon_{i,t,t+h}$ in equation 2, are serially correlated. The clustering by state-time allows for the possiblity that errors are contemporaneously correlated across counties within the same state. This clustering will account for cross-county correlation stemming from unobserved statewide factors such as state economic and public

[7] The share of the population no longer susceptible to infection is a key variable in epidemiological SIR models. The level of this share is subject to considerable debate among epidemiologists, but most agree that it is many (perhaps 10) times larger than known cases per capita.

health policies. It will also account for geo-spatial correlation in measurement error, for example in weather data, to the extent such correlation is encompassed by state boundaries.

3.5 Data Sample

All regressions in the paper use the maximum data sample available for the variables used in that regression, with one restriction. I restrict the sample time period for each county to begin with the first date on which cases per capita exceeded one per 10,000 persons. This restriction excludes observations from time and places where the COVID-19 outbreak had not yet begun. The sample time period varies across regressions depending on the horizon (in the local projections regressions) and on the availability of the mobility variable used. For the local projections regressions, the further out the horizon, the fewer the time periods (t) available for estimation.

Data availability varies across the mobility variables, with the MEI data beginning in January and the Google Mobility data beginning in mid-February. All variables are available through late June as of the time of this writing. None of the mobility variables is available for all counties due to suppression of data (by Safegraph and Google) for counties with fewer mobile devices to mitigate privacy concerns. The county coverage varies across the Google mobility measures from about 1,000 counties for time spent at parks to roughly 2,800 for time spent at work. The MEI data covers about 3,000 counties. (There are 3,140 counties in the U.S..)

4 Main Results

4.1 Effects of Weather on Mobility

Before presenting the formal regression results, I begin with some non-parametric graphical evidence on the contemporaneous daily relationship between temperature and mobility. Temperature is the maximum daily high measured in degrees Fahrenheit. The panels in Figure 6 show bin-scatter plots with temperature on the x-axis and mobility on the y-axis. A bin-scatter plot is a common way to visualize correlations involving a large number of observations. Each variable is first residualized by regressing on county and date fixed effects. The x and y variables are then averaged within 100 bins corresponding to each percentile of the distribution of temperature values. (Hence, bins will be many degrees wide toward the lower and upper end of the temperature range and somewhat narrower than one degree in the middle of the range.) In each plot, a quadratic fit line (based on the raw, not binned, data) is added. These plots show a strong positive relationship between temperature and mobility,

Covid Economics 35, 7 July 2020: 257-299

Covid Economics 35, 7 July 2020: 257-299

conditional on county and time fixed effects.[8] The relationship generally is approximately linear. The tightness of the fit varies across mobility measures. It is tightest for time spent at home, at parks, at retail & recreation, and at grocery & pharmacy. It is somewhat weaker for time spent at work and at transit stations.

Table 1 displays the results from estimating equation 1 using alternative measures of mobility. Each column corresponds to a separate regression and the column heading indicates which mobility measure is used as the dependent variable. The mobility measures are the six Google Mobility Report variables ("At...") and the Dallas Fed Mobility & Engagement Index (MEI). I multiply the time spent at home measure by -1 to facilitate comparison with the other measures. It should thus be interpreted as time spent *away* from home. The sample time period and number of counties are indicated at the bottom of the table and vary primarily depending on the availability of the mobility data, as noted in the previous section. Recall that county*date observations in which cases did not yet exceed one per 10,000 population are excluded, thus each regression uses an unbalanced panel. The weather variables are interacted with a weekday vs weekend indicator to allow for the possiblity that weather affects mobility differently on weekdays, when most workers work, than on weekends, when most workers have off.

The results show a number of clear patterns. First, weekday temperatures have a strong positive effect on all measures of mobility. Not surprisingly, the effect is strongest, by an order of magnitude, for time spent at parks. The coefficient of **1.2** (standard error of **0.08**) for parks implies that each degree of higher temperature on a weekday increases mobility at parks by **1.2** percentage points. (Recall that the Google mobility variables are measured in percentage change relative to Jan. 3 – Feb. 6.) The weekend temperature effect for parks is smaller, at **0.7**, but still very statistically and economically significant. Weekday temperature also has a positive and strongly statistically significant effect on the other mobility measures, though it is quantitatively smaller. The coefficients vary from about **0.06** (for At Work) to **0.16** (for At Grocery & Pharmacy). Second, weekend temperatures have a positive effect on some mobility measures (parks, retail & recreation, grocery & pharmacy) but a negative effect on others (work, transit, home, and MEI). The negative effect, which is especially strong for At Work, may reflect that some workers with flexibility regarding weekend work may opt to choose leisure over work on weekends with pleasant weather (and/or on hot days when outside work is less pleasant). Third, precipitation has a strong negative effect on mobility. This is true for all measures of mobility, though as with temperature, the effect is largest for time spent at parks. The negative effect is true for precipitation on both

[8]Results for MEI are shown in Appendix Figure A1. The relationship is also strongly positive on average though it peaks around 70° F and turns negative above that.

Covid Economics 35, 7 July 2020: 257-299

Figure 6: Relationship Between Temperature and Mobility

(a) At Residence (Inverted)

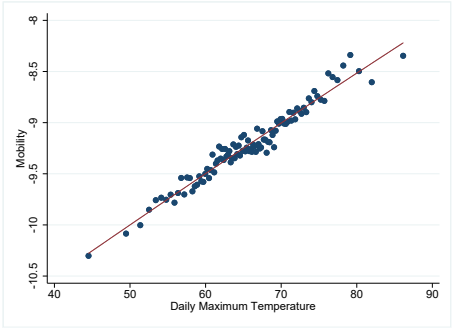

(b) At Workplace

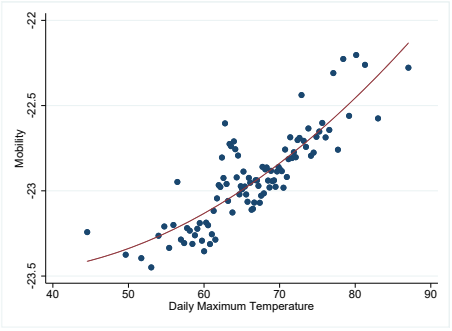

(c) At Transit Stations

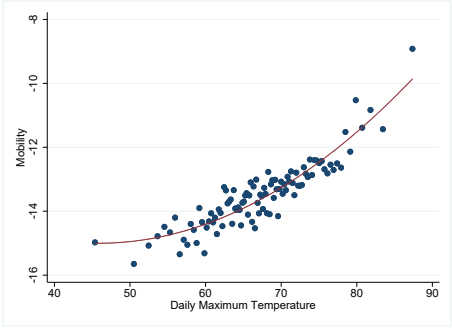

(d) At Parks

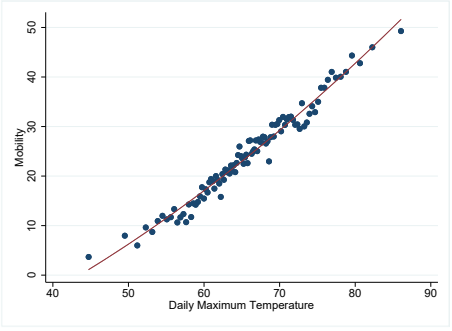

(e) At Retail and Recreation

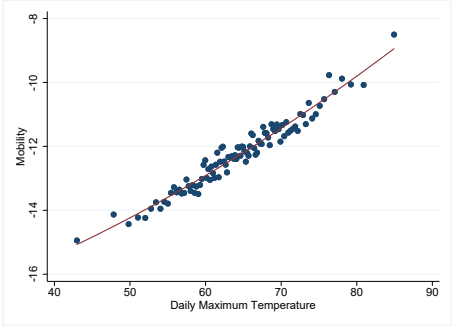

(f) At Grocery and Pharmacy

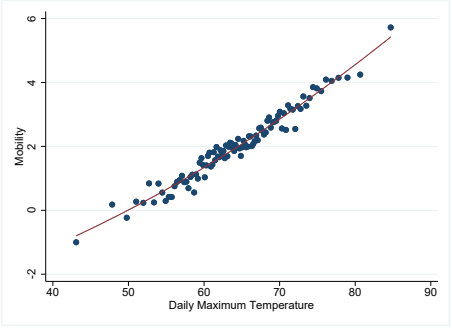

Note: Bin scatterplots, using 100 bins after residualizing on county and date fixed effects. Solid line is quadratic fit line on the full (non binned) sample of residuals.

Table 1: Effect of Weather on Mobility

	(1) At Home (Inverted)	(2) At Work	(3) At Transit	(4) At Parks	(5) At Retail & Rec	(6) At Grocery & Pharmacy	(7) MEI
Max Daily Temp – Weekday	0.0669***	0.0564***	0.127***	1.207***	0.148***	0.156***	0.102***
	(0.00322)	(0.00561)	(0.0136)	(0.0779)	(0.0130)	(0.0112)	(0.0149)
Max Daily Temp – Weekend	-0.0190***	-0.0514***	-0.0347*	0.700***	0.0889***	0.133***	-0.0515**
	(0.00617)	(0.0104)	(0.0205)	(0.0967)	(0.0176)	(0.0179)	(0.0241)
Precipitation – Weekday	-0.0193***	-0.00574	-0.0438***	-0.445***	-0.0273***	-0.0326***	-0.101***
	(0.00298)	(0.00370)	(0.00782)	(0.0626)	(0.00942)	(0.00832)	(0.0118)
Precipitation – Weekend	-0.0491***	-0.0485***	-0.140***	-0.834***	-0.0839***	-0.0946***	-0.173***
	(0.00704)	(0.0103)	(0.0212)	(0.130)	(0.0196)	(0.0204)	(0.0266)
Snowfall – Weekday	-0.0827**	-0.107**	-0.299***	-1.288***	-0.338***	-0.350***	-0.0936
	(0.0352)	(0.0418)	(0.103)	(0.296)	(0.0908)	(0.0884)	(0.127)
Snowfall – Weekend	-0.329***	-0.579***	-1.137***	-3.201***	-0.715***	-0.573**	-1.210***
	(0.0708)	(0.134)	(0.253)	(0.927)	(0.187)	(0.239)	(0.250)
Cases Growth	0.000500	0.00245***	0.000112	-0.0175*	-0.00504***	0.00411**	-0.00328*
	(0.000445)	(0.000635)	(0.00241)	(0.0105)	(0.00159)	(0.00165)	(0.00177)
Testing Growth	-0.0303*	-0.0204	0.000301	-0.398	-0.182**	0.0371	-0.578***
	(0.0156)	(0.0286)	(0.0614)	(0.293)	(0.0796)	(0.0642)	(0.0855)
Observations	96841	197064	82194	51545	128825	118581	214134
Adjusted R^2	0.919	0.884	0.841	0.656	0.843	0.756	0.874
Earliest date	29feb2020	29feb2020	29feb2020	29feb2020	29feb2020	29feb2020	29feb2020
Latest date	23jun2020	23jun2020	23jun2020	23jun2020	23jun2020	23jun2020	20jun2020
# of days	116	116	116	116	116	116	113
# of counties	1308	2532	1060	854	2083	2006	2705

