


**CENTRE FOR
ECONOMIC
POLICY
RESEARCH**

CEPR PRESS



**COVID ECONOMICS
VETTED AND REAL-TIME PAPERS**

**AGE-SPECIFIC DECONFINEMENT
STRATEGIES**

Christian Gollier

**THE ECONOMIC EFFECT OF
CONTAINMENT MEASURES**

Pragyan Deb, Davide Furceri,
Jonathan D. Ostry and Nour Tawk

RAINFALL MATTERS

Rolly Kapoor, Haedong Aiden Rho,
Kinpritma Sangha, Bhavyaa Sharma,
Ajay Shenoy and Guanghong Xu

AIR TRAVEL RESTRICTION

Olivier Hubert

FORECASTING CORONAVIRUS

Andrew Harvey and Paul Kattuman

CHILD CARE

Umair Ali, Chris M. Herbst,
and Christos A. Makridis

**REMITTANCES: THE CASE OF
THE PHILLIPINES**

Enerelt Murakami, Satoshi Shimizutani
and Eiji Yamada

**HOW DID GERMANY FLATTEN
THE CURVE?**

Enzo Weber

**ISSUE 24
1 JUNE 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.

Copyright for the papers appearing in this issue of *Covid Economics: Vetted and Real-Time Papers* is held by the individual authors.

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.

Chair of the Board

Sir Charlie Bean

Founder and Honorary President

Richard Portes

President

Beatrice Weder di Mauro

Vice Presidents

Maristella Botticini

Ugo Panizza

Philippe Martin

Hélène Rey

Chief Executive Officer

Tessa Ogden

Editorial Board

Beatrice Weder di Mauro, CEPR

Charles Wyplosz, Graduate Institute Geneva and CEPR

Viral V. Acharya, Stern School of Business, NYU and CEPR

Abi Adams-Prassl, University of Oxford and CEPR

Jérôme Adda, Bocconi University and CEPR

Guido Alfani, Bocconi University and CEPR

Franklin Allen, Imperial College Business School and CEPR

Michele Belot, European University Institute and CEPR

David Bloom, Harvard T.H. Chan School of Public Health

Nick Bloom, Stanford University and CEPR

Tito Boeri, Bocconi University and CEPR

Alison Booth, University of Essex and CEPR

Markus K Brunnermeier, Princeton University and CEPR

Michael C Burda, Humboldt Universitaet zu Berlin and CEPR

Luis Cabral, New York University and CEPR

Paola Conconi, ECARES, Universite Libre de Bruxelles and CEPR

Giancarlo Corsetti, University of Cambridge and CEPR

Fiorella De Fiore, Bank for International Settlements and CEPR

Mathias Dewatripont, ECARES, Universite Libre de Bruxelles and CEPR

Jonathan Dingel, University of Chicago Booth School and CEPR

Barry Eichengreen, University of California, Berkeley and CEPR

Simon J Evenett, University of St Gallen and CEPR

Maryam Farboodi, MIT and CEPR

Antonio Fatás, INSEAD Singapore and CEPR

Francesco Giavazzi, Bocconi University and CEPR

Christian Gollier, Toulouse School of Economics and CEPR

Rachel Griffith, IFS, University of Manchester and CEPR

Timothy J. Hatton, University of Essex and CEPR

Ethan Ilzetzki, London School of Economics and CEPR

Beata Javorcik, EBRD and CEPR

Sebnem Kalemli-Ozcan, University of Maryland and CEPR Rik Frehen

Erik Lindqvist, Swedish Institute for Social Research (SOFI)

Tom Kompas, University of Melbourne and CEBRA

Miklós Koren, Central European University and CEPR

Anton Korinek, University of Virginia and CEPR

Philippe Martin, Sciences Po and CEPR

Warwick McKibbin, ANU College of Asia and the Pacific

Kevin Hjortshøj O'Rourke, NYU Abu Dhabi and CEPR

Evi Pappa, European University Institute and CEPR

Barbara Petrongolo, Queen Mary University, London, LSE and CEPR

Richard Portes, London Business School and CEPR

Carol Propper, Imperial College London and CEPR

Lucrezia Reichlin, London Business School and CEPR

Ricardo Reis, London School of Economics and CEPR

Hélène Rey, London Business School and CEPR

Dominic Rohner, University of Lausanne and CEPR

Paola Sapienza, Northwestern University and CEPR

Moritz Schularick, University of Bonn and CEPR

Paul Seabright, Toulouse School of Economics and CEPR

Flavio Toxvaerd, University of Cambridge

Christoph Trebesch, Christian-Albrechts-Universitaet zu Kiel and CEPR

Karen-Helene Ulltveit-Moe, University of Oslo and CEPR

Jan C. van Ours, Erasmus University Rotterdam and CEPR

Thierry Verdier, Paris School of Economics and CEPR

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.

<i>American Economic Review</i>	<i>Journal of Econometrics*</i>
<i>American Economic Review, Applied Economics</i>	<i>Journal of Economic Growth</i>
<i>American Economic Review, Insights</i>	<i>Journal of Economic Theory</i>
<i>American Economic Review, Economic Policy</i>	<i>Journal of the European Economic Association*</i>
<i>American Economic Review, Macroeconomics</i>	<i>Journal of Finance</i>
<i>American Economic Review, Microeconomics</i>	<i>Journal of Financial Economics</i>
<i>American Journal of Health Economics</i>	<i>Journal of International Economics</i>
<i>Canadian Journal of Economics</i>	<i>Journal of Labor Economics*</i>
<i>Economic Journal</i>	<i>Journal of Monetary Economics</i>
<i>Economics of Disasters and Climate Change</i>	<i>Journal of Public Economics</i>
<i>International Economic Review</i>	<i>Journal of Political Economy</i>
<i>Journal of Development Economics</i>	<i>Journal of Population Economics</i>
	<i>Quarterly Journal of Economics*</i>
	<i>Review of Economics and Statistics</i>
	<i>Review of Economic Studies*</i>
	<i>Review of Financial Studies</i>

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

Covid Economics

Vetted and Real-Time Papers

Issue 24, 1 June 2020

Contents

Cost-benefit analysis of age-specific deconfinement strategies <i>Christian Gollier</i>	1
The economic effects of Covid-19 containment measures <i>Pragyan Deb, Davide Furceri, Jonathan D. Ostry and Nour Tawk</i>	32
God is in the rain: The impact of rainfall-induced early social distancing on Covid-19 outbreaks <i>Rolly Kapoor, Haedong Aiden Rho, Kinpritma Sangha, Bhavyaa Sharma, Ajay Shenoy and Guanghong Xu</i>	76
Spatial distancing: air traffic, Covid-19 propagation, and the cost efficiency of air travel restrictions <i>Olivier Hubert</i>	111
Time series models based on growth curves with applications to forecasting coronavirus <i>Andrew Harvey and Paul Kattuman</i>	126
The impact of Covid-19 on the US child care market: Evidence from stay-at-home orders <i>Umair Ali, Chris M. Herbst, and Christos A. Makridis</i>	158
The potential impact of the Covid-19 pandemic on the welfare of remittance-dependent households in the Philippines <i>Enerelt Murakami, Satoshi Shimizutani and Eiji Yamada</i>	183
Which measures flattened the curve in Germany? <i>Enzo Weber</i>	205

Cost-benefit analysis of age-specific deconfinement strategies¹

Christian Gollier²

Date submitted: 24 May 2020; Date accepted: 27 May 2020

I calibrate a Multi-Risk SIR model on the covid pandemic to analyze the impact of the age-specific confinement and PCR testing policies on incomes and mortality. Two polar strategies emerge as potentially optimal. The suppression policy would crush the curve by confining 90% of the population for 4 months to eradicate the virus. The flatten-the-curve policy would reduce the confinement to 30% of the population for 5 months, followed by almost one year of free circulation of the virus to attain herd immunity without overwhelming hospitals. Both strategies yield a total cost of around 15% of annual GDP when combining the economic cost of confinement with the value of lives lost. I show that hesitating between the two strategies can have a huge societal cost, in particular if the suppression policy is stopped too early. Because seniors are much more vulnerable, a simple recommendation emerges to shelter them as one deconfines young and middle-aged people in order to build our collective herd immunity. By doing so, one reduces the death toll of the pandemic together with the economic cost of the confinement, and the total cost is divided by a factor 2. I also show that expanding the mass testing capacity to screen people sent back to work has a large benefit

1 I thank Nour Meddahi, Moritz Meyer-ter-Vehn, Laurent Miclo, Arnold Migus, Stefan Pollinger, the members of the French Academy of Medicine, and participants to the TSE webinar series on covid-19 for helpful comments. The research leading to these results has received the support from the ANR Grant Covid-Metrics.

2 Toulouse School of Economics, University of Toulouse-Capitole

Copyright: Christian Gollier

under various scenarios. This analysis is highly dependent upon deeply uncertain epidemiologic, sociological, economic and ethical parameters.

1 Introduction

Because the covid-19 pandemic has huge economic consequences, it makes sense for economists to explore the dynamics of this virus with the aim of searching for efficient public policies. This reminds me what economists have been able to do in the field of climate economics over the last four decades. I am neither an epidemiologist nor a climatologist, but I believe that combining these fields with economics is important for the policy debate, given the key issue of the economic cost of fighting the coronavirus or climate change.

In this paper, I explore a Multiple-Risk Susceptible-Infected-Recovered (MR-SIR) model with heterogeneous citizens. People differ by their age, by the intensity of their social interactions, by their ability to transmit the virus, and by their probability to be hospitalized or to die. If infected, they can be symptomatic or not. Their reproduction rate is a function of whether they are quarantined, confined or freed to live their life. PCR tests can be used to detect infected people among the population of apparent susceptible people (asymptomatic contagious people belong to this category). In my model, individual reproduction numbers, and therefore the dynamics of the pandemic, depend upon the implemented public policy in terms of the intensity of confinement and PCR testing, which can be age-specific. I compare the merits of different intuitive public policies.

It is by now well recognized that the laissez-faire strategy is not an efficient solution, given the high mortality rate of SARS-Cov-2 compared to the standard flu. With a herd immunity attainable only with a rate of immunity around 80% and a mortality rate around 1%, this strategy would kill 0.8% of the population, not taking account of the excess mortality due to the overcrowding of hospitals under this scenario. In the absence of treatment or vaccine, two families of health policies remain, the "suppression" (or "crush-the-curve") strategy and the "flatten-the-curve" strategy. Suppression policies consist in imposing various rules such as strict social distancing, mass testing, confinement of susceptible people and quarantine of contagious people, with the aim to reduce the reproduction number as much as possible below unity to crush the curve of infection. To illustrate, China has implemented such a strategy around Wuhan, and Italy, Spain and France have used similar strategies until early May 2020. Because the dynamics of the pandemic has an exponential nature, following such a strategy until the full eradication of the virus in the population may take time that most pandemic models measure in months if not in years. The economic cost of the suppression strategy could therefore be huge. The termination of the pandemic also requires a specific method (testing-and-tracing) to eliminate the last clusters of infection.

The "flatten-the-curve" strategy consists in imposing much weaker restrictions in order to reduce the speed of propagation of the virus so that the initial wave of infection is manageable by the health care system. This is because the ICU capacity is limited, and its overcrowding is known to dramatically increase the infected-fatality ratio of the pandemic. Stabilizing the flow of hospitalizations requires a weaker confinement than in the suppression strategy, and is therefore less economically costly. Under this family of policies, the population converges towards herd immunity, whose level depends upon the intensity of the social restrictions. At some date along this asymptotic herd immunity, the prevalence rate will become so small that eradication can also be obtained with some form of testing-and-tracing method. In spite of the preservation of the health-care system, flatten-the-curve strategies are expected to impose a much larger death toll to the population since it requires that a large fraction of the population to be infected sometimes during the pandemic.

Age Class	Prob[hospitalized if infected]	Prob[deceased if infected]
0-18	0.10%	0.001%
19-64	2.00%	0.15%
65+	12.27%	3.65%

Table 1: Estimation of the hospitalization rate and of the infection-fatality proportion by age class in France. Source: Saltje et al. (2020) and INSEE.

To sum up, compared to crush-the-curve strategies, flattening the curve typically yields more lives lost but a smaller GDP loss. Comparing the two families of policies from a welfare point of view thus requires valuing lives. Within economics, this is a relatively consensual issue. Outside economics, valuing lives is vastly rejected, in spite of the fact that public institutions have been using a "value of statistical life" for at least four decades in the Western world.¹ In this paper, I follow the traditional approach used by health and environmental economists in which health effects of the public policies under scrutiny are translated into income equivalent. I discuss this specific issue in the context of the current pandemic in Gollier (2020).

A striking feature of covid-19 is its huge differential health impact on human being across different age classes, as shown in Table 1 for the case of France. If the objective is to minimize the death toll or to make sure hospitals are not submerged, it is tempting to protect senior people from the virus. The problem is that doing so marginally increases the mortality risk for the younger generations, thereby raising the question of the relationship between the value of life and age. In this paper, I value life to 60 years, 40 years and 20 years of annual GDP/cap respectively for the young, middle-aged and senior people. I show that a strict lockdown of the most vulnerable persons in our society during the most active phase of the pandemic is an efficient policy to reduce both the death toll and the economic cost of the pandemic, independent upon whether one pursue the suppression or the flatten-the-curve objective.

A few recent papers have supported similar age-targeted deconfinement strategies. Acemoglu, Chernozhukov, Werning and Whinston (2020) characterize two intertemporally optimal exit strategies from lockdown, one in which the policy is constrained to be uniform across age classes, and the other in which different age classes can be treated in a discriminated way. They claim that 2.7 million lives could be saved in the United States by maintaining a stricter confinement for the seniors.² Favero, Ichino and Rustichini (2020) compare different age-specific policies for Italy and come to the same conclusion of the overwhelming dominance of confining elderly people longer. Fischer (2020) and Wilder et al. (2020) also support a strong sheltering of the vulnerable persons. Brotherhood, Kircher, Santos and Tertilt (2020) explore the impact of various confinement policies on the incentive of different age classes to behave efficiently. All those models share the same fundamental structure of the MR-SIR framework that I use in this paper.

¹For more information, see for example Drèze (1962), Schelling (1968), Jones-Lee (1974), Shepard and Zeckhauser (1984), Murphy and Topel (2006), Viscusi (2009), US-EPA (2010) and Quinet (2013).

²The paper of Acemoglu et al. (2020) has the great advantage to describe the Pareto frontier of efficient policies. In this paper, I don't search for an intertemporally optimal solution, but I compare different intuitive strategies that have been used, or are expected to be used in the future.

What did I learn from this analysis beyond the high social value of discriminating the intensity of the confinement across different age classes? Several things in fact. First, I learned that in the family of reasonable policies, there exists two polar solutions, one with strong restrictions with the objective to crush the curve as soon as possible, and another one with much milder restrictions just to flatten the curve. Under the calibration of my MR-SIR model, the suppression strategy necessitates confining 90% of the entire population for 125 days. The flatten-the-curve strategy imposes the confinement of 30% of the entire population for six months, followed by almost a year of a low level of the virus prevalence in the population. These two very different strategies yields a similar total cost around 15% of annual GDP, which combines the economic cost of the confinement and the value of lives lost.

Second, the planner should be strong in her policy choice. Hesitation and trembling hand can have a high welfare cost. To illustrate this point, suppose that the country initially plan to follow the suppression strategy of confining 90% of the population during 125 days. But suppose that it changes her mind after t days of strong confinement to follow the flatten-the-curve strategy in which only 30% of the population remains confined until herd immunity. Figure ?? shows the relationship between the duration of the suppression strategy and the total cost of the pandemic. For a zero duration, we have the best flatten-the-curve policy, and for a 125-day duration, we have the best suppression strategy. But stopping short of the 125-day duration of the suppression policy will generate a high-cost second wave of the pandemic. I describe this dynamic in Figure 1 in the Appendix when abandoning the suppression policy after 95 days, one month short of what is necessary to eradicate the virus. Switching to the flatten-the-curve strategy a few days short of the 125 days dramatically increases the total cost of the pandemic, which increases from 15% of annual GDP to around 24%, because of the second wave of infection that this "trembling hand" policy generates.

Third, because of the coexistence of these two polar policies yielding similar total costs, an important international coordination problem emerges. If all countries select the suppression strategy except one which follows the flatten-the-curve strategy, this contrarian country will impose a huge negative externality on the international community. Indeed, the suppression strategy has this weakness to converge to a low rate of immunization, far away from the herd immunity level in the absence of social distancing. Therefore, because flattening the curve implies a much longer duration of the pandemic, the contrarian country imposes to other countries the risk to export the virus, triggering a new worldwide wave of the pandemic. The absence of coordination implies that countries implementing stronger confinement and quarantine rules will have to limit their interactions (trade, human mobility) with other countries, yielding potentially high additional costs that I do not count in this paper.

Fourth, I characterize the tradeoff between the strength and the duration of the confinement. In the family of flatten-the-curve policies, stricter confinement rules delay the herd immunization and increases the duration of the pandemic. In the family of suppression policies, stricter confinement rules speed up the eradication of the virus and reduce the duration of the pandemic. Thus, the relationship between the intensity of the confinement and the duration of the pandemic is hump-shaped. This observation has important policy implications. In particular, when implementing the suppression strategy, the stricter the confinement rules the better.