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Covid Economics 35, 7 July 2020: 257-299

weekdays and weekends, though it larger for weekends. This likely reflects that common weekend activities, like going to parks, going out to eat, going retail shopping, and going grocery shopping are less appealing when it is raining. Finally, snowfall, like rain, also has a strong negative effect on mobility, with larger effects on weekends. Again, the effect is largest for time spent at parks. Each centimeter of snowfall reduces mobility at parks by **3.2** percentage points.

Overall, Table 1 makes clear that weather and mobility are highly correlated and thus omitting either factor when studying the COVID-19 effects of the other is likely to result in substantial omitted variable bias.

4.2 Dynamic Impacts of Weather on COVID-19

I now turn to jointly estimating the impulse response functions (IRFs) of mobility and weather on COVID-19 cases and deaths using the local projections estimator described above (equation 2). As mentioned earlier, I estimate the IRFs at the weekly frequency on weekly-aggregated data in order to smooth over the day-of-week variability and high-frequency measurement error and also to reduce computational burden. The dependent variables are h-weeks ahead growth in cases and deaths, for $h = 1$ to 10 weeks. Recall that growth is defined as new cases (deaths) during week $t + h$ relative to total cases (deaths) as of week t.

Covid Economics 35, 7 July 2020: 257-299

I begin by discussing the estimated IRFs with respect to weather shocks. The results for case growth are presented in Figure 7. The point estimates are shown with the circles, while the inner and outer brackets display the 90 and 95% confidence intervals, respectively. The regressions underlying the panels on the left (a, c, and e) omit mobility, while those on the right (b, d, and f) include the mobility away from home measure. Comparing panels (a) and (b), which display the estimated IRF of temperature on subsequent COVID-19 case growth, clearly demonstrates the bias that occurs from omitting mobility. When mobility is omitted, temperature appears to have a positive and increasing effect on COVID-19 cases. When mobility is held fixed, temperature is found to have a significant negative effect on COVID-19 cases 3 to 6 weeks ahead.

The peak coefficient on temperature is approximately **-0.026** at 4 weeks ahead, with a 95% confidence interval of **-0.014** to **-0.038**. This coefficient implies that a one degree warmer temperature during a week lowers growth in COVID-19 cases 4 weeks later by **0.026** percentage points. Evaluated at mean COVID-19 case growth (**0.67**) and mean daily high temperature ($\approx$ **62**°), this implies an elasticity of about **2.4**. In other words, a 10% increase in daily high temperature for a week from **62.0**° to **68.2**° would predict a **24**% decrease in case growth four weeks later from **0.67**% to **0.51**% (holding mobility and the other regressors fixed).

Precipitation and snowfall are found to have significant effects on case growth at some horizons when mobility is omitted. Yet, the effects become statistically insignificant (or imprecise and weakly significant in the case of snowfall 10 weeks ahead) once mobility is controlled for.

The analogous weather results for growth in COVID-19 deaths are shown in Figure 8. For deaths, the results for weather are less dependent on controlling for mobility, though one can still see evidence of positive bias in the effects of temperature when mobility is omitted. When controlling for mobility, temperature has a negative and statistically significant effect on growth in deaths starting two weeks ahead and continuing for at least 10 weeks. Thus, the beneficial effect of temperature is much longer lasting for COVID-19 deaths growth than for cases growth. This result is likely explained by the fact that COVID-19 deaths have been shown to lag cases by one to four weeks. I find no effect of precipitation or snowfall on deaths growth at any horizon.

In sum, holding mobility fixed, temperature is found to have a negative and significant direct effect on COVID-19 cases for up to 6 weeks ahead and on deaths for at least 10 weeks ahead, while precipitation and snowfall have no consistent significant effects. It should be noted, however, that these results are based on a single measure of mobility, time spent away from home. As discussed in the next subsection, the negative and statistically significant

Covid Economics 35, 7 July 2020: 257-299

Figure 7: Dynamic Impacts of Weather on COVID-19 Case Growth – Weekly

Impulse Response Functions Estimated by Panel Linear Projections

(a) Max Temp - Omit Mobility

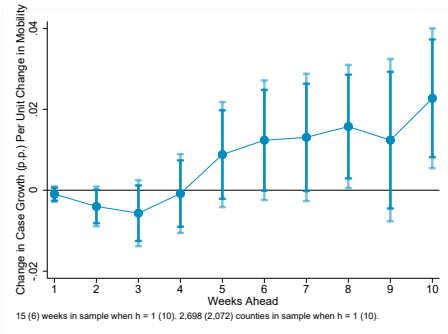

(b) Max Temp - Control for Mobility

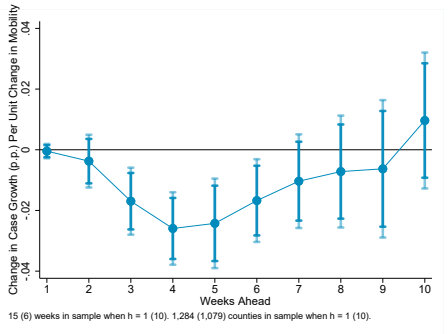

(c) Precip - Omit Mobility

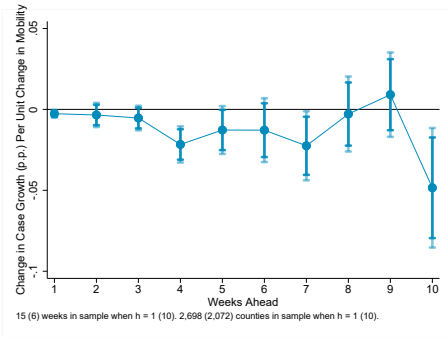

(d) Precip - Control for Mobility

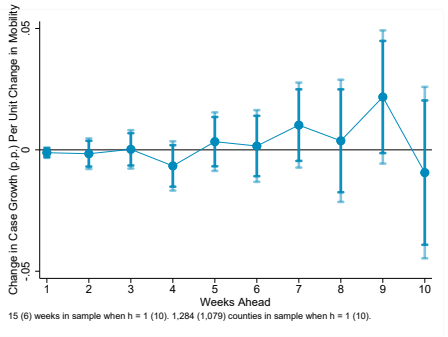

(e) Snowfall - Omit Mobility

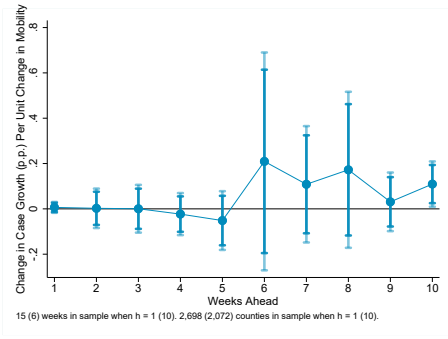

(f) Snowfall - Control for Mobility

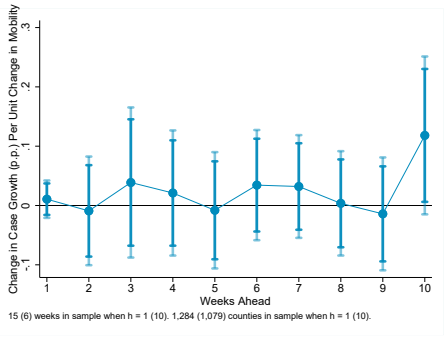

Note: Estimates of equation 2 in the text using panel local projections regressions. Shaded regions are 90% and 95% confidence intervals.

Covid Economics 35, 7 July 2020: 257-299

Figure 8: Dynamic Impacts of Weather on COVID-19 Deaths Growth – Weekly

Impulse Response Functions Estimated by Panel Linear Projections

(a) Max Temp - Omit Mobility

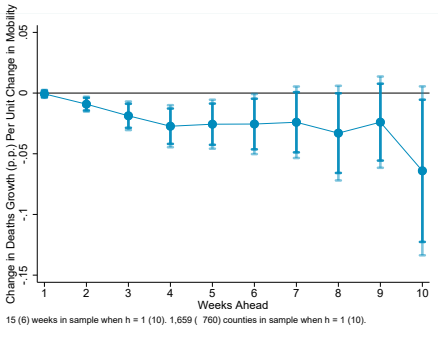

(b) Max Temp - Control for Mobility

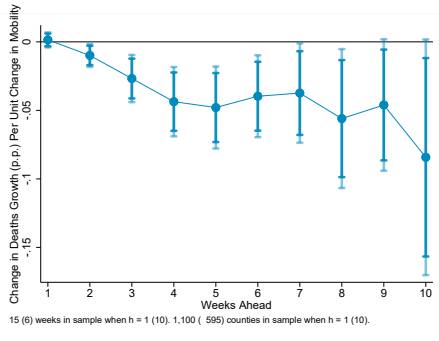

(c) Precip - Omit Mobility

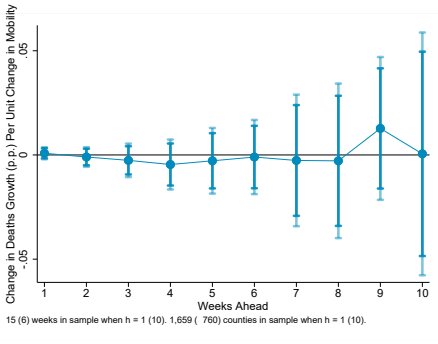

(d) Precip - Control for Mobility

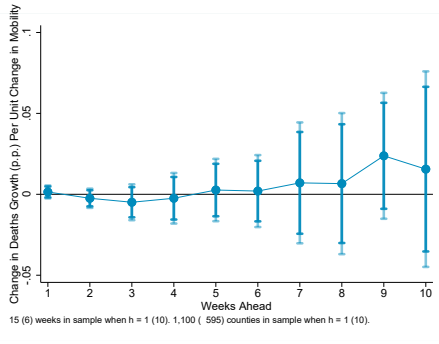

(e) Snowfall - Omit Mobility

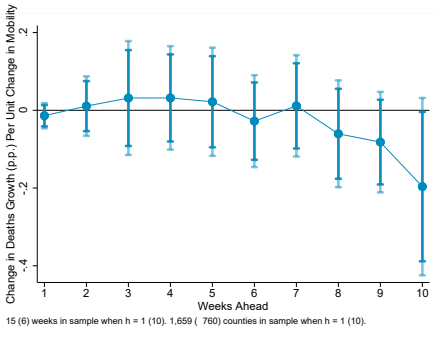

(f) Snowfall - Control for Mobility

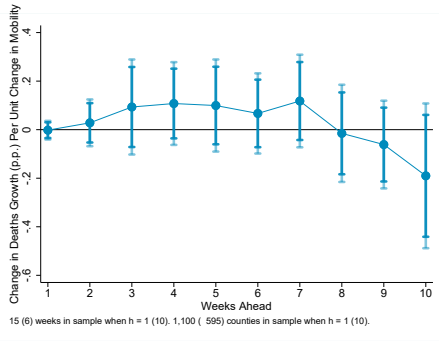

Note: Estimates of equation 2 in the text using panel local projections regressions. Shaded regions are 90% and 95% confidence intervals.

effect of temperature is not robust across all measures of mobility (at least for case growth) and thus should be viewed with some caution.

4.3 Dynamic Impacts of Mobility on COVID-19

I now arrive at the central results of the paper, the estimated impulse response functions for COVID-19 cases and deaths with respect to mobility shocks. These IRFs are the sequence of $\hat{\beta}^{h,0}$ from estimating equation 2 by OLS for each weekly horizon from $h = 1$ to 10 weeks ahead. The IRFs are estimated separately for each of the alternative measures of mobility. The results for case growth are plotted in Figure 9. They reveal two general findings. First, overall mobility, as measured by time spent away from home, has a large positive and significant effect on case growth from approximately 2 weeks ahead to 8 weeks ahead. The peak effect occurs 4 to 6 weeks ahead, with a coefficient of about **0.25**. That magnitude implies that a 1 percentage point increase (decrease) in mobility results in a **0.25** percentage point increase (decrease) in case growth. Evaluated at sample means, this effect implies an elasticity of **5.11**, meaning that a 10% increase in mobility raises cases growth 4 weeks ahead by **51**%.

Second, the effect of mobility on case growth varies considerably across specific types of mobility. The effects are large and statistically significant for time spent at workplaces, at retail & recreation, and at parks. The effect for time spent at parks is particularly notable given the concerns raised recently by public officials about potential surges in COVID-19 stemming from crowded parks and beaches around Memorial Day in many parts of the U.S.. The effect of time spent at transit stations and at grocery & pharmacy is small and statistically insignificant.

The IRFs of mobility shocks for growth in COVID-19 deaths are shown in Figure 10. As with cases, overall mobility, measured by time spent away from home, has a positive, statistically significant, and transitory effect on growth in deaths (see panel (a)). Not surprisingly, the peak effect of mobility on deaths occurs later than the peak effect on cases. For time spent away from home, the peak on deaths occurs 10 weeks ahead, though it is not statistically significant after 9 weeks. Recall that for cases the peak was found to be between 4 and 6 weeks. As with cases, I also find significant positive effects on deaths of mobility for time spent at workplaces, at retail & recreation, and at parks (though it is not statistically significant at all horizons) – again with peak effects later than for cases. However, in contrast to the results for cases, I also find significant positive effects on deaths for transit stations and for pharmacy & grocery.

The full regression results underlying these IRFs are provided, for a single selected hori-

Covid Economics 35, 7 July 2020: 257-299

Figure 9: Dynamic Impacts of Mobility on COVID-19 Case Growth – Weekly

Impulse Response Functions Estimated by Panel Linear Projections

Mobility Measured by Google Mobility Reports

(a) At Residence (Inverted)

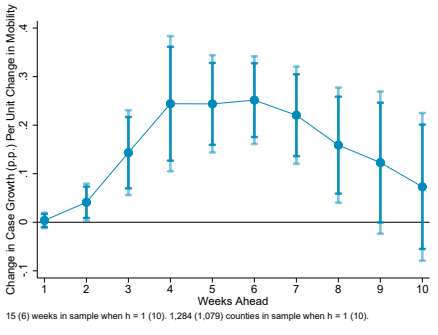

(b) At Work

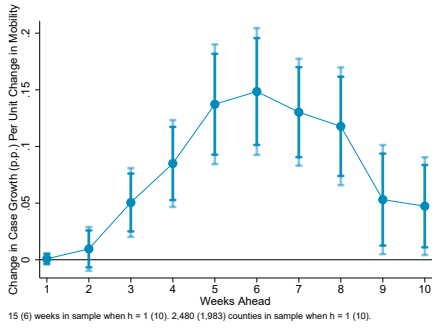

(c) At Transit Stations

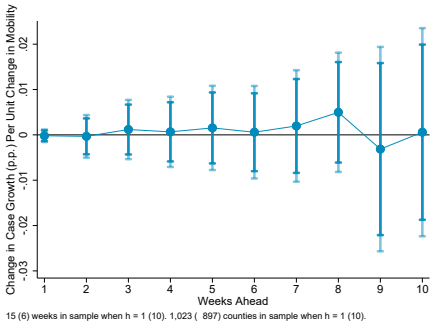

(d) At Parks

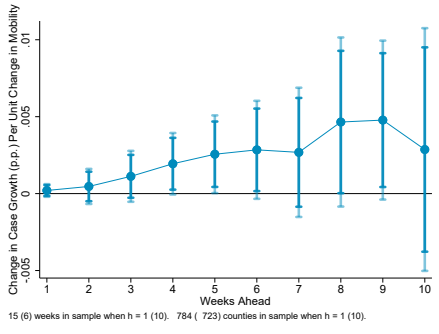

(e) At Retail and Recreation

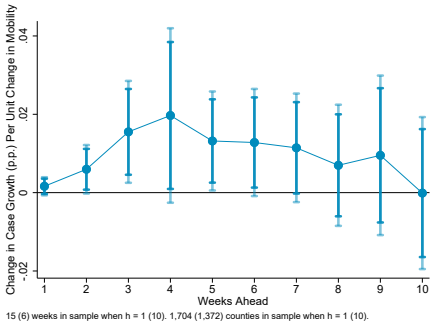

(f) At Grocery and Pharmacy

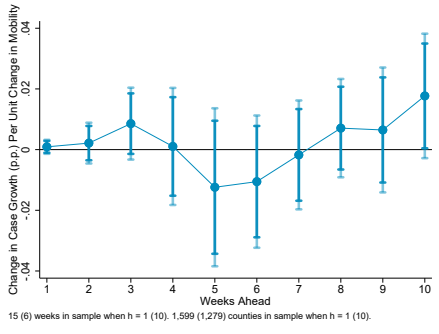

Note: Estimates of equation 2 in the text using panel local projections regressions. Shaded regions are 90% and 95% confidence intervals.

Covid Economics 35, 7 July 2020: 257-299

Figure 10: Dynamic Impacts of Mobility on COVID-19 Deaths Growth – Weekly

Impulse Response Functions Estimated by Panel Linear Projections

Mobility Measured by Google Mobility Reports

(a) At Residence (Inverted)