Figure 1: Dynamics of the pandemic when following the following strategy. After 21 days of free circulation of the virus, 90% of the population is confined for 95 days. Rather than maintaining this strong confinement until day 125 where the virus is fully eradicated, one shifts to the best flatten-the-curve strategy one month short of this duration. Under this new flatten-the-curve strategy, a weak 30% confinement is established as soon as soon as the bed occupancy in hospital attains 80% of capacity during the second wave. Full exit to lockdown is decided when this occupancy rate goes below 20%.

Fifth, there is an obvious point to be made on the dominance of the mass testing strategy over the mass confinement strategy, as claimed at the beginning of the pandemic by the World Health Organization ("Test, test, test!").³ When the prevalence rate is 2%, it makes little sense to confine everyone, just to make sure that these 2% will have a small reproduction rate. If one could test all people exiting from lockdown, we could reduce the propagation of the virus without imposing the huge cost of freezing the economy. I make this important point in Section 7, since my model is able to simulate the impact of uniform or age-specific testing strategies.

This pandemic simulation exercise heavily relies on the implicit assumptions of the MR-SIR model and on the calibration of its parameters. Thus, my analysis is not more reliable than the estimation of the epidemiologic, sociological, ethical and economic parameters that feed it. As I write this paper, the uncertainty that surrounds many of these parameters should impose circumspection. Who knows the rate of asymptomatic cases, the mortality rate for the young, the immunization of the recovered, the impact of the weather, or the date of arrival of a mass vaccine? This paper, as most others on the same subject, is based on the absurd assumption of known parameters. My agenda in climate economics over the last two decades has been to determine the impact of the uncertainty surrounding the parameters of climate integrated-assessment models on the optimal climate policy. My agenda for the next two months is to do the same thing for the covid policy, based on the uncertainty-free model presented in this paper.

³<https://www.reuters.com/article/us-healthcare-coronavirus-who/test-test-test-who-chiefs-coronavirus-message-to-world-idUSKBN2132S4>

2 The MR-SIR model

In the spirit of Acemoglu et al. (2020) and Favero et al. (2020), I examine a simple SIR model with multiple classes of individuals. The whole population is partitioned in J classes $j \in \{1, \dots, J\}$. The size of class j is denoted N_j . In this paper, I partition the population by age classes, since age is an important characteristics that influence the covid mortality risk and the intensity of social interactions. Each person is either Susceptible, Infected, Recovered or Death, i.e., the health status of a person belongs to $\{S, I, R, D\}$. This implies that $S_{j,t} + I_{j,t} + R_{j,t} + D_{j,t} = N_j$ at all dates $t \geq 0$, where $S_{j,t}$ for example measures the number of persons in class j that are susceptible at date t . The total number of infected persons in the whole population at date t is denoted $I_t = \sum_j I_{j,t}$, and similarly for S_t , R_t and D_t . Although, there is no certainty about this aspect of the covid pandemic at this stage, I assume that recovered people are immunized. They are also all detected as such at no cost.

A susceptible person can be infected by meeting an infected person. Following the key assumption of all SIR models, this number of new infections is assumed to be proportional to the product of the densities of infected and susceptible persons in the population, weighted by the intensity of their social interaction. Using the Multi-Risk (MR) version of SIR modeling, and with no further justification, this is quantified as follows:

$$S_{i,t+1} - S_{i,t} = - \left(\sum_{j=1}^J \beta_{ijt} I_{j,t} \right) S_{i,t}. \quad (1)$$

I will soon describe how β_{ijt} , which measures the intensity of the risk of contagion of a susceptible person in class i by an infected person in class j at date t , is related to the social interactions between these two groups and by the confinement and testing policy. Once infected, a person of class i quits this health state at rate γ_i , so that the dynamics of the infection satisfies the following equation:

$$I_{i,t+1} - I_{i,t} = \left(\sum_{j=1}^J \beta_{ijt} I_{j,t} \right) S_{i,t} - \gamma_i I_{i,t}. \quad (2)$$

There are two exit doors to the infected status: One can either recover or die. The mortality rate among the infected persons of class i at date t is denoted π_i . It may be a function of the total number of infected people in the whole population. So, we have

$$R_{i,t+1} - R_{i,t} = (1 - \pi_i(I_t)) \gamma_i I_{i,t} \quad (3)$$

$$D_{i,t+1} - D_{i,t} = \pi_i(I_t) \gamma_i I_{i,t}. \quad (4)$$

The pandemic starts at date $t = 0$ with ϵ_j infected persons and $N_j - \epsilon_j$ susceptible persons in class j , $j \in \{1, \dots, J\}$.

Consider an infected person in class j . How many persons will he infect in period $[t, t+1]$? The answer is $\sum_i \beta_{ijt} S_{i,t}$. Because on average such a person remains infected during $1/\gamma_j$ periods, we can conclude that the reproduction number in class j at date t equals

$$\mathfrak{R}_{jt} = \frac{\sum_{i=1}^J \beta_{ijt} S_{i,t}}{\gamma_j}. \quad (5)$$

A fraction κ_j of the infected people in class j are asymptomatic. In the absence of PCR test, they cannot be identified within the population $\hat{S}_{j,t} = S_{j,t} + \kappa_j I_{j,t}$ of "apparently susceptible" persons in class j . All symptomatic persons are quarantined, whereas all recovered persons are freed from any constraint. What about the apparently susceptible persons? I examine two policy instruments. A fraction a_{jt} of $\hat{S}_{j,t}$ is PCR tested. I assume no error in the test. People tested positive are quarantined, whereas the others are freed from any constraint for period $[t, t + 1]$. Another fraction b_{jt} remains confined during that period. The remaining fraction $1 - a_{jt} - b_{jt}$ is freed from any lockdown constraint during the period in spite of their uncertain health state. As in Acemoglu et al. (2020), I assume that the confinement is imperfect in the sense that a fraction $1 - \theta_j$ of the confined people in class j ignores the constraints. As a consequence, among susceptible persons in class i , the true number of susceptible persons in that class who behave as confined is $\theta_i b_{it}$ at date t . Among the infected persons in class j at date t , a proportion

- $\kappa_j a_{jt} + 1 - \kappa_j$ is quarantined;
- $\kappa_j \theta_j b_{jt}$ is effectively confined;
- $\kappa_j (1 - a_{jt} - \theta_j b_{jt})$ is freed from the lockdown, or behaves in that way.

The intrinsic degrees $(\beta_q, \beta_c, \beta_f)$ of contagion of an infected person depends upon whether she is quarantined (q), confined (c) or freed (f), with $\beta_q < \beta_c < \beta_f$. Let α_{ij} denote the relative intensity of social interactions between a susceptible person of class i and an infected person of class j . Thus, the intensity of the contagion of a susceptible person in class i by an infected person in class j at date t equals

$$\beta_{ijt} = \alpha_{ij} \left(\beta_q (\kappa_j a_{jt} + 1 - \kappa_j) + \beta_c \kappa_j \theta_j b_{jt} + \beta_f \kappa_j (1 - a_{jt} - \theta_j b_{jt}) \right) (1 - \theta_i b_{it}) \quad (6)$$

For example, in the laissez-faire policy ($a = b = 0$), all susceptible persons work, together with all asymptomatic infected persons, whereas the symptomatic ones are quarantined. Thus, equation (6) this simplifies to

$$\beta_{ijt} = \alpha_{ij} (\beta_q (1 - \kappa_j) + \beta_f \kappa_j). \quad (7)$$

An important feature of equation (6) is that the intensity of the contagion between age classes i and j is a quadratic form of the confinement intensities b_i and b_j . In the case of a uniform confinement rule, the intensity of contagion is quadratic in the intensity b of confinement. This is due to the fact that the lockdown reduces the interaction from both sides, infected and susceptible.

How can we compare different policies in relation to their welfare impacts? Two dimensions should be taken into account. First, life is valuable, so death has a welfare cost. Let me associate a cost ℓ_j to the death of a person in age class j . The pandemic has also an economic cost associated to the deaths, quarantines, confinements and testing during the pandemic. I assume that quarantined people are unable to work. A fraction ξ_j of confined people in class j can telework. The value loss of a person in class j that cannot work is denoted w_j . For workers, w_j can be interpreted as their labour income. For young people, it's the lost human capital due to the reduced quality of their education during lockdown. For the retired people, it's the value of their contributions to the common good. We must also take account of the

economic cost of mass testing. In total, assuming a unit cost of PCR test equaling p , the economic loss of the pandemic in class j is measured as follow:

$$W_j = p \int_0^T a_{jt} \hat{S}_{j,t} dt + w_j \left((1 - \xi_j) \int_0^T b_{jt} S_{j,t} dt + \int_0^T (\kappa_j (a_{jt} + (1 - \xi_j) b_{jt}) + 1 - \kappa_j) I_{j,t} dt + \int_0^T D_{j,t} dt \right), \quad (8)$$

where T is the time horizon of the social planner, for example the discovery of a vaccine or a treatment. The total loss is thus equal to

$$L = \sum_{j=1}^J (\ell_j D_{j,T} + W_j). \quad (9)$$

A key dimension of the health policy during a pandemic is the risk of overcrowding the health care system facing limited capacities in beds, ICUs or respirators for example. I summarize this capacity problem by a capacity limit on covid beds in hospitals. I assume that a fraction $h_j < 1 - \kappa_j$ of infected people in class j faces an acute version of the virus and requires a bed to receive an efficient health care. If the overall bed capacity \bar{h} is not overwhelmed, then the mortality rate among infected people in class j is limited to π_{0j} . But of hospitals are overcrowded, then this mortality rate jump up to $\pi_{1j} > \pi_{0j}$:

$$\pi_j(I_{1t}, \dots, I_{Jt}) = \begin{cases} \pi_{0j} & \text{if } \sum_i h_i I_{it} \leq \bar{h}, \\ \pi_{1j} & \text{if } \sum_i h_i I_{it} > \bar{h}. \end{cases}$$

Finally, I assume that the pandemic can be obliterated by an aggressive testing-and-tracing strategy if the global infection rate in the whole population goes below some threshold I_{min} .

3 Calibration of the MR-SIR model

Many of the parameters of this model remain highly uncertain, so caution should remain a cardinal virtue when interpreting its results. I consider a daily frequency for this analysis, and I calibrate the model on French data. There are $J = 3$ age classes, young (0-18), middle-aged (19-64), and senior (65+). I normalize the French population of 67 million people to unity. The size of the population in the different age classes is $N = (0.227, 0.568, 0.205)$. At date $t = 0$, we assume that 1% of the population is infected, uniformly across all age classes. I also assume the arrival of a vaccine 1.5 years after the beginning of the pandemic. The minimum rate of infection below which the virus can be obliterated in the population is assumed to be $I_{min} = 0.05\%$.

To simplify the calibration and given the limited available data, I assume that recovery rates, asymptomatic rates, contagion rates, and telework rates are assumed to be age-independent. The daily recovery rate $\gamma_i = \gamma = 1/18$ is assumed to be the same across age classes. This corresponds to the observation that infected people remain sick for 2 or 3 weeks on average. I also assume that the daily contagion rate under business-as-usual is 0.6, and that this rate goes to $\beta_q = 0$ for quarantined people, which is compatible with a basic reproduction number $R_0 = 2.5$ in the absence of policy. In the case of confinement, the contagion

rate is equal to $\beta_c = 0.1$. As I write this paper, the rate of asymptomatic cases is particularly difficult to calibrate. The Center for Evidence-Based Medicine has estimated this rate somewhere between 5% and 80%.⁴ He, Lau, Wu et al. (2020) found a 95% confidence interval of [25%, 69%] for the proportion of asymptomatic cases. More recently, the US Center for Disease Control and Prevention (CDC) has edicted 5 scenarios of the pandemic with two plausible levels of the rate of asymptomatic, 0.2 and 0.5, with a central assumption at 0.35.⁵ I assumed a $\kappa = 35\%$ rate of asymptomatic people. The social contact matrix across age classes has been estimated in France by Béraud, Kazmierczak, Beutels, Levy-Bruhl, Lenne, Mielcarek et al. (2015). We approximate their results by the following contact matrix:

$$\alpha_{..} = \begin{pmatrix} 2 & 0.5 & 0.25 \\ 0.5 & 1 & 0.25 \\ 0.25 & 0.25 & 0.5 \end{pmatrix} \quad (10)$$

Social interactions go down with age, within and across age classes. From this set of information, I can estimate the reproduction number at date 0 for the three age classes in the absence of any public policy by using equations (5) and (7). It yields

$$\mathfrak{R}_{y0} = 2.99 \quad \mathfrak{R}_{m0} = 2.77 \quad \mathfrak{R}_{s0} = 1.14. \quad (11)$$

Older people have a smaller reproduction number under the laissez-faire because of their reduced social interactions. The population-weighted reproduction number in the laissez-faire is thus $\bar{\mathfrak{R}}_0 = 2.48$, in line with the central scenario of the CDC. This is smaller than the often stated number $\mathfrak{R} = 3...5$ that has been estimated at the beginning of the pandemic. This reduction is made to account for the improved protective behavior coming from the understanding of the basic prevention efforts (no hand-shaking, mask bearing, ...) that have been observed more recently.

As in Acemoglu et al. (2020), I assume an efficiency rate of confinement of 75%, i.e., one confined person in four behaves without any lockdown constraint. I also assume that a proportion of 50% of confined people are able to telework, whereas the others are totally unable to generate any economic activity. This is in line with the estimation of a 5.8% of GDP loss in France during the first quarter of 2020.⁶ I assume an economic loss of a full confinement by a middle-aged person equaling $1/N_m$. This means that a 100% confinement of the middle-aged people without any telework capability during one year would generate a 100% GDP loss. In this calibration, telework halves this loss. I also assume that confining a young or a senior person yields no economic loss. This is in line with the fact that GDP does not take account of most contributions of these two age classes to the wealth of the nation.

The health care capacity in France is 6 hospital beds per 1,000 inhabitants. I assume that half of this capacity can be mobilized for the pandemic, so that $\bar{h} = 0.3$. The hospitalization rate among infected people is assumed to be 0.1%, 2% and 12%, respectively for the young, middle-aged and senior class. This has been estimated for France in mid-April 2020 by Salje et al. (2020). This study has also produced the following statistics about mortality rates: π_{0j} equals 0.001%, 0.15% and 3.65% respectively for $j \in \{y, m, s\}$. These mortality rates of covid are multiplied by a factor 5 if the hospital capacity is overwhelmed. There is no

⁴<https://www.cebm.net/covid-19/covid-19-what-proportion-are-asymptomatic/>

⁵<https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html>

⁶<https://www.insee.fr/fr/statistiques/4485632>

	Value	Description
γ	1/18	Daily recovery rate
β_q	0	Daily contagion rate of quarantined persons
β_c	0.1	Daily contagion rate of confined persons
β_f	0.6	Daily contagion rate of working persons
κ	0.35	Proportion of asymptomatic positives
θ	0.75	Efficiency rate of confinement
ξ	0.5	Proportion of telework
T_{vac}	1.5×365	Days before mass vaccination
I_{min}	0.05	Extinction threshold of the pandemic (in %)
\bar{h}	0.3	Hospital capacity (in %)
N	(22.7, 56.8, 20.5)	Distribution of population (in %)
π_0	(0.001, 0.15, 3.65)	Normal mortality rate (in %)
π_1	$5\pi_0$	Crisis mortality rate (in %)
h	(0.1, 2.0, 12.0)	Hospitalization rate (in %)
ϵ	(1, 1, 1)	Initial fraction of infection (in %)
w	(0, 176, 0)	Economic loss of confinement (in % of GDP/cap)
ℓ	(60, 40, 20)	Value of life lost (in years of GDP/cap)
α_1	(2, 0.5, 0.25)	Intensity of transmission from young
α_2	(0.5, 1, 0.25)	Intensity of transmission from adult
α_3	(0.25, 0.25, 0.5)	Intensity of transmission from senior
p	7	cost of mass PCR testing (in % of GDP/cap)

Table 2: Benchmark calibration of the MR-SIR model.

credible data to support this assumption, which is used to justify the classical "flattening the curve" policy.

It remains to calibrate the value of lives. I discuss this critical issue in Gollier (2020), remarking in particular that the absence of any democratic debate on this issue over the last five decades during which Western governments used a "value of statistical life" for policy evaluation is problematic. In this paper, I value a life lost at 60, 40 and 20 annual GDP/cap, respectively for the young, middle-aged and senior classes.⁷ For example, if one percent of the population dies, all aged 65 years or more, this has the same effect on welfare than a 20% reduction in annual GDP. This is equivalent to 60% if the 1% death toll is borne by the young generation. Robinson, Raich, Hammitt and O'Keeffe (2019) have surveyed studies using differentiated values of life for children compared to adult. They conclude that "the ratio of values for children to values for adults ranges from 0.6 to 2.9; however, most estimates are greater than 1.5." My benchmark calibration reflects this meta-observation. Balmford, Bateman, Bolt, Day and Ferrini (2019) used a contingent valuation method to estimate the value ratio of child to adult, with a value of 2.37 in the baseline case. This is also in line with my calibration. This benchmark calibration is summarized in Table 2.

Finally, I assume that no policy is implemented during the first 3 weeks of the pandemic.

⁷Greenstone and Nigam (2020) have a similar age-dependent valuation system.

4 Constant uniform confinement policies

In this section, I explore different no-test confinement policies in which the intensity of the confinement is uniform in the population. The simplest version of this family of lockdown strategies takes the form of imposing a uniform strength b of confinement until the threshold prevalence rate I_{min} is attained to obliterate the virus in the population. The outcome of such policies is described in Table 3 for different lockdown intensities $b \in [0, 0.9]$.




Figure 2: Dynamics of the pandemic in the laissez-faire strategy ($b = 0$).

Let me start with the laissez-faire policy ($b = 0$), whose outcome is described by the first line in Table 3 and by Figure 2. Only the symptomatic persons are quarantined under this policy. At the 22^d day of the pandemic, the reproduction number is down to $\bar{R}_{22} = 2.05$, because the share of susceptible persons has already been reduced to 92%. After two months of the pandemic, the wave of infection peaks, in which more than one-fourth of the young and middle-aged people are simultaneously sick. Because seniors have intrinsically less social contacts, they are much less infected. Because of this huge wave of infection, hospitals remain overcrowded during almost three months, implying a catastrophic death toll in which 8.5% and 0.6% of respectively the senior and the middle-aged people die. It takes 220 days for the pandemic to die out thanks to the herd immunity. At that time, respectively 94%, 92% and 53% of the young, middle-aged and senior people are immunized, meaning herd immunity is obtained with an immunization rate of 84%. The economic loss amounts to as little as 3.73% of annual GDP, due to the revenue loss of quarantined and deceased people during the period. But when valuing the lives lost, the total loss equals a dismal 51% of annual GDP.

This laissez-faire policy illustrates the necessity of "flattening the curve" with the objective to preserve some bed capacity in hospitals. This is done by imposing some lockdown in the

Covid Economics 24, 1 June 2020: 1-31

b	\mathfrak{R}_{22}	t^*	\bar{R}_{t^*}	D_y	D_m	D_s	W	L
0.0	2.05	220	84.0	0.001	0.32	1.74	3.73	51.42
0.1	1.78	237	79.5	0.001	0.28	1.43	4.33	44.26
0.2	1.53	260	73.3	0.001	0.22	1.02	5.37	34.67
0.3	1.29	291	64.9	0.000	0.06	0.27	7.13	15.05
0.4	1.08	328	52.7	0.000	0.05	0.20	10.87	16.91
0.5	0.88	349	37.9	0.000	0.04	0.13	16.13	20.25
0.6	0.71	304	24.8	0.000	0.02	0.08	18.82	21.44
0.7	0.55	228	17.1	0.000	0.02	0.06	17.29	19.08
0.8	0.41	176	13.2	0.000	0.01	0.05	15.35	16.75
0.9	0.29	146	11.1	0.000	0.01	0.04	14.20	15.39
1.0	0.19	127	9.9	0.000	0.01	0.03	13.53	14.61

Table 3: Outcome of the uniform confinement policy as a function of the strength b of the confinement. \mathfrak{R}_{22} is the reproduction number at the beginning of the confinement period (22^d day of the pandemic). t^* is the duration of the pandemic (in days), and \bar{R}_{t^*} is the overall immunity rate (in %) at that date. D_j is the share of population who dies in age class j (in % of whole population). W and L are respectively the economic loss and the total loss (in % of annual GDP).

Covid Economics 24, 1 June 2020: 1-31




Figure 3: "Flatten the curve" strategy: Dynamics of the pandemic with a $b = 30\%$ lockdown after 3 weeks of laissez-faire.




Figure 4: "Crush the curve" strategy: Dynamics of the pandemic with a $b = 90\%$ lockdown after 3 weeks of laissez-faire.

economy. However, doing so implies additional costs due to the confinement itself, but also to the increased time to eradicate the virus. Let me illustrate this in Figure 3 with the policy in which a $b = 30\%$ confinement is imposed after 3 weeks of laissez-faire. It now takes 291 days to eradicate the virus. Because social interactions are reduced, herd immunity is now obtained with a 65% immunization rate.⁸ Combining this with the fact that hospitals are not overcrowded implies that the rates of mortality is reduced to 1.3% for senior people and to 0.1% for the middle-aged people. The economic cost goes up to 7.13% of annual GDP due to the lockdown, but the total cost of the pandemic is now limited to 15.0% of annual GDP.

Should we go farther in the intensity of the lockdown once we have have flattened the curve enough to escape hospital overcrowding? The public planner faces here a dilemma. Strengthening the lockdown increases the economic cost due to the reduced labour and the increased duration of the lockdown. But it reduces the death toll. However, this marginal benefit is much reduced because the mortality rate is now much less sensitive to b due to the fact that the bed capacity constrained is slack. As a consequence, for $b \in [0.29, 0.57]$, the marginal economic cost of increasing the intensity of confinement is larger than its marginal benefit in terms of reduced lives lost.

There is a limit to this argument. A very strict confinement is able to suppress the virus in a short period of time. When the confinement intensity is larger than 57%, the sensitivity of the pandemic duration to the severity of the confinement is so large that any marginal increase in confinement *reduces* the economic cost of the policy. This is probably due to the fact that the reproduction number is a quadratic function of b , whereas the economic cost is linear in b . Because increasing b also reduces the death toll, the total cost of the policy

⁸Observe that the timing of the termination strategy is crucial here. I assume that the testing-and-tracing search for the 0.01% remaining contagious people in the population at day $t^* = 291$ can be quarantined before exiting from the lockdown.




Figure 5: Total cost as a function of the uniform intensity of confinement. The dashed curves correspond respectively to the economic cost (hump-shaped) and to the value of lives lost (decreasing).

is decreasing in the strength of the confinement for large b . Consider for example the case case of a confinement intensity of $b = 90\%$, as illustrated in Figure 4. With such a strong confinement, the reproduction number goes down to $\mathfrak{R}_{22} = 0.29$ on the day the confinement is implemented. This implies a fast reduction in the prevalence rate. The pandemic is terminated within 146 days, with a rate of immunity that is limited to 11.1% at that time. The economic loss is relatively large, but the global mortality rate is reduced to 0.05% of the whole population. This yields a total loss of 15.39%, similar to the total loss associated with the much milder confinement rate $b = 30\%$.

To sum up, one can consider two possible uniform confinement strategies, corresponding to two local minima for total cost in Figure 4. One can go for a "flatten-the-curve" strategy, or for a "suppression" strategy. The flatten-the-curve strategy is supported by a limited confinement intensity that is just large enough to escape hospital overcrowding. But it requires a large fraction of the population to be infected and to recover in order to build herd immunity. The suppression strategy is implemented with a very high confinement intensity to obliterate the virus in the population at the fastest possible speed. There is a set of bad policies between these two strategies in which economic costs and the mortality rate are larger. Hesitating between the flatten-the-curve strategy and the suppression strategy can be very costly. The current relatively high tolerance to breaking the strong confinement rules in some countries could bring us in this dismal outcome, with a longer duration of the pandemic, more lives lost, and larger economic losses. This is the curse of the hesitant confinement. Another illustration of the necessity of maintaining a time-consistent policy is documented in Figure 1

when the planner abandon the suppression strategy too early to switch to a flatten-the-curve strategy.

In this section, I examined constant strategies. This may not be optimal. For example, when crushing the curve of infection, actual policies around the world have tended to impose a strict confinement strategy at the beginning of the policy period to relax it once the rate of prevalence is considered to be low enough. Could this kind of time-varying strategy generates a better outcome compared to constant strategies? To answer this question, let me start with suppression strategies. Observe that under these strategies, the rate of change in the number of susceptible persons in each age class is almost zero for most of the strict confinement policy, as seen in Figure 4. In that case, the dynamic of the pandemic simplifies to

$$I_{i,t+1} - I_{i,t} = \left(\sum_{j=1}^J \hat{\beta}_{ijt} I_{j,t} \right) - \gamma_i I_{i,t}, \tag{12}$$

with $\hat{\beta}_{ijt} = \beta_{ijt} S_i$. In that case, as observed by Pollinger (2020), the dynamic of infection is purely exponential, with $I_{t+1} = A(x_t) I_t$ where I_t is the vector of age-specific numbers of infection, and $A(x_t)$ is a 3x3 matrix that is a time-independent function of the policy variables $x_t = (a_t, b_t)$. In that case, we obtain the following result.

Proposition 1. *Consider the set of suppression strategies in which the rate of change in the number of susceptible persons is almost zero. In this set, the cost-minimizing strategy is characterized by almost constant age-specific intensities of confinement and testing.*

Proof: In this proof, I denotes a 3x1 vector of age-specific numbers of infection, and x is a vector of policy variables that affect the 3x3 infection matrix $A(x)$ together with the convex per-period cost $w(x)$ of the policy. Consider all policies which imply the same age-specific rates of prevalence I_T at some date $T > 0$. I look for the dynamic strategy (x_1, \dots, x_T) which is feasible in the sense that the rate of infection at date T is I_T and that minimizes the cumulated economic cost:

$$\min_{x_1, \dots, x_T} \sum_{t=1}^T w(x_t) \quad \text{s.t.} \quad \left(\prod_{t=1}^T A(x_t) \right) I_0 = I_T. \tag{13}$$

Let A be the solution of the equation $A^T I_0 = I_{min}$. Obviously, the solution of the following static problem

$$x^* = \arg \min_x w(x) \quad \text{s.t.} \quad A(x) = A \tag{14}$$

characterizes the optimal solution $(x_1, \dots, x_T) = (x^*, \dots, x^*)$ of the dynamic problem (14). ■

This means that, at least for strategies aimed at crushing the curve of infection, limiting the search of policies to constant confinement and testing policies, potentially age-specific, is not restrictive. Of course, this result cannot be generalized to strategies aimed at flattening the curve in which the susceptible numbers S_{it} vary widely over time, so that matrix A is not a time-independent function of the policy. Also, when considering an objective that combines the economic loss and the value of lives lost, it may be optimal to start with a stricter confinement rule that can be relaxed over time. However, because under our calibration

the total loss is mostly entirely concentrated in the economic cost of the confinement, this effect of lives lost is expected to be marginal. The optimal crush-the-curve policy yields an almost constant confinement effort all along the duration of the pandemic. Proposition 1 can be interpreted as a special case of the model by Pollinger (2020) extended to MR-SIR. Pollinger characterizes the dynamic confinement strategy that minimizes total cost in the standard SIR model when the rate of change of S_t is almost zero. In his model, the dynamics of infection is such that $\partial I_t / \partial t = a(x)I_t - X(I_t)$, where $X(I)$ is an exogenous extraction technology from the pool of infected people. In my model, X equals zero everywhere except at $I = I_{min}$ where $X(I)$ equals I . The more comprehensive objective function and the richer set of extraction technologies considered by Pollinger yield much broader properties of the optimal solution. Under a weak condition on the extraction function X , Pollinger shows that the optimal intensity of confinement is decreasing with time.

5 Time-variable uniform confinement policies

I also examine non-discriminating confinement strategies that fluctuate over time. The lack of understanding of the dynamics of the pandemic and the difficulty to behave responsibly by a fraction of the population puts a strong pressure on the political system to exit from lockdown, or at least to relax the constraints, when the hospitalization rate goes down. I therefore explore the welfare impacts of "stop-and-go" strategies defined by four parameters: $\bar{b} \geq \underline{b}$ and $\bar{I} \geq \underline{I}$. After 3 weeks of an unconstrained circulation of the virus in the population, one starts with implementing a high confinement \bar{b} . As long as the prevalence rate I_t is larger than a minimum threshold \underline{I} , this strict rate of confinement is maintained. Once this threshold is attained, the lockdown is exited and a lower rate of confinement $\underline{b} \leq \bar{b}$ is imposed. It is maintained as long as the prevalence rate remains lower than \bar{I} . Once this threshold is attained, the larger confinement rate \bar{b} is reimposed, and so on until the prevalence rate attains I_{min} to obliterate the virus. Notice that uniform policies are special cases of stop-and-go policies in which \underline{b} and \bar{b} are the same.

In this section, I fix the trigger points (\underline{I}, \bar{I}) at respectively 20% and 80% of the bed capacity in the country. This guarantees that hospitals are never overcrowded. It also allows for reducing the intensity of the lockdown when a strict lockdown is unnecessary to flatten the curve.

Consider for example the stop-and-go policy with $\underline{b} = 0.2$ and $\bar{b} = 0.8$. Figure 6 describes the dynamic of the pandemic under this policy. It entails 3 waves of infection. After the first wave due to the laissez-faire of the first three weeks, a strict lockdown of 26 days is imposed that reduces the prevalence rate. Then, for 44 days, the intensity of the lockdown is weakened, which generates a second wave of the pandemic. This triggers a second strict lockdown that lasts for 26 days. Finally, the lockdown is weakened once again, which triggers a third wave. However, given the high rate of immunity prevailing in the population, no more strict lockdown need to be imposed because the hospitalization rate remains manageable. The pandemic lasts for 475 days. The economic loss is limited to 12.26%, and the mortality rate equals 0.1% and 1.2% respectively for the middle-aged and senior classes. The total loss is 19.36%.




Figure 6: "Stop-and-go" strategy: Dynamics of the pandemic with $\underline{b} = 20\%$ and $\bar{b} = 80\%$. I assume the trigger points $\underline{I} = 0.2\bar{h}$ and $\bar{I} = 0.8\bar{h}$.

In Table 4, I describe other stop-and-go strategies. For the sake of completeness, I added to this table the suppression policy $\underline{b} = \bar{b} = 0.9$ that I identified in the previous section as a potentially good suppression solution, yielding a total loss of 15.39% of annual GDP after a pandemic of 146 days. In terms of least-cost efficiency, it competes with the flatten-the-curve strategy consisting in maintaining a weak confinement $b = 30\%$ for 270 days. In Table 4, I examine another flatten-the-curve strategy in which that confinement intensity $\bar{b} = 0.3$ is maintained only during the period of time necessary to tip over the initial contagion wave. When the hospitalization rate goes back to 20% of bed capacity, a full relaxation of lockdown ($\underline{b} = 0$) is decided. This happens on day 165 of the pandemic, which dies out on day 386 without a second wave. It yields a record low total loss of 14.28% of annual GDP. This total cost is similar in level to the one described in the previous section, but its composition is different, with a smaller economic cost and a larger mortality rate. This two-stage strategy looks similar to the optimal uniform strategy described by Acemoglu et al. (2020). They show that the optimal uniform strategy under their calibration imposes a confinement that peaks at 50% of the population for a short period of time at the beginning of the pandemic, then stabilizes at around 30% for several months, and finally goes to zero when the infection rate vanishes. Notice however that the duration of their 30% confinement is close to one year, whereas I predict that half a year would be enough.

Although the crush-the-curve and flatten-the-curve strategies documented by the last two lines in Table 4 yield similar total costs, the suppression strategy suffers from an important comparative weakness due to the low rate of immunity that prevails at the end of the pandemic. If only 11.1% of the population got immunity in a world where the virus continues to circulate in other regions, implementing the suppression strategy is possible only if frontiers remain sealed as long as the virus survives elsewhere. The associated cost of this closure of the frontiers are not counted in my model.

\underline{b}	\bar{b}	t^*	\bar{R}_{t^*}	D_y	D_m	D_s	W	L
0.2	0.8	475	59.7	0.000	0.06	0.24	12.26	19.36
0.4	0.6	473	43.4	0.000	0.04	0.16	18.55	23.36
0.3	0.9	360	60.9	0.000	0.06	0.25	10.83	18.10
0.5	0.9	451	30.2	0.000	0.03	0.10	23.75	26.98
0.9	0.9	146	11.1	0.000	0.01	0.04	14.20	15.39
0.0	0.3	386	68.4	0.000	0.06	0.30	5.74	14.28

Table 4: Outcome of the stop-and-go confinement policy as a function of the strict and weak strengths (\underline{b}, \bar{b}) of the confinement. Notation and units are as in Table 3.

Covid Economics 24, 1 June 2020: 1-31




Figure 7: Total loss in stop-and-go strategies as a function of (\underline{b}, \bar{b}). I assume in this section the trigger points $\underline{I} = 0.2\bar{h}$ and $\bar{I} = 0.8\bar{h}$.




Figure 8: Optimal flatten-the-curve strategy: Dynamics of the pandemic under the Mico-Spiro-Weibull strategy of minimal daily confinement compatible with the health care capacity constraint.

In Figure 7, I describe the total cost as a function of (\underline{b}, \bar{b}) . The last two lines of Table 4 describe the two local minima of this bivariate function. This figure also illustrates the course of hesitation between these two polar strategies of suppression and flatten-the-curve. A convex combination of these two policies if for example $(\underline{b}, \bar{b}) = (0.5, 0.9)$, which implies a total loss of 26.98% as described in Table 4.