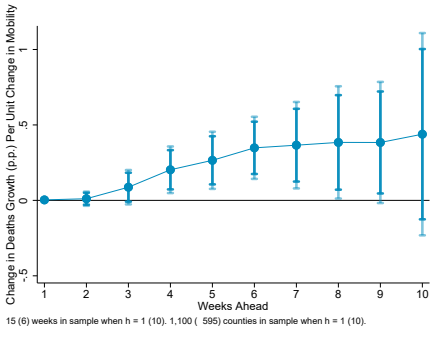

(b) At Work

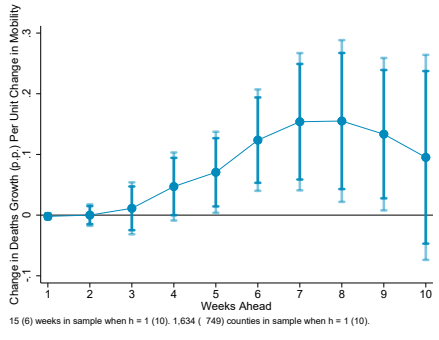

(c) At Transit Stations

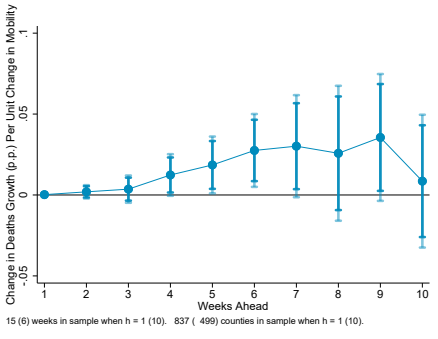

(d) At Parks

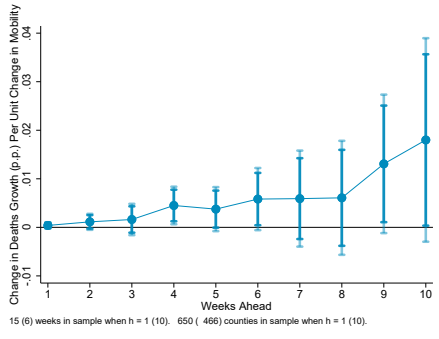

(e) At Retail and Recreation

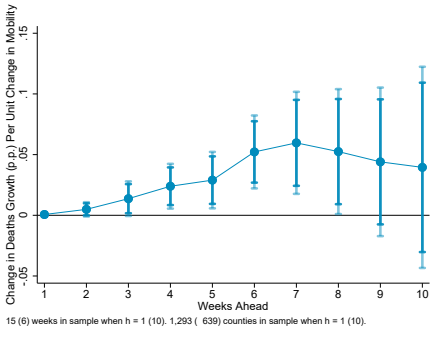

(f) At Grocery and Pharmacy

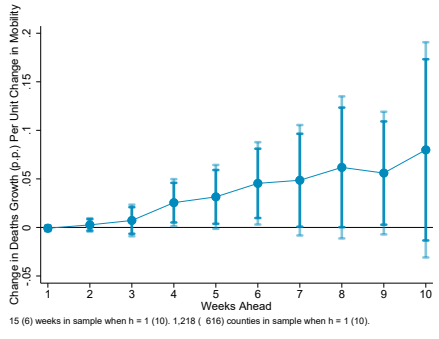

Note: Estimates of equation 2 in the text using panel local projections regressions. Shaded regions are 90% and 95% confidence intervals.

Table 2: OLS Estimates of Effect on 4-Week Ahead Cases Growth

	(1) At Home (Inv.)	(2) At Work	(3) At Transit	(4) At Parks	(5) At Retail & Rec	(6) At Grocery & Pharmacy	(7) MEI
Mobility	0.244***	0.0850***	0.000666	0.00194*	0.0197*	0.00104	0.00672
	(0.0710)	(0.0196)	(0.00395)	(0.00102)	(0.0114)	(0.00984)	(0.00526)
L.Mobility	0.0479	0.0404*	0.00698	0.000238	-0.0154	-0.00307	0.00647
	(0.0647)	(0.0232)	(0.00479)	(0.00124)	(0.0118)	(0.0108)	(0.00798)
Cases Growth Rate	-0.0233	-0.0452***	-0.0137	0.0187	-0.0232*	-0.0199	-0.0462***
	(0.0175)	(0.0129)	(0.0144)	(0.0175)	(0.0136)	(0.0141)	(0.0119)
L.Cases Growth Rate	0.0273	0.00595	0.00316	0.000409	0.0153	0.0170	0.00644
	(0.0220)	(0.00461)	(0.00458)	(0.00488)	(0.0112)	(0.0129)	(0.00396)
Testing Growth Rate	0.240	1.351	4.702	6.729**	0.0618	-0.551	1.988
	(2.465)	(1.926)	(3.126)	(2.817)	(2.238)	(2.394)	(1.964)
Cases Per Capita	-380.9***	-377.1***	-297.0***	-197.2***	-379.8***	-424.3***	-393.3***
	(95.19)	(57.00)	(55.57)	(50.41)	(82.78)	(95.09)	(50.71)
Max Daily Temp	-0.0259***	-0.00523	-0.00772	-0.00794*	-0.00782	-0.00254	-0.00103
	(0.00610)	(0.00551)	(0.00559)	(0.00437)	(0.00515)	(0.00528)	(0.00510)
Precipitation	-0.00662	-0.0163***	-0.0159**	-0.00159	-0.00963*	-0.0146**	-0.0201***
	(0.00518)	(0.00543)	(0.00682)	(0.00555)	(0.00537)	(0.00590)	(0.00579)
Snowfall	0.0211	0.0185	-0.0622	-0.0251	-0.0462	-0.0187	-0.0136
	(0.0538)	(0.0452)	(0.0551)	(0.0448)	(0.0439)	(0.0424)	(0.0477)
Observations	10830	20378	8759	6429	12899	11879	21464
Adjusted R^2	0.145	0.175	0.153	0.123	0.134	0.129	0.165
Dep. Variable Mean	0.667	0.698	0.680	0.637	0.688	0.707	0.693
Mobility Mean	-13.949	-32.992	-23.406	19.499	-24.263	-3.301	-63.189
Elasticity at means	5.11	4.02	.02	.06	.690	0	.61

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Covid Economics 35, 7 July 2020: 257-299

zon, in Tables 2 and 3. Table 2 shows the results for cases growth 4 weeks ahead and Table 3 shows the results for deaths growth 6 weeks ahead. Beyond the results for mobility already discussed, one can see here that current cases per capita strongly predicts lower subsequent case growth. This finding is consistent with the notion that cases per capita proxies for the share of the population no longer (or less likely to be) susceptible to infection; as this share rises, the capacity for future cases falls. I also find current case growth tends to negatively predict future case growth, though the effect is not consistently significant across specifications with different mobility measures. Current testing growth tends to have a positive coefficient, but it is only statistically significant in 1 of the 7 specifications.

The results are qualitatively similar for future deaths growth (Table 3), though growth of cases and testing both have larger and more statistically significant effects while cases per capita's effect is smaller.

The coefficients on temperature in Tables 2 and 3 also warrant some discussion. The temperature coefficient in each table for the specification measuring mobility by time spent away from home – column 1 – corresponds to the IRF at $h = 4$ in panel (b) of Figure 7 for cases growth and at $h = 6$ in panel (b) of Figure 8 for deaths growth. The tables show that the statistically significant negative effect of temperature on cases is not robust across specifications using different measures of mobility, suggesting the evidence on temperature's

Covid Economics 35, 7 July 2020: 257-299

Table 3: OLS Estimates of Effect on 6-Week Ahead Deaths Growth

	(1) At Home (Inv.)	(2) At Work	(3) At Transit	(4) At Parks	(5) At Retail & Rec	(6) At Grocery & Pharmacy	(7) MEI
Mobility	0.348***	0.124***	0.0275**	0.00582*	0.0522***	0.0454**	0.0355***
	(0.105)	(0.0426)	(0.0115)	(0.00327)	(0.0153)	(0.0216)	(0.0121)
L.Mobility	0.181	0.0625	-0.00314	0.00488	0.0119	-0.000318	0.0252*
	(0.110)	(0.0438)	(0.0119)	(0.00396)	(0.0157)	(0.0134)	(0.0141)
Deaths Growth Rate	-0.00863	0.00494	-0.0125	-0.0152	-0.00913	-0.00770	0.00324
	(0.0213)	(0.0164)	(0.0264)	(0.0312)	(0.0205)	(0.0210)	(0.0167)
L.Deaths Growth Rate	0.00366	0.00154	0.00384	0.00618*	0.00341	0.00357	0.00173
	(0.00230)	(0.00159)	(0.00311)	(0.00343)	(0.00211)	(0.00229)	(0.00158)
Cases Growth Rate	0.0806**	0.0644***	0.115***	0.148**	0.0812**	0.0791**	0.0672***
	(0.0353)	(0.0214)	(0.0421)	(0.0599)	(0.0328)	(0.0354)	(0.0215)
Testing Growth Rate	33.61***	28.44***	36.39***	40.39***	33.81***	34.49***	29.89***
	(10.36)	(9.355)	(11.35)	(12.25)	(10.40)	(10.68)	(9.476)
Cases Per Capita	-147.3**	-57.07	-140.3**	-174.9**	-139.3**	-158.4**	-60.41*
	(64.68)	(37.04)	(67.48)	(86.10)	(64.16)	(66.49)	(35.23)
Max Daily Temp	-0.0397***	-0.0264**	-0.0276*	-0.0370**	-0.0259*	-0.0356**	-0.0246*
	(0.0152)	(0.0130)	(0.0166)	(0.0185)	(0.0151)	(0.0156)	(0.0128)
Precipitation	0.00200	-0.00147	0.00309	0.00157	-0.00251	0.00225	0.00313
	(0.0113)	(0.00915)	(0.0133)	(0.0146)	(0.0114)	(0.0107)	(0.00890)
Snowfall	0.0669	-0.00117	0.00215	-0.0167	0.0242	0.0112	-0.00225
	(0.0843)	(0.0644)	(0.0780)	(0.0874)	(0.0771)	(0.0708)	(0.0662)
Observations	5434	7414	4422	3730	5865	5578	7506
Adjusted R^2	0.100	0.072	0.083	0.102	0.089	0.085	0.073
Dep. Variable Mean	0.684	0.552	0.765	0.863	0.655	0.678	0.548
Mobility Mean	-15.346	-36.804	-29.378	15.236	-29.342	-6.247	-73.172
Elasticity at means	7.81	8.24	1.06	.1	2.34	.42	4.74

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

negative effect on cases should be viewed with some caution. Temperature has a more robust negative effect on growth in deaths, as can be seen in Table 3.

As noted in Section 3.2, one can also estimate the dynamic effects of mobility on *cumulative* growth in cases or deaths out to any given horizon. Tables 4 and 5 report the effects of current mobility and other variables on cumulative growth in cases and deaths, respectively, over the subsequent 8 weeks. 90% confidence intervals are shown below each coefficient. The 8-week ahead cumulative effect is positive in all regressions. For cumulative growth in deaths, it is statistically significant (below the 10% level) for all measures of mobility except parks and grocery & pharmacy. For cumulative growth in cases, it is significant for mobility measured by time spent away from home, by time spent at work, and by the MEI. The effects of mobility are, in general, quantitatively large. In particular, the coefficient on time spent away from home in Table 4 of **2.440** implies that a one percentage point increase (decrease) in that mobility measure leads to an increase (decrease) in cumulative case growth over the subsequent 8 weeks of about **2.4** percentage points, or about one-third of average 8-week case growth in the sample of **7.1**% (shown at the bottom of the table). The lower bound of the 90% confidence interval is **1.657**, implying an effect on 8-week case growth of **23**% of its sample average. The effect magnitudes, expressed as elasticities, are shown at the bottom of the table. For example, the **0.244** coefficient implies an elasticity of **5.6**, implying that a 10% increase in mobility leads to **56**% higher cumulative case growth 8 weeks ahead. The elasticities are even higher for cumulative deaths growth (Table 5). For instance, the elasticity with respect to time spent away from home is found to be **8.7**.