Let me conclude this section on time-variable uniform confinement strategies with an analysis of the optimal flatten-the-curve strategy examined by Mico, Spiro and Weibull (2020). These authors solve the problem of minimizing the economic cost of confinement in the standard SIR pandemic under a capacity constraint of the health care system. They show that it is optimal to impose no restriction as long as the capacity constraint is not binding, and to impose the minimal degree of confinement compatible with the capacity constraint when it is binding. I describe the outcome and the dynamics of the pandemic under my benchmark calibration of the MR-SIR model in Table 5 and Figure 5. The virus circulates freely in the population for the first 34 days of the pandemic before the health care capacity constraint binds. Then, a 45% confinement is imposed, that goes down gradually for 66 days until the population is fully deconfined. The virus is eradicated after 288 days. The economic cost is minimal at 4.43% of annual GDP, but the total cost is evaluated at 14.90%. It is useful to compare this policy to the best constant flatten-the-curve strategy with $b = 30\%$ for 291 days, as described in Table 3 and Figure 3. Because the confinement is weaker and shorter in the Mico-Spiro-Weibull strategy (they call it "filling-the-curve"), the economic cost is much smaller. But the rate of immunization is larger, in particular among the elderly whose immunity rate goes up to 49% (from 35%). Notice in particular that during the partial confinement period, the prevalence rate continues to grow in the senior class, which is compensated by a strong reduction of the prevalence rate in the other two age

t^*	\bar{R}_{t^*}	D_y	D_m	D_s	W	L
288	78.0	0.000	0.07	0.38	4.43	14.90

Table 5: Outcome of the Miclo-Spiro-Weibull confinement policy of minimal daily confinement compatible with the health care capacity constraint, as described in Figure 5.

classes. This unfortunate dynamics implies an increase in the number of lives lost. In terms of total cost, the two strategies yield a similar outcome.

6 Age-specific policies

Up to now in this paper, I only considered policies that are non-discriminated across age classes. There are at least two reasons why one should consider confinement policies that are age-sensitive. First, the case-fatality rate is extraordinarily different at different ages. Therefore, it may be interesting to expose more less vulnerable people if, for example, herd immunity is the final outcome of the pandemic. This suggests that the intensity of the confinement should be made increasing with age. The second reason is that the economic cost of confinement is also very different across age classes. This suggests that middle-aged people should receive priority in the exit of the lockdown. Both reasons justifies a strong sheltering of the seniors. But these two reasons offer contradictory arguments for which of the young or middle-aged people should exit from the lockdown first. Adults are more valuable to send back to work, but they are more vulnerable to the virus.

Table 6 explores different possible age-specific policies. As a follow-up of the stop-and-go strategies considered in the previous section, I now assume two age-specific confinement intensities, $\underline{b} = (0, 0, 1)$ and $\bar{b} = (b_y, b_m, 1)$, that are triggered when the hospitalization rate passes through one of the two thresholds $(\underline{I}, \bar{I}) = (0.2\bar{h}, 0.8\bar{h})$. Seniors remain 100% confined until the end of the pandemic. After three weeks of the free circulation of the virus, the young and the middle aged people are confined with an intensity b_y and b_m , respectively. Both young and middle-aged people are fully deconfined when the hospitalization rate reach 20% of the bed capacity. Consider first the case in which the middle-aged class is fully deconfined, whereas 80% of the young class remain in lockdown for 115 days before exiting it. The dynamic of covid-19 in this case is described in Figure 9. The pandemic is eradicated after 368 days, so that the seniors remain under full confinement for one year under this policy. The confinement of the junior and senior classes flattens the curve enough so that hospitals are not overwhelmed. Because the working class is never confined, the economic cost is limited to 3% of annual GDP. All in all, the total cost of the pandemic is contained at 8.58% of annual GDP.

One may find sheltering the seniors for more than one year as socially, morally or psychologically unacceptable. An alternative solution would be to exit seniors at the same time as the young generation, after 115 days of lockdown. That would increase the mortality of the

Covid Economics 24, 1 June 2020: 1-31




Figure 9: Age-specific policy with young confinement: Dynamics of the pandemic with $b_y = 80\%$ and $b_m = 0\%$. Seniors are 100% sheltered.




Figure 10: Age-specific policy with middle-aged confinement: Dynamics of the pandemic with $b_y = 0\%$ and $b_m = 0.18\%$. Seniors are 100% sheltered.

b_y	b_m	t^*	\bar{R}_{t^*}	D_y	D_m	D_s	W	L
1.00	0.00	401	68.6	0.000	0.08	0.13	3.01	8.58
0.80	0.00	368	68.1	0.000	0.07	0.13	3.00	8.52
0.25	0.10	256	70.8	0.000	0.07	0.13	3.62	9.11
0.00	0.18	250	71.3	0.000	0.07	0.13	4.16	9.59

Table 6: Outcomes of the age-specific confinement policy as a function of the young and middle-aged confinement strengths (b_y, b_m) of the confinement, assuming a 100% confinement of the seniors. Notation and units are as in Table 3.

seniors from 0.13% to 0.19% of the entire population (46,000 additional deaths in the case of France), and it would increase the total cost from 8.52% to 9.91%.

If the objective is to flatten the hospitalization curve and at the same time to build herd immunity, it would be intuitive to do it by exiting the young first from the lockdown, since their hospitalization rate is 20 times smaller than the middle-aged, and their mortality rate is 150 times smaller. The tradeoff comes from the fact that it is more expensive to confine middle-aged people. Consider for example the bottom line of Table 6, which documents the pandemic dynamics when young are not lockdown, and 18% of middle-aged people are confined for 112 days. This is enough to protect hospitals from the covid wave, as shown in Figure 10. Because the same fraction of adults have been infected at the end of the pandemic, the two strategies yield the same global death toll. But because of the additional cost of confining some workers, the total cost of the adult confinement is larger than when the young generation is initially confined to flatten the curve.

In this study, I discriminate the deconfinement strategy only on the base of age. But we know now that comorbidities have a large impact on the mortality rate too. For example, in New York state, 86% of the 5,489 reported COVID-19 deaths before 6 April 2020 involved at least one comorbidity, according to the state's department of health.⁹ Adding some of these comorbidities such as diabetes (37.3% of the New York deaths) and obesity in the individual characteristics of the discriminated deconfinement strategy could considerably reduce the death toll of age-specific strategies.

The main message of this section – together with all papers reviewed in the introduction dealing with this issue – is that there is a large social benefit of using an age-specific confinement strategy to flatten the curve. Whether this should be organized through a large confinement of the young or a weak confinement of the middle-age is not key. What is key is the necessity to protect the seniors, in particular during the strong wave of contagion happening in the first semester of the pandemic. Building herd immunity by sending the young to school and the middle-aged to work would be a demonstration of intergenerational solidarity towards the elderly people (Gollier (2020)).

⁹<https://www.the-hospitalist.org/hospitalist/article/220457/coronavirus-updates/comorbidities-rule-new-yorks-covid-19-deaths>

7 The economic value of mass testing

What is the value of the PCR test information in this framework? Suppose that it would require a person to be tested every week to detect his infection before becoming contagious. Let me estimate the unit cost of the test at 50 euros. This is a conservative estimation.¹⁰ Thus, testing a person on a weekly frequency by using an efficient PCR test would yield an annual cost of 2600 euros, or approximately 7% of annual GDP/cap for the calibration on French data. Thus, I calibrate $p = 0.07$ to estimate the economic cost W of the policy, as defined by equation (8).

Suppose that the current policy is a brutal suppression strategy with a 100% confinement of the population until the virus is obliterated in the population. This will take 106 days for the policy to attain this objective, with a total cost of 14.61% of annual GDP. Suppose that a testing capacity is obtained that allows for testing a fraction a of the population every week. Consider a policy in which a fraction a of the apparently susceptible population is tested every week-end. Among these tested people, those who are tested negative go back to work for a week, and those who are tested positive are quarantined. The fraction $b = 1 - a$ of \hat{S}_t which is not tested is confined for the week. Thus under this policy, only the tested negatives go back to work. Suppose that this policy is implemented until the virus is obliterated. Consider for example the case of a testing strategy $a = 50\%$. It will take 136 days to fight the pandemic under this policy. The total testing cost will equal 1.02% of GDP, whereas the confinement cost will amount to 7.56% of GDP, yielding a total cost of 9.70% when counting the value of lives lost. This total cost is reduced by one-third compared to the 100% confinement of the brutal confinement strategy. I describe the relationship between the testing intensity a and the total cost of the pandemic in Figure 11. If one could test the whole population every week, the pandemic could have been obliterated in 103 days, for a total cost of 2.80% of annual GDP.¹¹ This is the least-cost policy among the set of policies that I have examined in this paper. The bottom line of this comparative statics exercise is that developing a mass testing capacity could reduce the total cost of the suppression policy by a factor 5.

Developing a mass testing capacity is also useful to implement the flatten-the-curve strategy. Remember that the best uniform strategy to flatten the curve is to impose a stop-and-go confinement policy with $\underline{b} = 0\%$ and $\bar{b} = 30\%$, implying a total cost of 14.28% of annual GDP (see Table 4). This confinement intensity is just enough to suppress the risk of overwhelming the bed capacity in hospital. We can reproduce a similar dynamics by replacing the 30% uniform confinement by a 36% testing, with the same result to maintain the bed capacity constraint slack. This test-only strategy reduces the total cost of the pandemic to 12.18% of annual GDP.

A similar exercise can be made to examine the age-discriminated flatten-the-curve strategy in which the confinement intensity equals 0%, 18% and 100% respectively for the young, the middle-aged and the senior class. As documented in Table 6, this yields a total cost of 9.59%.

¹⁰This is the current price of covid PCR tests on the free market, which contains a mark-up. Increasing return to scale should also affect the cost negatively in the future.

¹¹A mass testing experiment has been put in place in the city of Vò in Lombardy. Although the city was one of the most important covid cluster in March, the city is now considered free of the virus.




Figure 11: Value of testing when the susceptible population is either tested or confined on a weekly basis ($b = 1 - a$). The total cost and the costs of testing, confinement and lives lost are measured in percents of annual GDP, as a function of the testing intensity a .

Consider an alternative strategy in which the young and the old generations continue to be respectively fully unlocked and fully sheltered. For the middle-aged people, suppose that 23% are tested and 77% go to work without testing. The total cost is reduced to 8.64%.

8 Sensitivity analysis

Given the high uncertainty that surrounds several parameters of the model, it is crucial to perform various robustness exercises. This work is summarized in Table 7. Different sets of rows document the effect of changing one parameter on three strategies: the suppression strategy $b = 0.9$; the strategy of uniform confinement that is just enough to flatten the curve given the health care capacity; and the best age-specific strategy that flattens the curve with a 100% confinement of the seniors and a partial confinement of the other two age classes. The first row summarizes the best policies with the set of parameters described in Table 2. Under this benchmark case, sheltering the seniors is the obvious winner, followed by the strategy of uniformly flattening the curve. The next row report the impact of uniformly doubling the value of life. This does not affect the structure of the policies under scrutiny, but it changes the weights in the objective function. Increasing the value of life makes more attractive policies that preserve more lives, i.e., suppression and sheltering the old. Suppression now clearly dominates flattening the curve under this alternative calibration. I also report on the impact of reducing the value of lives lost for the elderly, considering the fact that many covid fatalities have other co-morbidity factors and a relatively limited remaining life expectancy. Reducing the value of lives lost of people aged 65+ years to 10 GDP/cap has the effect to make the fatten-the-curve policy much more favorable.

I could have reported other robustness checks in which the relative values of lives lost

across ages are altered compared to the benchmark. For example, a standard practice is to use a value of statistical life that is independent of age (US-EPA (2010), Greenstone and Nigam (2020)). Reducing the value of life of the young relative to the senior reinforces the recommendation to shelter the old generation under age-specific confinement strategies. In fact, there is no reasonable valuation system that could reverse this result. Gollier (2020) clarifies this point in a simple static analysis.

In the next row, I examine the impact of raising the rate of asymptomatic from 0.35 to 0.45. More asymptomatic cases implies more contagion since it reduces the number of infected people that are quarantined in the absence of test. This raises the mortality rate. If one follows the suppression policy, the economic cost also increases because of the increased duration of the pandemic. Under the flatten-the-curve strategy, the rate of confinement must be increased from $\bar{b} = 0.30$ to 0.51 in order to preserve hospitals. This also increases the economic cost. Under the old-sheltering strategy, the increased rate of asymptomatic cases forces the planner to confine all young people and 30% of the working age class.

In the benchmark calibration, one in four confined persons behaves as if not confined. In the next row, I increase this inefficiency of the lockdown to $2/4$. Of course, this deteriorates the attractiveness of the suppression strategy, in which all costs basically double. Because the uniform flattening of the curve relies less on confinement, it is much less affected by the increased lockdown inefficiency. In fact, the misbehavior observed in the population can be compensated by an increased intensity of confinement from $\bar{b} = 0.3$ to 0.43, thereby marginally increasing the economic cost of the pandemic. Concerning the strategy of sheltering the seniors, the reduced efficiency of the confinement also necessitates increasing the intensity of the lockdown for the younger generations. To escape hospital overcrowding, one needs to fully confine the young class, together with some people from the working age class. This has a sizeable impact on the total cost.

In the benchmark calibration, the intensity of social interactions goes down with age, as expressed by equation (10). Suppose alternatively that senior people behaves as the middle-aged class:

$$\alpha_{..} = \begin{pmatrix} 2 & 0.5 & 0.5 \\ 0.5 & 1 & 0.5 \\ 0.5 & 0.5 & 1 \end{pmatrix} \quad (15)$$

Obviously, this is no good news, as this raises the reproduction numbers. If applying the suppression policy, the duration of the pandemic will be longer, and the mortality rate will be larger, in particular in the senior class. However, because the virus is contained in the early stage of the pandemic the global impact of these increased social interactions remains marginal. We cannot say the same thing when flattening the curve. Because senior people are more integrated in the population, they contribute much more to the building of herd immunity. This massively increases their mortality rate, which goes up to 2.2% of their age class. Notice that the increased contagion implies that the uniform confinement rate must be increased to $\bar{b} = 0.41$ to flatten the curve. This implies that flattening the curve is now dominated by the strategy of suppression. Notice also that this increased integration of the seniors makes it more difficult to shelter them under the third strategy.

Next, I examine the effect of a 50% increase in the mortality risk of the senior people. Because the three policies that I examine in this table are rather protective of the old generation, the effect of this increased mortality risk is relatively limited in the number of lives lost

Parametrization	Policy	Duration (days)	Mortality (% of pop)	Econ cost (% of GDP)	Total cost (% of GDP)
Baseline	Crush curve	146	0.049	14.20	15.39
	Flatten curve	386	0.363	5.74	14.28
	Old sheltered	368	0.200	3.00	8.52
$\ell = (120, 80, 40)$ (\uparrow value life)	Crush curve	146	0.049	14.20	16.59
	Flatten curve	386	0.363	5.74	22.82
	Old sheltered	368	0.200	3.00	14.03
$\ell = (60, 40, 10)$ (\downarrow value life)	Crush curve	146	0.049	14.20	15.01
	Flatten curve	386	0.363	5.74	11.29
	Old sheltered	368	0.200	3.00	7.26
$\kappa = 0.45$ (\uparrow asympt.)	Crush curve	173	0.097	15.31	17.69
	Flatten curve	433	0.429	6.92	16.91
	Old sheltered	366	0.236	4.54	10.76
$\theta = 0.5$ (\downarrow lockdown eff.)	Crush curve	304	0.107	27.76	30.38
	Flatten curve	367	0.368	7.01	15.68
	Old sheltered	398	0.272	4.52	11.37
$\alpha_3 = (0.5, 0.5, 1)$ (s behaves as m)	Crush curve	155	0.086	14.91	16.88
	Flatten curve	498	0.512	7.44	18.99
	Old sheltered	429	0.293	3.52	10.84
$\pi_{03} = 5.47\%$ (\uparrow s mortality)	Crush curve	146	0.068	14.20	15.78
	Flatten curve	386	0.512	5.74	17.26
	Old sheltered	368	0.263	3.00	9.76
$w_y = w_s = 0.4$ (y, s create value)	Crush curve	146	0.049	16.72	17.92
	Flatten curve	386	0.363	6.87	15.41
	Old sheltered	368	0.200	7.45	12.96

Table 7: Sensitivity analysis.

and economic terms. Finally, similarly, increasing the economic cost of confinement of the young and seniors generations from 0 to 0.4, i.e., approximately one-fourth of the economic cost of the middle-age confinement, does not change much the results.

Clearly, this sensitivity analysis shows that the recommendation to shelter the seniors during the peak of infection is robust to changes in the parameters of the MR-SIR model. On the contrary, this also shows that there is no clear winner between the crush-the-curve strategy and the flatten-the-curve strategy.

9 Concluding remarks

When only 1% of the population is infected, a strong confinement of the whole population is a rather inefficient way to fight the virus. To reduce the contagion by one person, one hundred persons are confined, yielding a large economic impact. Obviously, our ability to test people in order to deconfine the negatives would be a much better strategy.¹²

In the absence of a testing capacity, one possible strategy is to suppress the virus by a strong degree of confinement during a few months. In this context, increasing the strength of the confinement has a large effect on reducing the duration of the lockdown. It happens that the net effect of strengthening the confinement is to reduce the total cost of the pandemic. Therefore, if the objective is to suppress the virus as soon as possible, it makes no sense to reduce the confinement before the full eradication of the virus in the population. Because this strategy implies a very low rate of immunization until the end of the pandemic, stopping short of full eradication exposes the country to a restart almost from the beginning of the pandemic. Following this strategy thus requires a lot of time-consistency and a strong political resistance to the lobbies and to some increasingly impatient citizens.

Most SIR modelers explore an alternative strategy whose objective is to reach herd immunity in a relatively long horizon. Reaching herd immunity without controlling the speed of the pandemic exposes hospitals to a tsunami of infections, implying a terrible death toll. This is why reaching herd immunity may be an efficient strategy only if the infection curve is flattened via a weak confinement at the beginning of the immunization process. The economic impact of the policy is small, but the mortality is severe. Depending upon how Society value lives lost, this strategy can dominate the suppression strategy.

When implementing a non-discriminatory confinement policy, the senior class bears most of the burden of deaths, in spite of the lower intensity of their social interactions. This is because of the extremely large case-fatality proportion faced by people aged 70 and older. An efficient intergenerational solidarity would be to ask the younger generations to build Society's herd immunity by living their life, participate to the economy, get infected and recover from covid-19. More vulnerable people should be sheltered during this high-contagion period of the pandemic. This sheltering of the vulnerable increases the mortality of the young and of the middle-aged. This is the price they could pay to express their solidarity to the elderly. The mortality risk differential is so big that targeting herd immunity, flattening the curve and sheltering the old generation is a no brainer, with a total cost far smaller than any uniform

¹²If our PCR test capacity is limited, Gollier and Gossner (2020) suggested to use a standard testing protocol in which several individual samples are pooled and tested with a single test. If the objective is to maximize the number of people to unlock, the optimal group size is approximately equal to the inverse of the rate of prevalence.

policy. Confine a fraction of the young to flatten the curve, get the working age class back to work to reduce the economic cost, and shelter the senior class to reduce the death toll. Acemoglu et al. (2020) make a very similar final recommendation.

Bibliography

Acemoglu, D., V. Chernozhukov, Ivan Werning and Michael Whinston, (2020), A multi-risk SIR model with optimally targeted lockdown, NBER WP 27102.

Alvarez, F., D. Argente and F. Lippi, (2020), A simple planning problem for COVID-19 lockdown, University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2020-34.

Balmford, B., I.J. Bateman, K. Bolt, B. Day, S. Ferrini, (2019), The value of statistical life for adults and children: Comparisons of the contingent valuation and chained approaches, *Resource and Energy Economics* 57, 68-84.

Béraud G, S. Kazmierczak, P. Beutels, D. Levy-Bruhl, X. Lenne, N. Mielcarek et al., (2015), The French connection: The first large population-based contact survey in France relevant for the spread of infectious diseases, *PLoS ONE* 10,

Brotherhood, L., P. Kircher, C. Santos and M. Tertilt, (2020), An economic model of covid-19 epidemic: The importance of testing and age-specific policies, mimeo.

Drèze, J., (1962), L'utilité sociale d'une vie humaine, *Revue Francaise de Recherche Opérationnelle*, 23, 93-118.

Favero, C., A. Ichino and A. Rustichini, (2020), Restarting the economy while saving lives under covid-19, WP Bocconi University.

Ferguson, N.M. et al., (2020), Impact of non-pharmaceutical interventions (npis) to reduce covid-19 mortality and healthcare demand, *CEPR Covid Economics* 2, 60-66.

Fischer, C., (2020), External costs and benefits of policies to address COVID-19, mimeo.

Gollier, C., (2020), If herd immunity is the objective, on whom should it be built?, *Covid Economics* 16, 98-114.

Gollier, C., and O. Gossner, (2020), Group testing against Covid-19, *Covid Economics* 1, 32-42.

Greenstone, M., and V. Nigam, (2020), Does social distancing matter?, *Covid Economics* 7, 1-22 (May 2020).

He, X., E.H.Y. Lau, P. Wu et al., (2020), Temporal dynamics in viral shedding and transmissibility of COVID-19, *Nature Medicine* 26, 672-675. <https://doi.org/10.1038/s41591-020-0869-5>

Jones-Lee M. (1974) The value of changes in the probability of death and injury, *Journal of*

Political Economy 82, 835-849.

Miclo, L., D. Spiro and J. Weibull, (2020), Optimal epidemic suppression under an ICU constraint, mimeo, TSE.

Murphy, K. M. and R. H. Topel, (2006), The value of health and longevity, *Journal of Political Economy* 114, 871-904.

Pindyck, R.S., (2020), Covid-19 and the welfare effects of reducing contagion, mimeo, MIT.

Pollinger, S., (2020), Optimal tracing and social distancing policies to suppress COVID-19, mimeo, TSE.

Quinet, E., (2013), L'évaluation socioéconomique des investissements publics, *Commissariat Général à la Stratégie et à la Prospective*, Paris.

Robinson, L.A., (2007), How US government agencies value mortality risk reductions, *Review of Environmental Economics and Policy* 1, 283-299.

Robinson, L.A., W.J. Raich, J.K. Hammitt and L. O'Keeffe, (2019), Valuing children's fatality risk reduction, *Journal of Benefit Cost Analysis* 10, 156-177.

Salje, H., C. Tran Kiem, N. Lefrancq, N. Courtejoie, P. Bosetti, et al., (2020), Estimating the burden of SARS-CoV-2 in France. [pasteur-02548181](https://doi.org/10.1101/2020.05.05.20088181)

Schelling, T., (1968), The life you save may be your own, in "Problems in public expenditure analysis" (Chase, S.B., ed.), Washington, D.C., The Brookings Institution, 127-162.

Shepard, D.S., and R.J. Zeckhauser, (1984), Survival versus Consumption, *Management Science* 30, 423-39.

US-EPA, (2010), Valuing mortality risk reductions for environmental policy: A white paper, SAB-EEAC Review Report.

Viscusi, W.K., (2009), The devaluation of life, *Regulation & Governance* 3, 103-127.

Wilder, B., M. Charpignon, J. Killian, O. Han-Ching, A. Mate, S. Jabbari, A. Perrault, A. Desai, M. Tambe and M. Majumder, (2020), The role of age distribution and family structure on COVID-19 dynamics: A preliminary modeling assessment for Hubei and Lombardy, mimeo.

The economic effects of Covid-19 containment measures¹

Pragyan Deb,² Davide Furceri,³ Jonathan D. Ostry⁴
and Nour Tawk⁵

Date submitted: 23 May 2020; Date accepted: 25 May 2020

Many countries around the world have implemented stringent containment measures to halt the spread of the 2019 coronavirus disease (Covid-19) and limit the number of fatalities. Though crucial to slow the course of the pandemic, these measures entail large short-term economic costs. This paper tries to quantify these effects using daily data on real-time containment measures implemented by countries around the world as well as daily indicators of economic activity such as Nitrogen Dioxide (NO₂) emissions, international and domestic flights, energy consumption, maritime trade, and retail mobility indices. Results suggest that containment measures have had, on average, a very large impact on economic activity—equivalent to a loss of about 15 percent in industrial production over a 30-day period following the implementation of containment measures. Using a novel database on discretionary fiscal and monetary policy measures implemented by countries in response to the crisis, we find that these policy measures have been effective in mitigating some of these costs. Finally, we find that among different types of containment measures, while stay-at-home requirements and workplace closures are the most effective in curbing both infections and deaths, they are also those associated with the largest economic costs.

- 1 The views expressed in this paper are those of the authors and do not necessarily represent those of the IMF or its member countries. We would like to thank Naihan Yang for excellent research assistance.
- 2 Economist, International Monetary Fund.
- 3 Deputy Division Chief, International Monetary Fund and University of Palermo.
- 4 Deputy Director, International Monetary Fund and CEPR.
- 5 Economist, International Monetary Fund.

Copyright: Pragyan Deb, Davide Furceri, Jonathan D. Ostry and Nour Tawk

I. INTRODUCTION

Many countries around the world have enacted stringent containment measures and non-pharmaceutical interventions (NPIs) to halt the spread of the coronavirus and limit the number of fatalities, and in a bid avoid overwhelming the medical system and to buy time while effective treatments and vaccines are developed. Interventions have ranged from improved diagnostic testing and contact tracing, isolation and quarantine for infected people, and importantly, measures aimed at reducing mobility and creating social distancing (containment measures, hereafter).

Empirical evidence from China (Kraemer et al. 2020; Chinazzi et al. 2020; H. Tian et al. 2020) as well as for other countries in the world (Deb et al. 2020) suggest that these measures have been effective in flattening the pandemic “curve” and significantly reducing the number of fatalities. In particular, they find that countries that have put in place stringent measures, for example those implemented in countries such as China and Italy, as well as early intervention, such as in New Zealand and Vietnam, may have reduced the number of confirmed cases and deaths by more than 200 percent relative to the underlying country-specific path in the absence of interventions.

However, while these measures have contributed to saving lives, and have therefore provided the foundation for a stronger medium-term growth (see Barro et al. 2020), they have led to unprecedented economic losses in the short term. Quantifying these short-term economic effects and whether they vary across types of containment measure is of paramount importance for many policymakers around the world facing a painful short-term tradeoff between normalizing economic activity and minimizing health risks.

This paper tries to address these issues empirically. In particular, the paper has three main goals. The first is quantify the average economic effect—across countries and measures—of

containment measures. For this purpose, we assemble daily data on real-time containment measures implemented by countries around the world as well as a unique database containing daily data on several indicators of economic activity: Nitrogen Dioxide (NO₂) emissions—as explained in the next section, our main variable of interest; international and domestic flights; energy consumption; maritime trade; and retail mobility indices. Establishing the causal effect on economic activity is difficult, because the decision of countries to implement containment measures crucially depends on the evolution of the virus, which in turn may affect mobility and economic activity (Maloney and Taskin 2020). This implies that addressing causality requires the researcher to effectively control for this endogenous response which would otherwise bias estimates of the effect of containment measures. The use of daily data allows us to address this issue by controlling for the change in the number of infected cases (and deaths) occurring a day before the implementation of containment measures, as well as for lagged changes in daily economic indicators. Indeed, given lags in the implementation of interventions at daily frequency, this approach effectively controls for the endogenous response of containment measures to the spread of the virus. To further account for expectations about the country-specific evolution of the pandemic, we also control for country-specific linear time trends.¹ The results of this analysis suggest that containment measures have had, on average, a very large impact on economic activity—equivalent to a loss of about 15 percent in industrial production over the 30-day period following the implementation of the containment measure.

The second goal of the paper is to examine whether fiscal and monetary measures implemented by many governments and central banks around the world have been effective so far

¹ A remaining concern is that containment measures were announced before being implemented and, therefore, were anticipated. This may have resulted in reduced mobility ahead of the implementation of the containment measures. In the section on robustness checks, we show that the results are practically unchanged when we control for changes in mobility.

in mitigating the negative effects of containment measures. To answer this question, we use data provided by the *IMF Policy Tracker* which compiles discretionary fiscal and monetary measures implemented in response to COVID-19. The results suggest that macroeconomic stimulus deployed so far has been effective, with the negative effect of containment measures being much larger—equivalent to a loss in industrial production of about 22 percent—in countries that have provided limited fiscal and monetary policy stimulus.

Finally, the third goal of the paper is to examine which types of containment measure have resulted in larger economic costs and short-term tradeoffs between minimizing health risks and economic losses. For this purpose, we analyze the economic and virus transmission effects of the following containment measures: (i) school closures; (ii) workplace closures; (iii) cancellation of public events; (iv) restrictions on size of gatherings; (v) closures of public transport; (vi) stay-at-home orders; (vii) restrictions on internal movement; (viii) restrictions on international travel. While the results should be treated with caution since many of these measures were often introduced simultaneously as part of the country's response to limit the spread of the virus, preliminary evidence suggest that stay-at-home requirements and workplace closures, the two containment measures which are most effective in curbing both infections and deaths, are also the costliest in economic terms. In contrast, school closures and restrictions on international travel are the least costly but still successful in lowering COVID-19 infections, though less effective in reducing fatalities.

This paper contributes to two strands of literature. The first is on the use high-frequency daily economic indicators to monitor economic activity. Kumar and Muhuri (2019) employ a transfer learning-based approach to predict per capita GDP of a country using CO₂ emissions. Marjanovic et al. (2016) uses Extreme Machine Learning (EML) and Genetic Programming (GP)

to predict GDP based on CO₂ emissions. Other notable empirical approaches using novel high-frequency indicators include Small et al. (2010), who show that stable night lights data can act as a proxy for urban development. Cerdeiro, Komaromi, Lui and Saeed (2020) use raw Automatic Identification System (AIS) signals emitted by global vessel fleets to create real-time indicators of world seaborne trade. Glazer, Kim and Luca (2017) use Yelp data to measure economic activity in close to real time, and a granular level, across different US states and counties.

The second strand of literature this paper contributes to is on the potential economic effect of COVID-19 and containment measures, including based on past pandemic episodes. Barro, Ursua and Weng (2020) studied the effects of non-pharmaceutical interventions (NPIs) such as school closings, prohibition on public gathering and quarantine/isolation on death rates in the United States during the 1919 pandemic. They find that while NPIs have a significant effect on peak death rates, they had a more limited impact on the cumulative number of deaths, possibly because they were not enforced for long enough. They also find that the macroeconomic effects of the pandemic were quite large, with the economy of a typical state contracting by around 6 percent. Ma et al. (2020) draw lessons for the COVID-19 pandemic from examining the immediate and bounce-back effects of six past health crises: the 1968 Flu (also referred to as the Hong Kong Flu), SARS (2003), H1N1 (2009), MERS (2012), Ebola (2014), and Zika (2016). They find that real GDP is 2.4 percent lower the year of the outbreak in countries affected relative to those unaffected, and that it remains below its pre-shock levels for five years after the crisis despite bouncing back. They also find that fiscal policy plays an important role in mitigating the impact of a health crisis, with the negative impact on GDP being reduced in countries that deployed large first-year responses in government spending and health care. Coibion et al. (2020) use data from customized surveys from over 10,000 respondents to estimate the impact of COVID-19 on households'

spending and macroeconomic expectations in the United States. They find that aggregate consumer spending has declined substantially so far, especially in travel and clothing. They also find that households living in countries which enforced lockdowns earlier expect a higher unemployment rate over the next three to five years.²

The remainder of the paper is structured as follows. Section II describes stylized facts, data and econometric methodology. Section III presents our results on the effect of containment measures, and how these effects vary across countries, depending on fiscal and monetary measures deployed since the pandemic outbreak, and by type of containment measure. The last section concludes.

II. STYLIZED FACTS, DATA AND METHODOLOGY

A. Data

We assemble a comprehensive daily database of economic indicators, containment measures and COVID-19 infections and deaths. Table 1 provides the country coverage of each variable and key summary statistics.

Economic data

Nitrogen Dioxide (NO_2) emissions. We use daily data on Nitrogen Dioxide (NO_2) emissions from the Air Quality Open Data Platform of the World Air Quality Index (WAQI). Data available on WAQI is collected from countries' respective Environmental Protection Agencies (EPA). The database for NO_2 levels covers 62 countries in total, 57 of which are used for our analysis, with

² For theoretical studies examining the effect of containment measures on economic activity see, for example, Eichenbaum, Rebelo and Trabandt (2020) and references therein.

coverage beginning from January 1, 2020.³ The data is based on the median emissions reported by city-specific stations which are updated three times a day. Data on NO₂ pollution is provided in US EPA standards, which mandates that units of measure for NO₂ emissions be parts per billion (ppb).

We use NO₂ emissions as our main variable of interest for the empirical work in this paper, for three reasons: (i) NO₂ emissions are strongly correlated to lower-frequency economic variables which are used in macro-economic analysis, such as industrial production (see next section); (ii) emission levels can be directly linked to overall economic activity, and are not indicative of activity for specific sectors only (as flights would be for tourism, for instance); (iii) data are available on a daily frequency, covering a relatively large sample of 57 countries. That said, we present the effect of containment measures on the following set of daily indicators:

Flights. Flight data are collected from FlightRadar24, which provides real-time information on worldwide flights from several data sources, including automatic dependent surveillance-broadcast (ADS-B), (Multilateration) MLAT and radar data.⁴ The database covers international and domestic inbound and outbound flights data for over 200 countries, 84 of which are used in our analysis. Data coverage is on a daily frequency and begins on January 1, 2020. Data for total flights is calculated by summing daily domestic and international flights.

Energy consumption. We use daily data on energy consumption for 35 countries in Europe from ENTSO-E's transparency platform. The platform provides hourly total load of electricity

³ *COVID-19 Worldwide Air Quality Data*. Accessed May 7, 2020. <https://aqicn.org/data-platform/COVID-19/report/>

⁴ <https://www.flightradar24.com/how-it-works>

generated per market time unit by plants covered by Transmission System Operators (TSO) and Distribution System Operators (DSO) networks. Coverage in our sample begins from January 1, 2020.

Maritime imports and exports indices. For maritime import and export indices, we use data from Cerdeiro, Komaromi, Lui and Saeed (2020), who build real-time indicators of world seaborne trade using raw Automatic Identification System (AIS) signals emitted by global vessel fleets through their transponders. They use machine-learning techniques to transform AIS data, which contain information on vessels' speed, location, draught, etc., into import and export maritime indices. Their database produces import and export indices for 22 countries. Data coverage begins on January 1, 2020.

Retail and transit-station mobility. We collect data on retail and transit-station mobility from Google Mobility Reports. The reports provide daily data by country and highlight the percent change in visits to places related to retail activity (restaurants, cafes, shopping centers, movie theaters, museums, and libraries), or public transport (subways, buses, train stations etc.). The data for each day is reported as the change relative to a baseline value for that corresponding day of the week, and the baseline is calculated as the median value for that corresponding day of the week, during the 5-week period between January 3rd and February 6th, 2020. Daily data are available for 73 countries in our dataset, with coverage beginning from February 15, 2020.