Lastly, it is interesting to consider the adjusted-R^2's, shown at the bottom of the tables. The regressors in the model, including the county and week fixed effects, explain only as much as **16**% of the variation in subsequent 8 weeks cumulative case growth and as much as **15**% of variation in subsequent 8 weeks cumulative deaths growth. In other words, **84** to **85**% of the variation in these COVID-19 outcomes is not readily explainable by time-invariant county characteristics, common national time-varying factors, current case and testing growth, current cases per capita, weather, and mobility.

In sum, overall mobility is found to have a large positive effect on subsequent growth in COVID-19 cases and deaths. The effects become significant around 2 weeks ahead and persist through around 8 weeks for cases and around 10 weeks for deaths. The peak effect occurs 4 to 6 weeks ahead for cases and around 8 to 9 weeks ahead for deaths. Looking across subcategories of mobility, the effects are clearest for time spent at workplaces and retail & recreation, though there is also evidence of an adverse effect on deaths growth for mobility at transit stations, at grocery & pharmacy, and at parks.

Covid Economics 35, 7 July 2020: 257-299

Covid Economics 35, 7 July 2020: 257-299

Table 4: Estimates of Cumulative Effect on Cases Growth over Subsequent 8 Weeks

	(1) At Home (Inv.)	(2) At Work	(3) At Transit	(4) At Parks	(5) At Retail & Rec	(6) At Grocery & Pharmacy	(7) MEI
Mobility	2.440***	1.115***	0.00809	0.0208	0.0932	-0.0306	0.114*
	[1.657,3.224]	[0.864,1.366]	[-0.0608,0.0770]	[-0.00350,0.0452]	[-0.00731,0.194]	[-0.163,0.101]	[0.00646,0.221]
L.Mobility	0.640	0.266	0.0975*	0.00175	-0.132	0.0394	0.0384
	[-0.100,1.381]	[-0.0575,0.589]	[0.00349,0.191]	[-0.0242,0.0277]	[-0.274,0.00924]	[-0.0966,0.175]	[-0.0636,0.141]
Cases Growth Rate	-0.130	-0.245***	-0.103	0.0948	-0.151*	-0.123	-0.234***
	[-0.279,0.0192]	[-0.376,-0.115]	[-0.268,0.0610]	[-0.104,0.294]	[-0.291,-0.0119]	[-0.264,0.0176]	[-0.368,-0.101]
L.Cases Growth Rate	0.0804*	0.0299**	0.0409	-0.0174	0.0556*	0.0586*	0.0286**
	[0.00355,0.157]	[0.00702,0.0527]	[-0.00462,0.0864]	[-0.0646,0.0297]	[0.00701,0.104]	[0.00405,0.113]	[0.00927,0.0480]
Testing Growth Rate	10.43	13.40	29.41	45.71**	9.052	5.724	17.13
	[-16.26,37.11]	[-6.237,33.04]	[-3.865,62.69]	[15.91,75.50]	[-15.21,33.32]	[-21.30,32.75]	[-3.830,38.09]
Cases Per Capita	-2011.3***	-2363.8***	-1945.8***	-1170.9**	-2570.8***	-2235.3***	-2567.5***
	[-2919.4,-1103.2]	[-3056.2,-1671.4]	[-2839.4,-1052.1]	[-1933.7,-408.1]	[-3439.6,-1702.1]	[-3143.1,-1327.5]	[-3265.1,-1869.9]
Max Daily Temp	-0.173***	0.0356	-0.0943	-0.0419	-0.000749	-0.0131	0.0473
	[-0.275,-0.0716]	[-0.0320,0.103]	[-0.195,0.00603]	[-0.117,0.0328]	[-0.0878,0.0863]	[-0.106,0.0795]	[-0.0207,0.115]
Precipitation	0.0827	0.00868	-0.0323	0.0249	-0.0232	-0.00414	-0.0853
	[-0.0224,0.188]	[-0.0730,0.0903]	[-0.152,0.0872]	[-0.0978,0.148]	[-0.119,0.0725]	[-0.102,0.0939]	[-0.175,0.00431]
Snowfall	-0.0809	0.413	-0.377	-0.0737	-0.0624	-0.171	0.389
	[-0.613,0.452]	[-0.297,1.123]	[-0.968,0.215]	[-0.539,0.392]	[-0.511,0.386]	[-0.693,0.352]	[-0.359,1.138]
Observations	5949	10904	4981	3897	7066	6541	11361
Adjusted R^2	0.136	0.164	0.089	0.073	0.127	0.085	0.137
Dep. Variable Mean	7.111	7.386	7.356	6.707	7.292	7.345	7.390
Mobility Mean	-16.203	-36.998	-30.324	5.358	-33.241	-9.251	-76.487
Elasticity at means	5.56	5.58	.03	.02	.43	-.04	1.18

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Covid Economics 35, 7 July 2020: 257-299

Table 5: Estimates of Cumulative Effect on Deaths Growth over Subsequent 8 Weeks

	(1) At Home (Inv.)	(2) At Work	(3) At Transit	(4) At Parks	(5) At Retail & Rec	(6) At Grocery & Pharmacy	(7) MEI
Mobility	3.762***	1.583***	0.223*	0.0550	0.481***	0.318	0.382***
	[1.495,6.029]	[0.649,2.518]	[0.0203,0.425]	[-0.00918,0.119]	[0.181,0.780]	[-0.0619,0.697]	[0.180,0.584]
L.Mobility	3.221**	0.868*	-0.0000722	0.0451	0.123	0.0279	0.375**
	[0.827,5.614]	[0.0564,1.680]	[-0.196,0.196]	[-0.0234,0.114]	[-0.208,0.454]	[-0.191,0.247]	[0.0826,0.668]
Deaths Growth Rate	-0.0565	0.0309	-0.107	-0.117	-0.0669	-0.0516	0.00765
	[-0.340,0.227]	[-0.197,0.259]	[-0.446,0.231]	[-0.512,0.277]	[-0.341,0.207]	[-0.331,0.228]	[-0.226,0.241]
L.Deaths Growth Rate	0.0252	0.00994	0.0286	0.0460*	0.0266*	0.0248	0.0114
	[-0.00212,0.0526]	[-0.00930,0.0292]	[-0.00647,0.0637]	[0.00744,0.0846]	[0.00132,0.0519]	[-0.00235,0.0519]	[-0.00799,0.0308]
Cases Growth Rate	0.791**	0.661***	1.064**	1.407**	0.812**	0.818**	0.674***
	[0.184,1.398]	[0.254,1.068]	[0.337,1.791]	[0.399,2.415]	[0.216,1.408]	[0.177,1.459]	[0.250,1.097]
Testing Growth Rate	223.9***	198.1***	250.4***	270.0***	234.3***	240.4***	215.2***
	[113.4,334.5]	[92.45,303.8]	[124.3,376.5]	[136.0,404.0]	[118.5,350.1]	[119.2,361.7]	[106.7,323.8]
Cases Per Capita	-1649.3**	-1062.2**	-2065.0***	-2183.7***	-1822.4**	-2060.2***	-1140.7**
	[-2763.5,-535.1]	[-1837.7,-286.7]	[-3309.8,-820.2]	[-3526.9,-840.6]	[-2983.2,-661.7]	[-3269.9,-850.4]	[-1927.3,-354.2]
Max Daily Temp	-0.519***	-0.305**	-0.319*	-0.437**	-0.292**	-0.377**	-0.320**
	[-0.806,-0.231]	[-0.524,-0.0857]	[-0.590,-0.0475]	[-0.723,-0.150]	[-0.525,-0.0588]	[-0.630,-0.124]	[-0.532,-0.108]
Precipitation	0.101	0.101	-0.0373	-0.00683	-0.0507	0.0198	0.0377
	[-0.119,0.320]	[-0.0807,0.284]	[-0.324,0.250]	[-0.306,0.292]	[-0.301,0.199]	[-0.199,0.239]	[-0.147,0.222]
Snowfall	0.0160	-0.352	-0.0735	-0.286	0.0405	-0.165	-0.207
	[-1.118,1.150]	[-1.194,0.490]	[-1.162,1.015]	[-1.467,0.895]	[-1.003,1.084]	[-1.130,0.801]	[-1.137,0.723]
Observations	3500	4636	2905	2564	3786	3623	4692
Adjusted R^2	0.148	0.112	0.118	0.141	0.123	0.119	0.109
Dep. Variable Mean	7.256	5.972	8.054	8.695	6.906	7.137	5.930
Mobility Mean	-16.771	-40.198	-34.146	8.432	-34.570	-10.026	-81.850
Elasticity at means	8.70	10.66	.940	.05	2.41	.45	5.27

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5 Extensions

5.1 The Role of Public Health Policies

To the extent that public health policies, generally known as non-pharmaceutical interventions (NPIs), affect COVID-19 outcomes, their effects are likely to work primarily through the channel of affecting individuals' mobility/social-distancing behavior. Indeed, I find that the number of NPIs in place in a county (or state of the county where county NPI data is unavailable) has a strong reducing effect on mobility when estimating equation 1 with the number of NPIs added as a regressor. This is shown in Table 6, which provides the results from daily panel fixed effects regressions of each measure of mobility on the number of local NPIs in place as well the set of control variables used previously. I find that NPIs have a positive and significant effect on all measures of mobility. Indeed, it is clear to BOTH weather and policy interventions have large separate influences on mobility. However, it is possible that they have additional effects through other channels such as encouraging people to wash their hands, wear masks, and stay physically distant from others even when spending time in public places. Also, empirically, NPIs could be found to have direct effects on COVID-19 outcomes if mobility is incompletely measured.

I assess the effects of NPIs, holding measured mobility fixed, by adding the number of local NPIs in place as of date t to the estimations of equation 2 using cumulative growth in cases or deaths as the dependent variable. The results are shown in Tables 7 and 8. Consistent with the hypothesis that NPIs work entirely through the mobility channel, I find that NPIs have no significant effect on COVID-19 cases growth or deaths growth when the most relevant measures of mobility are included in the regression. In particular, NPIs have only a weakly significant effect in two of the seven specifications. In the regressions for cumulative deaths growth (Table 8), mobility has a significant effect in all specifications while NPIs are insignificant in all specifications. It is also worth noting that, comparing the results in Tables 7 and 8 to the analogous results in Tables 4 and 5 which omit NPIs, the coefficients on mobility are virtually unchanged by including NPIs. Lastly, note that the adjusted-R^2's are virtually unaffected by including NPIs.

In sum, it appears that the impact of public health NPIs works entirely through affecting individual mobility behavior.

5.2 Heterogeneous Mobility Effects

Lastly, I investigate possible heterogeneity in the average treatment effect of mobility on COVID-19 outcomes. I consider heterogeneity across several dimensions: the timing of their

Covid Economics 35, 7 July 2020: 257-299

Covid Economics 35, 7 July 2020: 257-299

Table 6: Effect of Non-Pharmaceutical Interventions on Mobility

	(1) At Home (Inverted)	(2) At Work	(3) At Transit	(4) At Parks	(5) At Retail & Rec	(6) At Grocery & Pharmacy	(7) MEI
Number of NPIs in place (out of 10)	-0.191***	-0.374***	-0.712***	-1.509***	-0.883***	-0.664***	-0.448***
	(0.0216)	(0.0388)	(0.118)	(0.511)	(0.103)	(0.0911)	(0.106)
Max Daily Temp – Weekday	0.0640***	0.0511***	0.121***	1.191***	0.135***	0.148***	0.0948***
	(0.00322)	(0.00571)	(0.0135)	(0.0779)	(0.0128)	(0.0109)	(0.0150)
Max Daily Temp – Weekend	-0.0214***	-0.0568***	-0.0391*	0.690***	0.0814***	0.128***	-0.0579**
	(0.00623)	(0.0105)	(0.0203)	(0.0959)	(0.0172)	(0.0176)	(0.0244)
Precipitation – Weekday	-0.0196***	-0.00653*	-0.0438***	-0.442***	-0.0267***	-0.0337***	-0.101***
	(0.00303)	(0.00376)	(0.00798)	(0.0626)	(0.00930)	(0.00840)	(0.0120)
Precipitation – Weekend	-0.0483***	-0.0445***	-0.138***	-0.828***	-0.0742***	-0.0868***	-0.167***
	(0.00667)	(0.00997)	(0.0203)	(0.127)	(0.0185)	(0.0195)	(0.0265)
Snowfall – Weekday	-0.0844**	-0.113***	-0.294***	-1.283***	-0.346***	-0.351***	-0.102
	(0.0347)	(0.0408)	(0.0992)	(0.299)	(0.0904)	(0.0863)	(0.126)
Snowfall – Weekend	-0.322***	-0.569***	-1.115***	-3.186***	-0.700***	-0.564**	-1.198***
	(0.0693)	(0.133)	(0.257)	(0.921)	(0.182)	(0.237)	(0.249)
Cases Growth	0.000579	0.00246***	-0.000208	-0.0183*	-0.00492***	0.00418**	-0.00310*
	(0.000443)	(0.000630)	(0.00243)	(0.0106)	(0.00159)	(0.00168)	(0.00180)
Testing Growth	-0.0255*	-0.00543	0.00533	-0.437	-0.152*	0.0812	-0.566***
	(0.0144)	(0.0269)	(0.0604)	(0.290)	(0.0790)	(0.0593)	(0.0864)
Observations	94179	192141	80563	51445	125224	115410	209196
Adjusted R^2	0.920	0.884	0.842	0.657	0.846	0.759	0.873
Earliest date	29feb2020	29feb2020	29feb2020	29feb2020	29feb2020	29feb2020	29feb2020
Latest date	23jun2020	23jun2020	23jun2020	23jun2020	23jun2020	23jun2020	20jun2020
# of days	116	116	116	116	116	116	113
# of counties	1265	2474	1037	850	2032	1956	2646

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Covid Economics 35, 7 July 2020: 257-299

Table 7: OLS Estimates of Effect on 4-Week Ahead Cases Growth

Controlling for Number of Local Non-Pharmaceutical Interventions in Place

	(1) At Home (Inv.)	(2) At Work	(3) At Transit	(4) At Parks	(5) At Retail & Rec	(6) At Grocery & Pharmacy	(7) MEI
Mobility	0.267***	0.0874***	0.0000754	0.00192*	0.0201	0.000465	0.00707
	(0.0786)	(0.0202)	(0.00399)	(0.00102)	(0.0127)	(0.0104)	(0.00554)
L.Mobility	0.0534	0.0410*	0.00676	0.000140	-0.0163	-0.00342	0.00635
	(0.0654)	(0.0234)	(0.00488)	(0.00124)	(0.0123)	(0.0110)	(0.00802)
Cases Growth Rate	-0.0245	-0.0455***	-0.0144	0.0186	-0.0236*	-0.0207	-0.0463***
	(0.0180)	(0.0131)	(0.0148)	(0.0175)	(0.0139)	(0.0144)	(0.0121)
L.Cases Growth Rate	0.0277	0.00598	0.00308	-0.0000288	0.0154	0.0171	0.00645
	(0.0223)	(0.00469)	(0.00461)	(0.00490)	(0.0114)	(0.0131)	(0.00403)
Testing Growth Rate	0.259	1.514	4.910	6.763**	0.202	-0.355	2.251
	(2.560)	(1.963)	(3.149)	(2.810)	(2.295)	(2.467)	(2.010)
Cases Per Capita	-381.9***	-377.3***	-296.3***	-195.8***	-380.1***	-425.1***	-393.4***
	(95.84)	(57.11)	(55.75)	(50.28)	(83.31)	(95.92)	(50.78)
Max Daily Temp	-0.0263***	-0.00401	-0.00750	-0.00798*	-0.00736	-0.00209	-0.000591
	(0.00617)	(0.00555)	(0.00560)	(0.00443)	(0.00514)	(0.00526)	(0.00507)
Precipitation	-0.00553	-0.0157***	-0.0163**	-0.00179	-0.00948*	-0.0148**	-0.0200***
	(0.00536)	(0.00549)	(0.00683)	(0.00554)	(0.00544)	(0.00595)	(0.00587)
Snowfall	0.0300	0.0229	-0.0610	-0.0280	-0.0463	-0.0189	-0.0132
	(0.0563)	(0.0455)	(0.0551)	(0.0448)	(0.0455)	(0.0436)	(0.0477)
# of NPIs in place	0.0256	-0.00177	-0.0660*	-0.0474*	-0.0313	-0.0334	-0.0324
	(0.0403)	(0.0283)	(0.0343)	(0.0266)	(0.0402)	(0.0371)	(0.0277)
Observations	10505	19922	8584	6414	12548	11575	20997
Adjusted R^2	0.146	0.175	0.154	0.124	0.134	0.129	0.166
Dep. variable mean	0.677	0.705	0.688	0.636	0.697	0.718	0.698
Mobility mean	-13.970	-33.020	-23.711	19.472	-24.549	-3.440	-63.203
Elasticity at means	5.5	4.09	0	.06	.71	0	.64

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Covid Economics 35, 7 July 2020: 257-299

Table 8: OLS Estimates of Effect on 6-Week Ahead Deaths Growth

Controlling for Number of Local Non-Pharmaceutical Interventions in Place

	(1) At Home (Inv.)	(2) At Work	(3) At Transit	(4) At Parks	(5) At Retail & Rec	(6) At Grocery & Pharmacy	(7) MEI
Mobility	0.364***	0.123***	0.0263**	0.00584*	0.0515***	0.0458**	0.0349***
	(0.106)	(0.0419)	(0.0112)	(0.00329)	(0.0155)	(0.0218)	(0.0123)
L.Mobility	0.183	0.0614	-0.00365	0.00491	0.0117	-0.00169	0.0250*
	(0.113)	(0.0444)	(0.0123)	(0.00397)	(0.0162)	(0.0138)	(0.0142)
Deaths Growth Rate	-0.00832	0.00509	-0.0130	-0.0151	-0.00933	-0.00781	0.00313
	(0.0218)	(0.0167)	(0.0267)	(0.0312)	(0.0209)	(0.0214)	(0.0170)
L.Deaths Growth Rate	0.00375	0.00156	0.00393	0.00620*	0.00342	0.00368	0.00175
	(0.00236)	(0.00162)	(0.00313)	(0.00344)	(0.00217)	(0.00236)	(0.00162)
Cases Growth Rate	0.0834**	0.0663***	0.118***	0.148**	0.0833**	0.0803**	0.0694***
	(0.0364)	(0.0218)	(0.0430)	(0.0597)	(0.0337)	(0.0361)	(0.0219)
Testing Growth Rate	33.72***	28.74***	36.41***	40.40***	33.95***	34.79***	30.12***
	(10.40)	(9.402)	(11.35)	(12.24)	(10.45)	(10.71)	(9.514)
Cases Per Capita	-146.2**	-56.65	-136.4**	-172.7**	-138.4**	-156.4**	-59.86*
	(64.95)	(36.55)	(67.61)	(86.26)	(64.34)	(66.64)	(34.79)
Max Daily Temp	-0.0408***	-0.0265**	-0.0271	-0.0371**	-0.0262*	-0.0361**	-0.0250*
	(0.0156)	(0.0133)	(0.0170)	(0.0188)	(0.0153)	(0.0161)	(0.0130)
Precipitation	0.00375	-0.000785	0.00107	0.000350	-0.00247	0.00141	0.00296
	(0.0118)	(0.00926)	(0.0135)	(0.0149)	(0.0118)	(0.0109)	(0.00925)
Snowfall	0.0629	-0.00393	0.00155	-0.0163	0.0213	0.00813	-0.00535
	(0.0835)	(0.0646)	(0.0790)	(0.0883)	(0.0769)	(0.0711)	(0.0664)
# of NPIs in place	0.00262	-0.0417	-0.0939	-0.0567	-0.0262	-0.0613	-0.0388
	(0.0614)	(0.0492)	(0.0724)	(0.0859)	(0.0582)	(0.0627)	(0.0486)
Observations	5309	7265	4357	3720	5739	5473	7349
Adjusted R^2	0.101	0.073	0.084	0.102	0.089	0.085	0.074
Dep. variable mean	0.696	0.559	0.774	0.865	0.664	0.687	0.553
Mobility mean	-15.368	-36.882	-29.649	15.190	-29.570	-6.378	-73.267
Elasticity at means	8.029999999999999	8.14	1.01	.1	2.29	.43	4.62