Containment measures

We use data from Oxford's COVID-19 Government Response Tracker⁵ (OxCGRT) for containment measures. OxCGRT collects information on government policy responses across eight dimensions, namely: (i) school closures; (ii) workplace closures; (iii) public event cancellations; (iv) gathering restrictions; (v) public transportation closures; (vi) stay-at-home orders; (vii) restrictions on internal movement; and (viii) international travel bans. The database scores the stringency of each measure ordinally, for example, depending on whether the measure is a recommendation or a requirement and whether it is targeted or nation-wide. We normalize each measure to range between 0 and 1 to make them comparable. In addition, we compute and aggregate a Stringency Index as the average of the sub-indices, again normalized to range between 0 and 1. The data start on January 1, 2020 and cover 151 countries/regions.

Fiscal and monetary policy measures

Data on fiscal stimulus (announced and implemented fiscal packages in percent of GDP) and monetary policy actions (change in policy rates) implemented in response to the COVID-19 pandemic are sourced from the IMF policy tracker. The survey is distributed to country authorities to provide information on policy measures implemented since the beginning of the pandemic, ranging from external, financial, fiscal, monetary, and other policy streams. Responses are collected and updated on a weekly basis. The coverage includes 195 IMF member countries.

⁵ "Coronavirus Government Response Tracker." Blavatnik School of Government. Accessed May 7, 2020. <https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker>.

COVID-19 infections and deaths

Data on infections and deaths are collected from the COVID-19 Dashboard from the Coronavirus Resource Center of Johns Hopkins University.⁶ Coverage begins from January 22, 2020. It provides the location and number of confirmed cases, deaths, and recoveries for 208 affected countries and regions.

⁶ COVID-19 Map, JHU Coronavirus Resource Center, Accessed May 7, 2020 <https://coronavirus.jhu.edu/map.html>.

Table 1. Summary Statistics

	Obs.	Mean	Min	Max	Std. Dev.	Source	Starting Date	N. of countries
						Air Quality Open Data Platform		
NO ₂ emissions (log)	9,170	2.0	-0.9	4.4	0.7		1-Jan-20	62
Total Flights (log)	29,997	3.4	0.0	10.8	2.1	FlightRadar24	1-Jan-20	217
Retail Mobility (%)	13,456	-11.7	-58.6	2.6	13.3	Google Mobility Index	15-Feb-20	132
Transit Station Mobility (%)	13,350	-12.2	-57.8	3.3	13.5	Google Mobility Index	15-Feb-20	131
						Cerdeiro, Komaromi, Lui and Saeed (2020)		
Maritime Import Index (log)	2,420	4.6	3.83	4.9	0.12		1-Jan-20	22
						Cerdeiro, Komaromi, Lui and Saeed (2020)		
Maritime Import Index (log)	2,310	4.6	4.21	5.1	0.12		1-Jan-20	22
Energy Consumption (log)	4,785	12.1	3.62	15.6	1.5	ENTSO-E	1-Jan-20	35
						Coronavirus Resource Center of JHU		
Confirmed Cases (log)	16,996	5.3	-0.9	14.3	3.0		21-Jan-20	208
						Coronavirus Resource Center of JHU		
Confirmed Deaths (log)	11,379	3.2	-1.9	11.5	2.5		22-Jan-20	176
Stringency of Measures Index (%)	24,626	0.4	0	1	0.4	OxCGRT.	1-Jan-20	158
Fiscal Stimulus (% of GDP)	14,290	3.3	0	12.1	3.1	IMF Policy Tracker	1-Jan-20	97
Policy rate cuts (bps)	25,377	76.2	0	1000	118.8	IMF Policy Tracker	1-Jan-20	172

B. Stylized Facts

In order to curb COVID-19 infection and fatality rates, governments worldwide put in place containment measures which have ranged from school closures and cancellations of public gatherings, to restrictions on internal movement and stay-at-home orders. The stringency of such measures effectively led to shutdowns of production, manufacturing, and transportation sectors, and to lockdowns of many cities for prolonged periods of time. This section provides a first look at the data to examine whether containment measures have played a role in the observed decline in economic activity, proxied by NO_2 emissions. To do so, we examine the levels of NO_2 emissions in four cities before and after the implementation of (national) containment measures to fight the COVID-19 outbreak: Wuhan (China), Rome (Italy), New York (United States), and Stockholm (Sweden).

Figure 1 presents the pattern of NO_2 emission (left scale) together with the evolution of the stringency indicator (right scale). It shows that emissions significantly declined in these four cities after containment measures have been put in place. In Wuhan, a dramatic fall in NO_2 levels coincided with the enforcement of the cordon sanitaire on January 22, 2020, and the implementation of the stringent containment measures in the days that followed. Measures which were put in place within a week of the cordon sanitaire included restrictions on internal movements and gatherings, stay-at-home orders, closures of public transport, and cancellations of public events. By the end of March, emissions were back on the rise, as public transport areas reopened, and restrictions on internal movement and stay-at-home requirements were relaxed (Figure 1A).

In Rome, the pace of decline in NO_2 emissions picked up significantly towards end-February (Figure 1B) as a result of containment measures introduced on February 23, 2020. Measures implemented were highly restrictive of internal movement, and, as in Wuhan, included


school and workplace closures, cancellation of public gatherings, restrictions on internal movements and gatherings, and stay-at-home orders. NO_2 levels fell even further following the official lockdown of Italy on March 9, and closures of public transport. There is a noticeable pick-up in NO_2 emissions in early May, a few days after four containment measures were relaxed (workplace closures, stay-at-home orders, and restrictions on internal movement and international travel), and one was lifted (closures of public transport).

In New York, containment measures were only tightened drastically by end-March. Initially, containment measures entailed restrictions on international travel, school closures and cancellations of public events. As the outbreak evolved, restrictions on internal movements and the size of gatherings were put in place. Closure of workplaces was the last type of containment measure to be enforced. Consequently, NO_2 emissions fell at a gradual pace and plateaued at their lowest levels only after all measures were enforced (Figure 1C).


Sweden's unique response to the COVID-19 pandemic has entailed limited containment measures. To-date, five containment measures have been implemented in the following order: restrictions on gatherings; school closures; restrictions on international travel; workplace closures; and restrictions on internal movement. However, with the exception of international travel restrictions, the other four containment measures implemented rank lowest in stringency: schools for younger children are open, bans on public gatherings are for crowds of over fifty people, and restaurants, cafes and pubs remain operational, but must enforce social distancing. Because of less stringent containment measures, it is perhaps unsurprising that NO_2 emissions have not declined significantly in Stockholm (Figure 1D). Summarizing, preliminary evidence seems to suggest that containment measures have led to a decline in economic activity, as reflected in lower emissions. The next section checks whether this descriptive evidence holds up to more formal tests.

Figure 1: Evolution of NO₂ emissions, selected cities


Panel A. NO₂ emissions, Wuhan (parts per billion (ppb))




Panel B. NO₂ emissions, Rome (parts per billion (ppb))



Panel C. NO₂ emissions, New York (parts per billion (ppb))



Panel C. NO₂ emissions, Stockholm (parts per billion (ppb))



Source: Air Quality Open Data Platform, OxCGRT Stringency Index and IMF Staff calculations.
 Note: levels of emissions are smoothed with a five-day moving average to remove excess volatility.

Covid Economics 24, 1 June 2020: 32-75

C. Methodology

This section describes the empirical methodology used to examine the causal effect of containment measures on economic activity. Establishing causality is difficult in this context because the decision of countries to implement containment measures crucially depends on the evolution of the virus, which in turn may affect mobility and economic activity (Maloney and Taskin 2020). This implies that addressing causality requires the researcher to effectively control for this endogenous response. Failure to control for possible reverse causality would result in biased estimates of the effect of containment measures.

We address this issue by controlling for the change in the number of infected cases and deaths the day before implementation of containment measures, as well as for lagged changes in daily economic indicators. Given lags in the implementation of interventions at daily frequency, this allows one to effectively control for the endogenous response of containment measures to the spread of the virus. To further account for expectations about the country-specific evolution of the pandemic, we also control for country-specific time trends.

Two econometric specifications are used to estimate the effect of containment measures on economic activity. The first establishes whether containment measures had, on average, significant effects. The second assesses whether these effects vary across countries depending on country-specific policy responses, such as the magnitude of the fiscal and monetary policy support.

We follow the approach proposed by Jordà (2005) to assess the dynamic cumulative effect of containment measures on economic activity, a methodology used also by Auerbach and Gorodnichenko (2013), Ramey and Zubairy (2018), and Alesina et al. (2019) among others. This procedure does not impose the dynamic restrictions embedded in vector autoregressions and is particularly suited to estimating nonlinearities in the dynamic response. The first regression we estimate is:

$$\Delta n_{i,t+h} = u_i + \theta_h c_{i,t} + X'_{i,t} \Gamma_h + \sum_{\ell=1}^{\ell} \psi_{h,\ell} \Delta n_{i,t-\ell} + \varepsilon_{i,t+h} \quad (1)$$

where $\Delta n_{i,t+h} = n_{i,t+h} - n_{i,t+h-1}$ and $n_{i,t}$ is the logarithm of the daily economic indicator (NO₂ emissions in the baseline) in country i observed at date t .⁷ $c_{i,t}$ denotes the OxCGRT Stringency Index. u_i are country-fixed effects to account for time-invariant country-specific characteristics. X is a vector of control variables which includes the amount of number of COVID-19 infections and deaths in country i observed at date t , daily temperature and humidity levels, and country-specific time trends.⁸

The second specification allows the response to vary with countries characteristics. It is estimated as follows:

$$\Delta n_{i,t+h} = u_i + \theta_h^L F(z_{i,t}) c_{i,t} + \theta_h^H (1 - F(z_{i,t})) c_{i,t} + X'_{i,t} \Gamma_h + \sum_{\ell=1}^{\ell} F(z_{i,t}) \psi_{h,\ell} \Delta n_{i,t-\ell} +$$

$$\sum_{\ell=1}^{\ell} (1 - F(z_{i,t})) \psi_{h,\ell} \Delta n_{i,t-\ell} + \varepsilon_{i,t+h}$$

$$\text{with } F(z_{it}) = \exp^{-\gamma z_{it}} / (1 + \exp^{-\gamma z_{it}}), \quad \gamma > 0 \quad (2)$$

where z is a country-specific characteristic normalized to have zero mean and a unit variance.

The weights assigned to each regime vary between 0 and 1 according to the weighting function $F(\cdot)$, so that $F(z_{it})$ can be interpreted as the probability of being in a given state of the economy. The coefficients θ_h^L and θ_h^H capture the impact of containment measures at each horizon h in cases of very low levels of z ($F(z_{it}) \approx 1$ when z goes to minus infinity) and very high levels of z

⁷ Given the large volatility in the daily economic indicators, we smooth their time series using a 5-day moving average. However, the results are very similar when using non-smoothed series (see Appendix Figure A1 for NO₂ emissions).

⁸ Since emissions are affected by climatic conditions, in the analysis using NO₂ as a dependent variable we include temperature and humidity levels as controls—the results, however, are almost identical excluding these variables. Data are collected from the Air Quality Open Data Platform and include humidity and temperature for each major city, based on the median of several stations, from January 1, 2020.

($1 - F(z_{it}) \approx 1$ when z goes to plus infinity), respectively. $F(z_{it})=0.5$ is the cutoff between low and high country-specific policy responses—that is, for example, low and high fiscal stimulus.

This approach is equivalent to the smooth transition autoregressive model developed by Granger and Teräsvirta (1993). The advantage of this approach is twofold. First, compared with a model in which each dependent variable would be interacted with a measure of country-specific characteristics, it permits a direct test of whether the effect of containment measures varies across different country-specific “regimes”. Second, compared with estimating structural vector autoregressions for each regime, it allows the effect of containment measures to vary smoothly across regimes by considering a continuum of states to compute impulse responses, thus making the functions more stable and precise.

Equations (1 and 2) are estimated for each day $h=0,\dots,30$. Impulse response functions are computed using the estimated coefficients θ_h , and the 95 percent confidence bands associated with the estimated impulse-response functions are obtained using the estimated standard errors of the coefficients θ_h , based on robust standard errors clustered at the country level.

Our sample consists of a balanced sample of 57 economies with at least 30 observation days after a significant outbreak (100 cases). The data cut-off date is May 26, 2020.

III. RESULTS


A. Baseline

Figure 2 shows the estimated dynamic response of NO_2 emissions to a unitary change in the aggregate containment stringency index over the 30-day period following the implementation of the containment measure, together with the 95 percent confidence interval around the point estimates. The left-hand panel shows the responses of daily change of NO_2 emissions while the

right-hand panel shows the cumulative response (which can be thought of as a proxy for lost output).

The results provide evidence that containment measures have significantly reduced the amount of NO₂ emissions. They suggest that in countries where stringent containment measures have been implemented, these may have reduced the amount of NO₂ emissions cumulatively by almost 99 percent 30 days after the implementation, relative to the underlying country-specific path in the absence of intervention.

Figure 2: Effect of Containment Measures on Total Nitrogen Dioxide (NO₂) Emissions



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta n_{it+h} = u_i + \theta_h c_{it} + X'_{it} \Gamma_h + \sum_{\ell=1}^{\ell} \psi_{h,\ell} \Delta n_{it-\ell} + \varepsilon_{it+h}$ where $\Delta n_{it+h} = n_{it+h} - n_{it+h-1}$ and n_{it} is the logarithm of NO₂ emissions in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; c_{it} is the index capturing the level of containment and mitigation measures; X is a matrix of time varying control variables and country specific linear time trend. Results are based on May 26 data.

B. NO₂ emissions and Industrial Production

In order to translate the drop in NO₂ emissions to losses in economic activity, we estimate the relationship between NO₂ emissions and industrial production indices using a monthly database of industrial production indices for 38 countries and monthly levels of NO₂ emissions from January 2019 to April 2020. The panel regression is estimated as follows:

$$\Delta IP_{i,t} = \alpha + \beta \Delta NO_2 + \mu_i + \varepsilon_{i,t} \quad (3)$$

where $\Delta IP_{i,t}$ is the monthly growth rate of industrial production, and ΔNO_2 is the monthly growth rate of NO₂ emissions. The results show that a one percent drop in NO₂ emissions is associated with a 0.015 percent decline in industrial production.⁹ Translating the estimated effect on NO₂ presented before, this implies that containment measures may have led to a 15 percent decline (month-on-month) of industrial production.¹⁰

C. Robustness checks

We conducted several robustness checks of our main finding. First, we included additional controls in the regressions that could be correlated with the level of emissions, such as daily time fixed effects. Second, we repeated the analysis adding changes in mobility as controls to account for the fact that in many cases, containment measures were anticipated and often announced before

⁹ $\Delta IP_{i,t} = 0.357 + 0.015 * \Delta NO_2$, with parenthesis denoting standard errors clustered at the country level.
(0.035) (0.006)

The results is consistent with previous studies highlighting a strong positive correlation between industrialization and emissions of pollutants, including NO₂ emissions (see, for example, Akimoto 2003; Cherniwchan 2012).


¹⁰ The percent effect is computed by $(e^{0.15}-1)*100$.

implementation. This may have reduced mobility ahead of the enforcement of containment measures, thus biasing the baseline estimates. Third, we repeated the analysis excluding China from the sample. This is due to the fact that containment measures were introduced first in China, therefore creating a risk that the longer-term (30 days) results may simply reflect the decline in economic activity in China. To further mitigate reverse causality, we use the contemporaneous change in NO_2 emissions as a control and estimate the impact only after one day of the implementation of containment measures. In all cases, the results are very similar to, and not statistically different from, the baseline (Figure 3).


Finally, another concern is related to the potential seasonality of NO_2 emissions. In particular, it could be the case that the level of emissions tends to systematically decline during the first months of the year—the main sample of our analysis. To check for this possibility, we estimate the relationship between NO_2 emissions and monthly fixed effects using a monthly database of 38 countries from January 2019 to May 2020. The results, not reported, show that (with the exception of July and October) monthly fixed effects are typically not statistically significant, suggesting that seasonality is not an important empirical issue in our analysis.

Figure 3: Robustness Checks (log-differences * 100)


Response to stringency of containment measures: with time fixed effects




Response to stringency of containment measures: with mobility controls



Response to stringency of containment measures: excluding China



Response to stringency of containment measures: contemporaneous dependent variable



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta n_{i,t+h} = u_i + \theta_h c_{i,t} + X'_{i,t} \Gamma_h + \sum_{\ell=1}^{\ell} \psi_{h,\ell} \Delta d_{i,t-\ell} + \varepsilon_{i,t+h}$ where $\Delta n_{i,t+h} = n_{i,t+h} - n_{i,t+h-1}$ and $n_{i,t}$ is the logarithm of the level of NO_2 emissions in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; $c_{i,t}$ is the index capturing the level of containment and mitigation measures; X is a matrix of time varying control variables and country specific linear time trend. Results are based on May 26 data.

Covid Economics 24, 1 June 2020: 32-75

D. Impact of containment on other indicators of economic activity


In this section, we examine whether containment measures have had an impact on other indicators of economic activity. Namely, we focus on the impact of stringency measures on: (i) total flights; (ii) energy consumption; (iii) maritime import indices; (iv) maritime export indices; (v) retail mobility indices; and (vi) transit indices. These variables can shed lights on the effect of containment measure on different sectors of the economy, such as tourism, trade, and retail consumption.

Results for equation (1) for each indicator are reported in Figure 4. They suggest that the impact of containment measures has been overwhelmingly adverse across all sectors, and most importantly tourism. Specifically, the results indicate that containment measures have reduced the total number international and domestic flights by more than 99 percent in the 30-day period following the implementation of containment measures. Total energy consumed has declined by more than 95 percent; maritime imports and exports have been reduced by over 40 percent, with a more pronounced impact on exports; retail and transit mobility have been reduced by more than 400 percent relative to country-specific paths in the absence of intervention.¹¹

¹¹ As for NO₂, the percent effects are computed as $(e^{0h}-1)*100$. We also find that energy consumption as well as flights are positively correlated with industrial production growth—both correlations are statistically significant at 5 percent.

Figure 4: Local projection response to indicators of economic activity

(log-differences * 100)



Note: Impulse response functions are estimated, using a sample of 119 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta e_{it+h} = u_i + \theta_h c_{it} + X'_{it} \beta_h + \sum_{\ell=1}^L \psi_{h,\ell} \Delta e_{it-\ell} + \varepsilon_{it+h}$ where $\Delta e_{it+h} = e_{it+h} - e_{it+h-1}$ and e_{it} is the logarithm of the economy activity indicator (depending on specification) in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; c_{it} the index capturing the level of containment and mitigation measures; X is a matrix of time varying control variables and country specific time trend. Results are based on May 26 data.

E. Role of macro policy responses in mitigating the fallout in economic activity

Governments and central banks around the world have implemented unprecedented economic measures in response to the COVID-19 pandemic. This section examines whether such measures have been effective in mitigating the negative effects of containment measures, using data on discretionary fiscal and monetary measures implemented in response to COVID-19 provided by the *IMF Policy Tracker*. We explore whether the average effect of containment measures varies depending on the magnitude of country policy responses deployed.

Fiscal stimulus

As of May 26th, 2020, more than 90 countries worldwide had deployed (or announced) fiscal measures to mitigate the impact of the pandemic. Fiscal packages have been heterogeneous in size, ranging from less than 1 percent of GDP, to as much 12 percent of GDP for economies such as El Salvador, Japan, Luxembourg, and Macao SAR (Figure 5, Panel A). On average, fiscal stimulus used in Advanced Economies (AEs) averaged at 5 percent of GDP, compared to 2.3 percent in Emerging Market and Developing Economies (EMDEs).

To examine the role of fiscal stimulus in mitigating the decline in NO₂ emissions, we estimate equation (2) with an interaction term which measures the amount of fiscal stimulus (as a percent of GDP) deployed since the beginning of the pandemic. The results in Figure 6 (top panel) show that containment measures have had a much larger adverse impact on economic activity in countries with relatively small fiscal packages—equivalent to a 22 percent decline in industrial production. In contrast, the impact is not statistically different from zero in countries that deployed large fiscal stimulus packages. Consistent with the evidence of Ma et al. (2020) on previous pandemics, this suggests that fiscal stimulus measures can play a crucial role during the COVID-19 pandemic to mitigate the economic fallout of the crisis.

Policy interest rate cuts



Policy rates have been reduced in 97 countries from January 2019 to-date (Figure 5B). With policy rates closer to zero-lower-bound in AEs, policy rate cuts were much larger in EMDEs: more than 10 EMDEs lowered their policy rates by over 200 bps, with Ukraine cutting its policy rate by 400 bps.

The results in Figure 6 (bottom panel) are obtained by estimating equation (2) using the cumulative policy rate cut as an interaction term. They suggest that in countries where central banks lowered policy rates more aggressively, the adverse impact of containment measures was mitigated to a greater extent. We find that the economic impact of containment measures is much more adverse in countries where monetary policy was not eased. In contrast, the impact of containment measures in countries with large cuts in policy rates is not statistically significant. The results suggest that monetary policy plays a significant role and may have helped in offsetting the economic fallout from the COVID-19 pandemic.

Figure 5: Policy Responses to the COVID-19 Pandemic

Panel A. Fiscal Stimulus (in percent of GDP)

Panel B. Policy Rate Cuts (in basis points)




Source: IMF Policy Tracker.


Figure 6: Interaction with Fiscal and Monetary Policy

(log-differences * 100)

Interaction with Fiscal Policy



Interaction with Monetary Policy



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on
$$\Delta n_{i,t+h} = u_i + u_t + \theta_h^L F(z_{it}) c_{it} + \theta_h^H (1 - F(z_{it})) c_{it} + X'_{i,t} \Gamma_h + \sum_{\ell=1}^L F(z_{it}) \psi_{h,\ell} \Delta n_{i,t-\ell} + \sum_{\ell=1}^L (1 - F(z_{it})) \psi_{h,\ell} \Delta n_{i,t-\ell} + \varepsilon_{it+h}$$
 with $F(z_{it}) = \frac{\exp(-\gamma z_{it})}{(1 - \exp(-\gamma z_{it}))}$, $\gamma > 0$ where $\Delta n_{i,t+h} = n_{i,t+h} - n_{i,t+h-1}$ and $n_{i,t}$ is the logarithm of NO₂ emissions in country i observed at date t and z is the country-specific characteristics normalized to have zero mean and a unit variance. The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2 \dots L$; $c_{i,t}$ is the index capturing the level of containment and mitigation measures; X is a matrix of time varying control variables and country specific linear time trend. Results are based on May 26 data.

F. Cost-effectiveness of different containment measures

In this section, we explore how different containment measures compare in terms of economic cost—through their impact on economic activity and effectiveness. Our purpose is to examine which types of containment measure resulted in larger short-term tradeoffs between minimizing health risks and economic losses. This can inform the discussion of how countries should open-up their economies as well as how best they can respond to any second wave of infections.

For this purpose, we analyze the effects on economic activity, infections, and deaths, of the following containment measures: (i) school closures; (ii) workplace closures; (iii) cancellation of public events; (iv) restrictions on size of gatherings; (v) closures of public transport; (vi) stay-at-home orders; (vii) restrictions on internal movement; and (viii) restrictions on international travel.

To estimate the effects of different containment measures on infections and deaths, we follow the approach used by Deb et al. (2020), and adapt equation (1) to the following:

$$\Delta d_{i,t+h} = u_i + \theta_h c_{i,t} + X'_{i,t} \Gamma_h + \sum_{\ell=1}^{\ell} \psi_{h,\ell} \Delta d_{i,t-\ell} + \varepsilon_{i,t+h} \quad (4)$$

where $\Delta d_{i,t+h} = d_{i,t+h} - d_{i,t+h-1}$ and $d_{i,t}$ is the logarithm of the number of deaths (infections), in country i observed at date t . $c_{i,t}$ denotes the OxCGRT Stringency Index. u_i are country-fixed effects to account for time-invariant country-specific characteristics (for example, population density, age profile of the population, health capacity, average temperature, etc.). X is a vector of

control variables which includes daily temperature and humidity levels, in addition to country-specific time trends.¹²

As noted earlier, estimating the overall effect of each measure is challenging, because many of the measures were introduced simultaneously. Following Deb et al. (2020), we use two alternative approaches to gauge the potential magnitude of the effect of each of measure. In the first, we introduce each measure one at a time in equations (1) and (4) respectively. Clearly, the problem with this approach is that the estimates suffer from omitted variable bias. In the second approach, we include them all together. While this approach addresses omitted variable bias, the estimates are likely to be less precise due to multicollinearity. We report results for the second approach in Appendix Figures A5-A6.

The results for the effects of different containment measures on economic activity and infections and deaths are summarized in Table 2 and reported in Appendix Figures A3-A4. They suggest that stay-at-home requirements and workplace closures are most effective in curbing infections and deaths; however, those measures are also associated with the largest economic losses. The results suggest that restrictions on internal movement, though costly in economic terms, are not as effective in curbing fatalities. Finally, less costly containment measures, such as school closures and restrictions on international travel, are nonetheless successful in lowering COVID-19 infections.

¹² As a robustness check, we used a dummy variable to identify the start and end of different containment and mitigation measures—this is similar to treating the containment measures as a shock. The results in Appendix Figure A2 are very similar to, and not statistically different, from the baseline.

Table 2. Cumulative effect of containment measure, 30 days after its introduction (log-differences * 100)

	NO ₂ emissions	Confirmed Cases	Confirmed Deaths
Workplace closures	-423	-101	-100
Cancellation of public events	-357	-138	-160
Stay-at-home requirements	-348	-101	-132
Restrictions on gathering size	-339	-104	-70
Closures of public transport	-331	-65	-81
Restrictions on internal movement	-311	-79	-91
School closures	-292	-79	-76
Restrictions on international travel	-233	-109	-151

Note: the results reported in Table 2 denote the cumulative local projection response to NO₂ emissions, confirmed cases, and confirmed deaths, to each type of containment measure. ' denotes that results are **not** significant 30 days after the introduction of containment measures. Estimates based on $\Delta n_{i,t+h} = u_i + \theta_n c_{i,t} + X'_{i,t} \Gamma_h + \sum_{\ell=1}^L \psi_{h,\ell} \Delta n_{i,t-\ell} + \varepsilon_{i,t+h}$ where $\Delta n_{i,t+h} = n_{i,t+h} - n_{i,t+h-1}$ and $n_{i,t}$ is the logarithm of NO₂ emissions (or infections/deaths) in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2 \dots L$; $c_{i,t}$ is the index capturing different types containment and mitigation measures, introduced one at a time; X is a matrix of time varying control variables and country-specific linear trend. Results are based on May 26 data.

IV. CONCLUSIONS

Containment measures, though crucial to halting the spread of COVID-19 and limiting the number of fatalities in the absence of effective therapies and vaccines, have resulted in large short-term economic losses. In this paper, we provide a first empirical assessment on the impact of COVID-19 containment measures on economic activity, through the use of a novel daily database of high-frequency indicators of economic activity, including Nitrogen Dioxide (NO₂) emissions, international and domestic flights, energy consumption, maritime trade, and retail mobility indices.

Results suggest that containment measures have had, on average, very large impacts on NO₂ emissions, with the decline in emissions levels equivalent to a loss of about 15 percent in industrial production over the 30-day period following the implementation of the containment measure. Results for other indicators of economic activity suggest that containment measures have had a very large adverse impact on flights worldwide, energy consumption, maritime trade, and retail and transit mobility.

Fiscal and monetary policy deployed during the COVID-19 crisis have played an important role in mitigating the impact of containment measures on economic activity: results suggest that short-term economic losses are greater in countries where less fiscal stimulus was deployed, and where monetary policy easing was more limited.

Among different types of containment measure, workplace closures and stay-at-home orders are the most effective in flattening COVID-19 related infections and deaths but are the costliest in terms of their impact on economic activity. Less costly containment measures, such as school closures and restrictions on international travel, are successful in reducing COVID-19 infections, but less effective in curbing fatalities.


References

- Alesina, A.F., Furceri, D., Ostry, J.D., Papageorgiou, C. and Quinn, D.P., 2019. Structural Reforms and Elections: Evidence from a World-Wide New Dataset (No. w26720). National Bureau of Economic Research.
- Akimoto H. 2003. Global Air Quality and Pollution. *Science*. 302(5651), pp.1719-17191
- Auerbach, A.J. and Gorodnichenko, Y., 2013. Output Spillovers from Fiscal Policy. *American Economic Review*, 103(3), pp.141-46.
- Cerdeiro D., Komaromi A., Lui Y. and Saeed M. 2020. World Seaborne Trade in Real Time: A Proof of Concept for Building AIS-based Nowcasts from Scratch. IMF Working Paper, Forthcoming.
- Cherniwchan J. Economic Growth, Industrialization, and the Environment. *Resource and Energy Economics*. 34(4), pp. 442-467
- Chinazzi et al., 2020. The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science*, 368, 395-400.
- Coibon O., Gorodnichenko Y., Weber M, (2020). The Cost of the COVID-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending. National Bureau of Economic Research Working Paper no.27141.
- Deb P., Furceri D., Ostry J., and Tawk N. (2020). The Effects of Containment Measures on the COVID-19 Pandemic. *Covid Economics: Vetted and Real-Time Papers*. 2020(19), pp.53-86.
- Eichenbaum M., Rebelo S., Trabandt M., 2020. The Macroeconomics of Epidemics. NBER Working Papers 26882, National Bureau of Economic Research, Inc.
- Glaeser E., Kim H., and Luca M. 2017. Nowcasting the Local Economy: Using Yelp Data to Measure Economic Activity. National Bureau of Economic Research Working Paper No. 024010.
- Granger, C.W.J. and Teräsvirta, T., 1993. *Modelling Nonlinear Economic Relationships* Oxford University Press. New York.
- Jordà, Ò., 2005. Estimation and Inference of Impulse Responses by Local Projections. *American economic review*, 95(1), pp.161-182.
- Kraemer et al., 2020. The effect of human mobility and control measures on the COVID-19 epidemic in China. *Science*, 368, 493-497.
- Kumar S., Pranab M., 2019. A Novel GDP Prediction Technique Based on Transfer Learning Using CO2 Emission Dataset. *Applied Energy*, 253- 113476.

- H. Tian et al., 2020. An investigation of transmission control measures during the first 50 days of the COVID-19 epidemic in China. *Science* 10.1126/science.abb6105 (2020).
- Ma C., Rogers J., Zhou S., 2020. Modern Pandemics: Recession and Recovery. Working paper.
- Marjanovic M., Milovancevic M., Mladenovic I., 2016. Prediction of GDP Growth Rate Based on Carbon Dioxide (CO₂) Emissions. *Journal of CO₂ Utilization*, 16, 212-217.
- Ramey, V.A. and Zubairy, S., 2018. Government spending multipliers in good times and in bad: evidence from US historical data. *Journal of Political Economy*, 126(2), pp.850-901.
- Small C., Elvidge C. Balk D, and Montgomery M. 2011. Spatial scaling of stable night lights. *Remote Sensing of Environment*, 115(2), pp.269-28


ANNEX

Figure A1: Local projection response of NO₂ emissions (unsmoothed) to containment measures
(log-differences * 100)




Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta n_{i,t+h} = u_i + \theta_h c_{i,t} + X'_{i,t} \beta_h + \sum_{\ell=1}^L \psi_{h,\ell} \Delta n_{i,t-\ell} + \varepsilon_{i,t+h}$ where $\Delta n_{i,t+h} = n_{i,t+h} - n_{i,t+h-1}$ and $n_{i,t}$ is the logarithm of NO₂ emissions in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; $c_{i,t}$ the index capturing the level of containment and mitigation measures; X is a matrix of time varying control variables and country specific time trend. Results are based on May 26 data.

Figure A2: Local projection response of NO₂ emissions to containment measures dummy
(log-differences * 100)



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures.
 Estimates based on $\Delta n_{it+h} = u_i + \theta_h c_{it} + X'_{it} \Gamma_h + \sum_{\ell=1}^h \psi_{h,\ell} \Delta n_{it-\ell} + \varepsilon_{it+h}$ where $\Delta n_{it+h} = n_{it+h} - n_{it+h-1}$ and n_{it} is the logarithm of NO₂ emissions in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; c_{it} is an index dummy capturing different types containment and mitigation measures, introduced one at a time; X is a matrix of time varying control variables and country specific linear trend. Results are based on May 26 data.


Figure A3: Local projection response of NO₂ emissions to different containment measures (log-differences * 100)



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures.
 Estimates based on $\Delta n_{i,t+h} = u_i + \theta_h c_{it} + X_{i,t} \beta_h + \sum_{\ell=1}^h \psi_{h,\ell} \Delta n_{i,t-\ell} + \varepsilon_{i,t+h}$ where $\Delta n_{i,t+h} = n_{i,t+h} - n_{i,t+h-1}$ and $n_{i,t}$ is the logarithm of NO₂ emissions in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; c_{it} is the index capturing different types containment and mitigation measures, introduced one at a time; X is a matrix of time varying control variables and country specific linear trend. Results are based on May 26 data.


Figure A4a: Local projection response of confirmed infections and deaths to different containment measures

(log-differences * 100)



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta n_{it+h} = u_i + \theta_h c_{it} + X'_{it} \beta_h + \sum_{\ell=1}^L \psi_{h,\ell} \Delta n_{it-\ell} + \varepsilon_{it+h}$ where $\Delta n_{it+h} = n_{it+h} - n_{it+h-1}$ and n_{it} the logarithm of the number of COVID-19 cases or deaths (depending on specification) in country i observed at date t . The model is estimated, at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; c_{it} is the index capturing different types containment and mitigation measures, introduced one at a time; X is a matrix of time varying control variables and country specific linear time trend. Results are based on May 26 data.


Figure A4b: Local projection response of confirmed infections and deaths to different containment measures
(log-differences * 100)



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta n_{i,t+h} = u_i + \theta_h c_{it} + X'_{i,t} \Gamma_h + \sum_{\ell=1}^L \psi_{h,\ell} \Delta n_{i,t-\ell} + \varepsilon_{i,t+h}$ where $\Delta n_{i,t+h} = n_{i,t+h} - n_{i,t+h-1}$ and $n_{i,t}$ the logarithm of the number of COVID-19 cases or deaths (depending on specification) in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; c_{it} is the index capturing different types containment and mitigation measures, introduced one at a time; X is a matrix of time varying control variables and country specific linear time trend. Results are based on May 26 data.