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

local onset of COVID-19 spread (first case), the timing of their initial plunge in mobility, the share of local population in nursing homes, the level of cases per capita at the time of the mobility shock, and the direction of the mobility change. To systematically investigate heterogeneity, for each dimension I construct a low/high (or early/late) indicator variable, interact it with current and lagged mobility, and add the interactions to the local projections specification (equation 2). I then plot the IRFs separately for high and low (early and late) observations. The results, using mobility measured by time spent away from home, are shown in Figure 11 for case growth and in Figure 12 for deaths growth.

Covid Economics 35, 7 July 2020: 257-299

Covid Economics 35, 7 July 2020: 257-299

Figure 11: Dynamic Impacts of Mobility on COVID-19 Cases Growth - Weekly

(a) Early (Blue) vs Late (Orange) Onset Counties

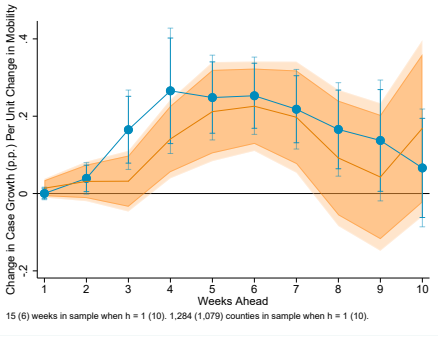

(b) Low (Blue) vs High (Orange) Cases p.c.

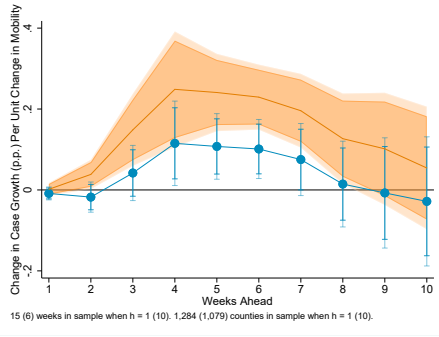

(c) Early (Blue) vs Late (Orange) Mobility Plunge Counties

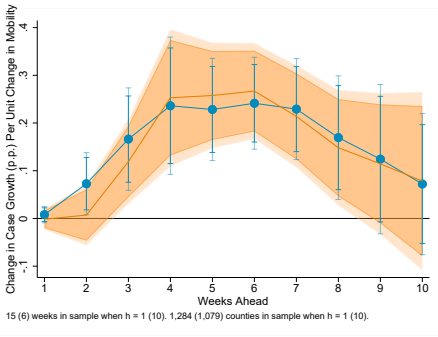

(d) Low (Blue) vs High (Orange) Nursing Home Population Share

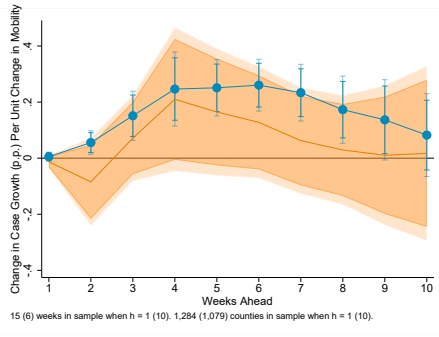

(e) Mobility Increases (Orange) vs Decreases (Blue)

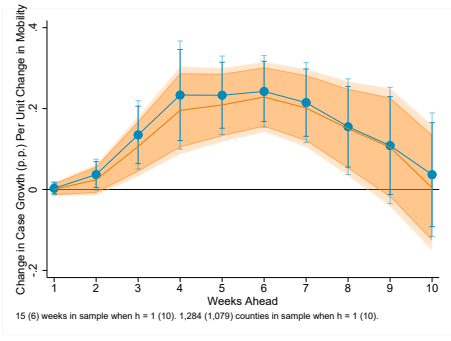

Note: Estimates of equation 2 in the text using panel local projections regressions. Shaded regions are 90% and 95% confidence intervals. Early (late) onset counties are those for which the date of their first COVID-19 case was below (above) the 50th percentile across counties.

The estimated IRFs are found not to differ across these dimensions with a couple of interesting exceptions. First, there is some evidence, especially for cases growth, that mobility has more deleterious effects when cases per capita are high. One possible explanation is that high cases per capita is correlated with high shares of the population that have undetected (possibly even asymptomatic) infections such that a given increase in overall population mobility (increasing the contact rate) interacts with a higher population share that is infectious. Second, the harmful effect of mobility on deaths from 2 to 5 weeks ahead is larger for counties which experienced their largest weekly drop in mobility at a relatively late date. This finding is consistent with the notion that for a given magnitude of social distancing, earlier action is more effective than later action.

Covid Economics 35, 7 July 2020: 257-299

5.3 The Power of Additional Data

As mentioned at the beginning of the paper, the ability to precisely estimate the full dynamic response of mobility, weather, and other variables on subsequent COVID-19 outcomes is made possible by the increasing availability of geographically granular, high-frequency, real-time data combined with a sufficient passage of time. To assess how the key results of the paper have evolved in recent weeks as data have accumulated, I re-estimate the mobility IRFs presented in Section 4.3 (based on equation 2) repeatedly using a series of expanding window samples. Specifically, I estimate the IRFs first using data through week 15 of 2020 (the week ending April 14), then using data through week 16, then using data through week 17, and so on until the latest week of available data (i.e., the same sample used in Section 4.3 and throughout the paper).

To summarize how the results have evolved with additional weeks of data, I plot the 4-week-ahead mobility elasticity (and its confidence interval) for both cases growth and deaths growth against the end-of-sample week. Recall the elasticity (evaluated at sample means) is the IRF coefficient times the ratio of mean mobility to the mean of the dependent variable (cases growth or deaths growth) and can be interpreted as the percentage effect on cases growth or deaths growth in response to a one percent increase in mobility. Plotting the IRF coefficient itself over rolling samples could be misleading because it will change mechanically if mean mobility or the mean of the dependent variable changes with the sample, which has been the case.

The results are shown in Figure 13. I focus on the three broadest measures of mobility: time spent away from home, time spent at work, and the MEI. The panels on the left are for cases growth, while the panels on the right are for deaths growth. The clearest finding is that estimates of mobility's impact have become increasingly precise – that is, the confidence

Covid Economics 35, 7 July 2020: 257-299

Figure 12: Dynamic Impacts of Mobility on COVID-19 Deaths Growth - Weekly

(a) Early (Blue) vs Late (Orange) Onset Counties

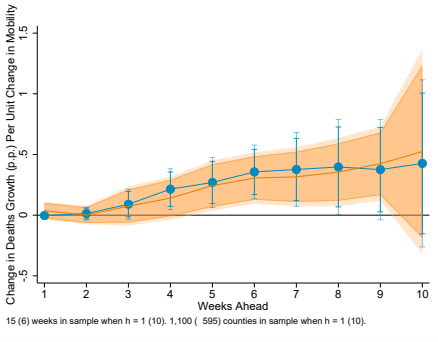

(b) Low (Blue) vs High (Orange) Cases p.c.

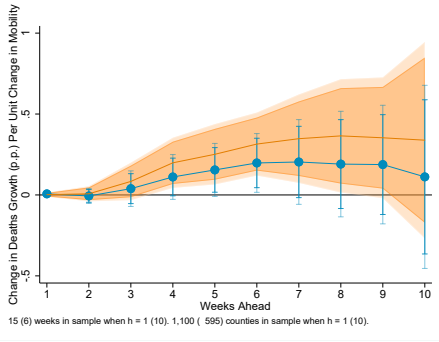

(c) Early (Blue) vs Late (Orange) Mobility Plunge Counties

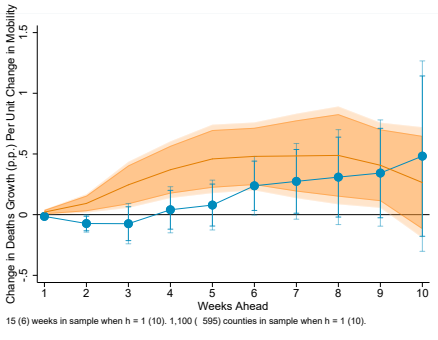

(d) Low (Blue) vs High (Orange) Nursing Home Population Share

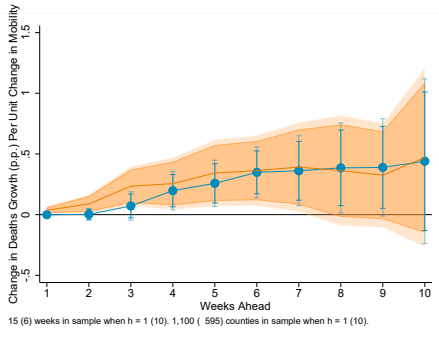

(e) Mobility Increases (Orange) vs Decreases (Blue)

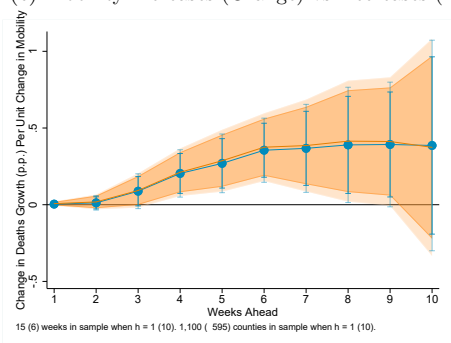

Note: Estimates of equation 2 in the text using panel local projections regressions. Shaded regions are 90% and 95% confidence intervals. Early (late) onset counties are those for which the date of their first COVID-19 case was below (above) the 50th percentile across counties.

intervals have become increasingly narrow – as more data have become available. One can also see that while the mobility elasticity shows no clear trend over expanding samples for cases growth, the elasticity has tended to shrink with additional data for deaths growth. One possible explanation for the decline in the impact of mobility on deaths is suggested by the anecdotal reports that movements in mobility have increasingly been driven by reduced social distancing of younger individuals, who are less to die from COVID-19 if infected. At any rate, these results highlight the importance of continual updates to the analysis done in this paper and similar research done by others going forward.

Covid Economics 35, 7 July 2020: 257-299

Figure 13: Rolling (Expanding-Window) Regressions
Impact of Mobility on COVID-19 Outcomes 4-Weeks Ahead

(a) Case Growth – At Home (Inv.)

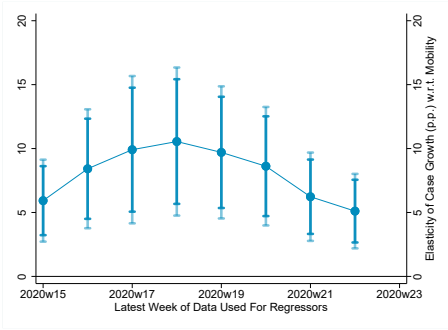

(b) Deaths Growth – At Home (Inv.)

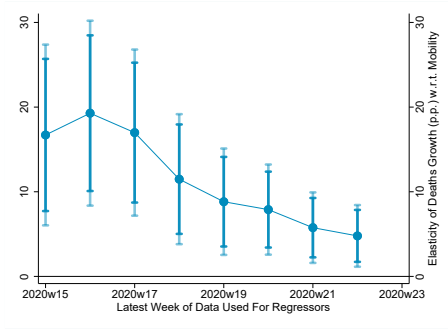

(c) Case Growth – Workplaces

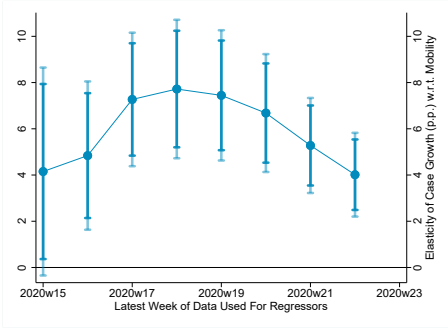

(d) Deaths Growth – Workplaces

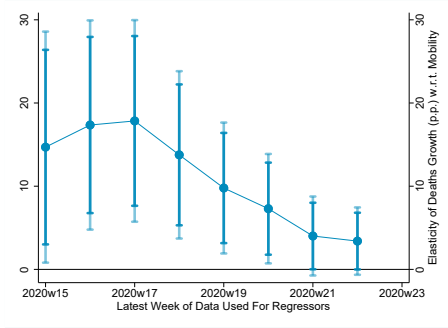

(e) Case Growth – MEI

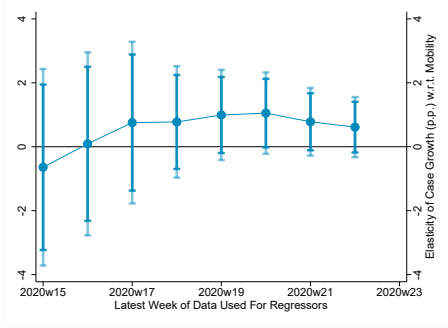

(f) Deaths Growth – MEI

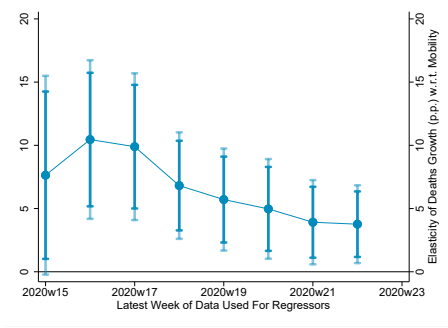

Note: Estimates of equation 2 in the text using panel local projections regressions. Shaded regions are 90% and 95% confidence intervals.

6 Conclusion

This paper sought to provide estimates of the full dynamic response of COVID-19 outcomes to exogenous movements in mobility. It uncovered several important findings. First, weather and mobility are highly correlated and thus omitting either factor when studying the COVID-19 effects of the other is likely to result in substantial omitted variable bias. Second, temperature is found to have a negative and significant effect on future COVID-19 cases and deaths, though the estimated effect is sensitive to which measure of mobility is included in the regression. Third, controlling for weather, overall mobility is found to have a large positive effect on subsequent growth in COVID-19 cases and deaths. The effects become significant around 2 weeks ahead and persist through around 8 weeks for cases and around 9 weeks for deaths. The peak effect occurs 4 to 6 weeks ahead for cases and around 8 to 9 weeks ahead for deaths. The effects are largest for mobility measured by time spent away from home and time spent at work. Fourth, I find that NPIs do affect future COVID-19 cases and deaths, but that their effects work entirely through, and not independent of, individuals' mobility behavior. Lastly, the dynamic effects of mobility were found to be generally similar across counties, though the effects are somewhat larger in places with high cases per capita and that reduced mobility relatively late.

This is a first attempt at estimating the full dynamic impact of mobility and weather on COVID-19 outcomes. As noted in the beginning of the paper, this dynamic estimation is increasingly feasible because of the availability of high-frequency, real-time data along with sufficient passage of time since the initial outbreaks in most of the U.S.. I plan to regularly revisit the analyses in this paper as further days and weeks of data become available.

Covid Economics 35, 7 July 2020: 257-299

Covid Economics 35, 7 July 2020: 257-299

References

Askitas, N., K. Tatsiramos, and B. Verheyden (2020): "Lockdown Strategies, Mobility Patterns and COVID-19," *Covid Economics*, (23), 263–302.

Atkinson, T., J. Dolmas, C. Koch, E. F. Koenig, K. Mertens, A. Murphy, and K.-M. Yi (2020): "Mobility and Engagement Following the SARS-Cov-2 Outbreak," *FRB of Dallas Working Paper*, (2014).

Badr, H., H. Du, M. Marshall, E. Dong, M. Squire, and L. M. Gardner (2020): "Social Distancing is Effective at Mitigating COVID-19 Transmission in the United States," *medRxiv*.

Carleton, T., J. Cornetet, P. Huybers, K. Meng, and J. Proctor (2020): "Ultraviolet radiation decreases COVID-19 growth rates: Global causal estimates and seasonal implications," *Available at SSRN 3588601*.

Courtemanche, C., J. Garuccio, A. Le, J. Pinkston, and A. Yelowitz (2020): "Strong Social Distancing Measures In The United States Reduced The COVID-19 Growth Rate: Study evaluates the impact of social distancing measures on the growth rate of confirmed COVID-19 cases across the United States.," *Health Affairs*, pp. 10–1377.

Desmet, K., and R. Wacziarg (2020): "Understanding Spatial Variation in COVID-19 across the United States," .

Hsiang, S., D. Allen, S. Annan-Phan, K. Bell, I. Bolliger, T. Chong, H. Druckenmiller, L. Y. Huang, A. Hultgren, E. Krasovich, et al. (2020): "The effect of large-scale anti-contagion policies on the COVID-19 pandemic," *Nature*, pp. 1–9.

Jamil, T., I. Alam, T. Gojobori, and C. M. Duarte (2020): "No evidence for temperature-dependence of the covid-19 epidemic," .

Jordà, Ò. (2005): "Estimation and inference of impulse responses by local projections," *American economic review*, 95(1), 161–182.

Kapoor, R., H. Rho, K. Sangha, B. Sharma, A. Shenoy, and G. Xu (2020): "God is in the Rain: The Impact of Rainfall-Induced Early Social Distancing on COVID-19 Outbreaks," *Available at SSRN 3605549*.

Nguyen, T. D., S. Gupta, M. Andersen, A. Bento, K. I. Simon, and C. Wing (2020): "Impacts of state reopening policy on human mobility," Discussion paper, National Bureau of Economic Research.

Soucy, J.-P. R., S. L. Sturrock, I. Berry, D. J. Westwood, N. Daneman, D. R. MacFadden, and K. A. Brown (2020): "Estimating effects of physical distancing on the COVID-19 pandemic using an urban mobility index," *medRxiv*.

Unwin, H. J. T., S. Mishra, V. C. Bradley, A. Gandy, M. Vollmer, T. Mellan, H. Coupland, K. Ainslie, C. Whittaker, J. Ish-Horowicz, et al. (2020): "State-level tracking of COVID-19 in the United States," *Imperial College London, Unpublished Report*.

Wilson, D. J. (2019): "Clearing the Fog: The Predictive Power of Weather for Employment Reports and their Asset Price Responses," *American Economic Review: Insights*, 1(3), 373–88.

Xu, R., H. Rahmandad, M. Gupta, C. DiGennaro, N. Ghaffarzadegan, H. Amini, and M. S. Jalali (2020): "The Modest Impact of Weather and Air Pollution on COVID-19 Transmission," *medRxiv*.

Covid Economics 35, 7 July 2020: 257-299

Covid Economics 35, 7 July 2020: 257-299

Appendix A – Supplemental Results

Figure A1: Relationship Between Temperature and MEI