Figure A4c: Local projection response of confirmed infections and deaths to different containment measures


(log-differences * 100)



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta n_{it+h} = u_i + \theta_h c_{it} + X'_{it} \Gamma_h + \sum_{\ell=1}^L \psi_{h,\ell} \Delta n_{it-\ell} + \varepsilon_{it+h}$ where $\Delta n_{it+h} = n_{it+h} - n_{it+h-1}$ and n_{it} the logarithm of the number of COVID-19 cases or deaths (depending on specification) in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; c_{it} is the index capturing different types containment and mitigation measures, introduced one at a time; X is a matrix of time varying control variables and country specific linear time trend. Results are based on May 26 data.

Figure A4d: Local projection response of confirmed infections and deaths to different containment measures


(log-differences * 100)



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta n_{it+h} = u_i + \theta_h c_{it} + X'_{it} \beta_h + \sum_{\ell=1}^L \psi_{h,\ell} \Delta n_{it-\ell} + \varepsilon_{it+h}$ where $\Delta n_{it+h} = n_{it+h} - n_{it+h-1}$ and n_{it} the logarithm of the number of COVID-19 cases or deaths (depending on specification) in country i observed at date t . The model is estimated, at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; c_{it} is the index capturing different types containment and mitigation measures, introduced one at a time; X is a matrix of time varying control variables and country specific linear time trend. Results are based on May 26 data.

Figure A5: Local projection response of NO₂ emissions to different containment measures (together)


(log-differences * 100)



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta n_{i,t+h} = u_i + \theta_h c_{it} + X'_{i,t} \Gamma_h + \sum_{\ell=1}^h \psi_{h,\ell} \Delta n_{i,t-\ell} + \varepsilon_{i,t+h}$ where $\Delta n_{i,t+h} = n_{i,t+h} - n_{i,t+h-1}$ and $n_{i,t}$ the logarithm of the number of COVID-19 cases or deaths (depending on specification) in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; c_{it} is the index capturing different types containment and mitigation measures, introduced altogether; X is a matrix of time varying control variables and country specific linear time trend. Results are based on May 26 data.

Figure A6a: Local projection response confirmed cases and deaths to different containment measures (together)


(log-differences * 100)



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta n_{i,t+h} = u_i + \theta_h c_{it} + X'_{i,t} \Gamma_h + \sum_{f=1}^L \psi_{h,f} \Delta n_{i,t-f} + \varepsilon_{i,t+h}$ where $\Delta n_{i,t+h} = n_{i,t+h} - n_{i,t+h-1}$ and $n_{i,t}$ the logarithm of the number of COVID-19 cases or deaths (depending on specification) in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $f = 1, 2, \dots, L$. c_{it} is the index capturing different types containment and mitigation measures, introduced altogether; X is a matrix of time varying control variables and country specific linear time trend. Results are based on May 26 data.

Figure A6b: Local projection response of confirmed cases and deaths of different containment measures (together)

(log-differences * 100)



Note: Impulse response functions are estimated using a sample of 57 countries using daily data from the start of the outbreak. The analysis is restricted to countries with a significant outbreak that has lasted at least 30 days. $t = 0$ is the date when the outbreak becomes significant (100 cases) in each country. The graph shows the response and confidence bands at 95 percent. The horizontal axis shows the response x days after the containment measures. Estimates based on $\Delta n_{i,t+h} = u_i + \theta_h c_{it} + X'_{it} \Gamma_h + \sum_{\ell=1}^L \psi_{h,\ell} \Delta n_{i,t-\ell} + \varepsilon_{it+h}$ where $\Delta n_{i,t+h} = n_{i,t+h} - n_{i,t+h-1}$ and $n_{i,t}$ the logarithm of the number of COVID-19 cases or deaths (depending on specification) in country i observed at date t . The model is estimated at each horizon $h = 0, 1, \dots, H$, with a lag structure $\ell = 1, 2, \dots, L$; c_{it} is the index capturing different types containment and mitigation measures, introduced altogether; X is a matrix of time varying control variables and country specific linear time trend. Results are based on May 26 data.

God is in the rain: The impact of rainfall-induced early social distancing on Covid-19 outbreaks¹

Rolly Kapoor,² Haedong Aiden Rho,³ Kinpritma Sangha,⁴
Bhavyaa Sharma,⁵ Ajay Shenoy⁶ and Guanghong Xu⁷

Date submitted: 26 May 2020; Date accepted: 27 May 2020

We test whether earlier social distancing affects the progression of a local COVID-19 outbreak. We exploit county-level rainfall on the last weekend before statewide lockdown. After controlling for historical rainfall, temperature, and state fixed-effects, current rainfall is a plausibly exogenous instrument for social distancing. Early distancing causes a reduction in cases and deaths that persists for weeks. The effect is driven by a reduction in the chance of a very large outbreak. The result suggests early distancing may have sizable returns, and that random events early in an outbreak can have persistent effects on its course.

- 1 We are grateful to SafeGraph for providing access to their data. We also appreciate helpful comments from Laura Giuliano, Justin Wolfers, and innumerable individuals on Twitter.
- 2 Ph.D. student, University of California, Santa Cruz.
- 3 Ph.D. student, University of California, Santa Cruz.
- 4 Senior Manager, Anlitiks, Inc.
- 5 Ph.D. student, University of California, Santa Cruz.
- 6 Assistant Professor, University of California, Santa Cruz.
- 7 Ph.D. student, University of California, Santa Cruz.

Copyright: Rolly Kapoor, Haedong Aiden Rho, Kinpritma Sangha,
Bhavyaa Sharma, Ajay Shenoy and Guanghong Xu

1 Introduction


As COVID-19 outbreaks have spiraled in some regions while sparing others, it remains an open question whether earlier social distancing played a key role. California's statewide lockdown began only 3 days earlier than New York's, yet California has suffered far fewer deaths. Six Bay Area counties within California (as well as Santa Cruz County) began their lockdown a few days earlier than the rest of the state. They have had a more favorable trajectory than other parts of the state despite suffering much earlier exposure to the disease.

But naïve comparisons between states risk conflating the impact of earlier distancing with differences in state characteristics. Figure 1 shows that states that issued earlier lockdowns have higher median incomes and more college degree holders, but fewer black and older residents. Even within a state, locales that issued earlier lockdowns may differ systematically in ways that may or may not be observable. For example, the Associated Press reports that the Bay Area lockdown had its roots in an association of local health officials that formed during the AIDS epidemic and has met regularly to discuss prior epidemics like Ebola and swine flu (Rodriguez, 19 April 2020). The presence of such an institution may have had other impacts on the local response to COVID-19 beyond the lockdown, making it difficult to isolate the effect of early social distancing. The problem of selection bias is compounded by the problem of measurement. It is possible that the states and counties that responded more quickly are also more active in testing for the disease, creating non-classical measurement error.

We sidestep these challenges by exploiting within-state variation in early social distancing induced by rainfall. We measure county-level rainfall on the last weekend before the county's home state went into mandatory lockdown. This key weekend is the last day that people had wide discretion in leaving home for reasons unrelated to work (dining at restaurants, for example). After controlling for average historical rainfall, temperature, and state fixed effects, rainfall on this specific weekend is plausibly exogenous. Counties that had heavy rainfall were exogenously induced to exercise a marginal degree of extra social distancing just a few days before counties that had less rainfall. We measure whether these counties had fewer COVID-19 cases and deaths in the weeks after the statewide lockdown.

We detect highly significant effects even two weeks after the statewide lockdown, many days

Figure 1
States that Lock Down Earlier are Systematically
Different on Baseline Characteristics



Note: The size of each circle is proportional to the number of states that shut down on that date. Demographics are from the 2014-2018 American Community Survey (5-year estimates). Dates of state-wide lockdown orders come from the Institute of Health Metrics and Evaluation. See Section 2.1 for details about the data.

after the crucial weekend. The two-stage least squares estimates imply that a 1 percentage point increase in the number of people leaving home causes an additional 14 cases and 1.3 deaths per 100,000 residents. These effects are all the more remarkable because the variation in social distancing induced by rainfall, though precise, is relatively small. But the impact of the initial reduction is propagated over time. We measure growing impacts that have not leveled off even 18 days after the lockdown, nearly 3 weeks after the crucial weekend. These effects appear to be driven by the right tail of the distribution. Counties where more people left home on the pre-shutdown weekend are no more likely to have a marginally higher case count, but are slightly more likely to have a big outbreak. This result is what might be expected given that differences in the number of infections on the eve of a statewide lockdown will either vanish or be drastically amplified depending on whether the county lowers the viral reproduction rate below 1 and avoids “superspreader” events.

Our paper joins a small but growing number of papers that study the impact of social distancing on COVID-19 transmission. Our research question is most similar to Pei et al. (2020), who use an epidemiological model to simulate COVID-19 trajectories in a counterfactual world where lockdowns had begun a few weeks sooner. Our study approaches this question using a natural experiment rather than a model. A few recent studies (Courtemanche et al., 2020; Fowler et al., 2020) use difference-in-differences designs to study the impact of statewide closures and lockdowns on transmission. Aside from exploiting an orthogonal source of variation, our study aims to answer a different question: whether marginal improvements in early distancing can affect medium-run outcomes.

Meanwhile, Brzezinski et al. (2020) use state-level rainfall and temperature as exogenous variation in non-mandated social distancing to study whether state governments are less likely to mandate social distancing where it is already being practiced.¹ Methodologically our study is most similar to Madestam et al. (2013), which measures the impact of rainfall on a single pivotal date (Tax Day 2010) to measure the long-run impacts of Tea Party protests. One major advantage to studying a one-time shock rather than panel variation is that we can fully trace differential trajectories across counties. And since that shock is on the weekend before statewide lockdown, it

¹ Since we exploit only within-state variation, their result is not a threat to our design. We verify in Section 3.4 and Appendix A.7 that county-level policy responses do not bias our results.

is the closest possible counterfactual to having a longer policy of social distancing.

Our results suggest that even small differences in the extent of early social distancing can have sizable impacts on the scale of the outbreak. As states begin to loosen their stay-at-home orders, health officials are considering plans to potentially return to lockdown if there are signs of a resurgence. Subject to caveats discussed in the final section, our results suggest moving even a few days more quickly could make a measurable difference. Our results also suggest that completely random events early in the course of a local outbreak can have surprisingly persistent effects on its size.

2 Research Design

2.1 Data

Weather : We measure rainfall by spatially merging weather stations from the Global Historical Climatology Network-Daily Database (Menne et al., 2012) to U.S. counties based on 2012 Census TIGER/Line shapefiles. We calculate county-level average precipitation and daily maximum temperatures. For each day in 2020 we calculate the average precipitation and max temperature for that same day-of-year from 2015—2019. We then take the inverse hyperbolic sine of all of these quantities. From here on we refer to these transformed quantities as simply current or historical rainfall and temperature.²

Social Distancing: Our primary measure of social distancing is the percentage of people that leave home, calculated using aggregated mobile phone GPS data provided by SafeGraph (SafeGraph, 2020a). The data report the total devices in SafeGraph's sample by block group, and the number that leave their home.³ We aggregate these two counts by county and calculate the percentage leaving home.

Leaving home is our first-stage regressor because keeping people home is the primary impact of rain on social distancing, and keeping people at home for an extra weekend is the most

² The inverse hyperbolic sine transformation $\log(x + \sqrt{x^2 + 1})$ is a convenient approximation to the natural logarithm that is well-defined when $x = 0$ and converges to $\log 2 + \log x$ as $x \rightarrow \infty$. Figure 2 suggests this transformation is successful in producing a roughly linear relationship.

³ SafeGraph defines "home" as the "common nighttime location of each mobile device over a 6 week period to a Geohash-7 granularity (153m x 153m)." Leaving home is defined as leaving that square.

natural analogy to locking down a few days sooner. But to better understand what activities people are deterred from doing when they stay home—and whether those who do leave change where they go—we draw on several other measures of social distancing. We use two measures of indoor exposure. The first is the Device Exposure Index (Couture et al., 2020a), which represents the number of people (cell phones) an average individual was exposed to in small commercial venues within the county. We also use SafeGraph’s Weekly Patterns data to compute a measure of “gatherings” based on whether more than 5 devices ping within a single indoor non-residential location within one hour (SafeGraph, 2020b). Since the SafeGraph sample represents roughly 6% of a typical county, 5 devices represent a large number of people. We rescale both measures by their daily average on the first full weekend in March, meaning a value of 100 denotes the same exposure or number of gatherings as the first weekend of March (which was before any local or state lockdown).

We also use several measures of long-distance travel. Using SafeGraph’s data we measure the percentage of devices that travel greater or less than 16 kilometers from home (among those that leave home). We also measure cross-county travel using the Location Exposure Index (Couture et al., 2020b). We measure the fraction of people in a county who were not present on any of the prior 14 days.⁴

COVID-19 Cases and Deaths : We measure daily (cumulative) COVID-19 cases/deaths by combining data from Johns Hopkins University and the CoronaDataScraper project (Center for Systems Science and Engineering (Johns Hopkins University); Corona Data Scraper (2020)). As described in detail in Appendix B, we manually corrected missing values by consulting county public health departments and local newspapers. All of these measures are cumulative cases and deaths rather than new cases and deaths. Our primary outcomes are the number of cases and deaths per 100,000 population, measured 14 days after the statewide lockdown.⁵

Demographics : We measure demographic characteristics (such as population size, median income, age profiles of the population) using the 2014-2018 five-year estimates from the Ameri-

⁴ For more information on the Device Exposure Index and the Location Exposure Index see Appendix B.1.

⁵ We choose these measures both because they are the measures most commonly used by policymakers to gauge the severity of an outbreak, and because they give the most accurate reflection of the number of infections relative to the number who could potentially be infected. We choose 14 days as our default horizon because this is the typical quarantine period for the disease, though Section 3.2 shows the impact at every horizon.

can Community Survey (Manson et al., 2019).

Lockdowns: Finally, we measure statewide lockdown dates using the Institute of Health Metrics and Evaluation's record of state policies as of 17 April 2020 (Institute for Health Metrics and Evaluation (2020)). The dataset has all shutdown dates up to 7 April. Any state that had not shut down by that date (or was not recorded as doing so by the Institute) is excluded from our study.

2.2 Instrument and Specifications

Defining the Instrument: We identify the last Saturday and Sunday before the day of the shutdown order. If the shutdown was announced on a Sunday we take only the Saturday of that weekend as the “weekend before.” If it is announced on a Saturday we take the prior weekend. We average rainfall and temperature (both current and historical) as well as social distancing across the days of this weekend. We compute baseline cases and deaths as those recorded for the day before this last weekend, and baseline growth in these measures as the average change in the inverse hyperbolic sine of each in the prior 7 days.

Specification, Identification, and Inference: We estimate first-stage, reduced form, and second-stage regressions of the form

$$D_i = \alpha_s + \gamma R_i + \tau_1 \bar{R}_i + \tau_2 T_i + \tau_3 \bar{T}_i + X_i \omega + u_i \quad (1)$$

$$Y_i = \zeta_s + \rho R_i + \xi_1 \bar{R}_i + \xi_2 T_i + \xi_3 \bar{T}_i + X_i \theta + v_i \quad (2)$$

$$Y_i = \kappa_s + \beta \hat{D}_i + \phi_1 \bar{R}_i + \phi_2 T_i + \phi_3 \bar{T}_i + X_i \vartheta + z_i \quad (3)$$

where i and s index counties and states, D is the percentage of people leaving home, Y is the outcome, α_s and κ_s are state fixed-effects, R and \bar{R} are current and historical rainfall, T and \bar{T} are current and historical temperature, and X is a vector of baseline and demographic control variables that vary across specifications, with the most basic specification having no controls.

We must control for historical rainfall because even within a state, counties that are typically rainy in March and April may be systematically different from those that are not (e.g. Santa Cruz versus San Diego in California). The instrument \bar{R}_i is thus excess or unexpected rainfall, which is plausibly uncorrelated with historical demographic characteristics. We control for temperature because some experts and politicians have hypothesized that it may directly impact COVID-

19 transmission.⁶ The identification assumption is that, after controlling for state fixed-effects, historical rainfall, and temperature, rainfall on the pre-shutdown weekend only affects endline case counts through its impact on the number of people leaving home. We show in Appendix A.1 that, as expected, rainfall is uncorrelated with baseline cases, deaths, and a host of demographic characteristics.

In all of these regressions β is the two-stage least squares estimate of the impact on the outcome of having 1 percentage point more people leave home on the weekend before the lockdown.⁷ Since there is spatial correlation in both rainfall and COVID-19 infections, we cluster standard errors using a 3°x 3° latitude-longitude grid.⁸

Additional Control Variables: Since rainfall is exogenous, the control variables X_i will not affect the consistency of the estimates. But they can make the estimates more precise by reducing the unexplained variation in social distancing and COVID-19 cases and deaths. Our basic specification includes nothing in X_i . Our preferred specification adds controls for baseline COVID-19 prevalence. We include the number of cases per 100,000 at baseline, the raw number of cases at baseline, and the growth rate of cases in the week prior to the pre-lockdown weekend.⁹ Our most comprehensive specification includes baseline controls as well as demographic characteristics.¹⁰

3 Results

3.1 Basic Estimates

First-Stage—Impact of Rainfall on Social Distancing: Column 1 in Panel A of Table 1 shows estimates of the first-stage (Equation 1). After controlling for historical rainfall and tempera-

⁶ Chin et al. (2020), for example, find that temperature affects virus stability in lab samples.

⁷ Since there is a single endogenous regressor and a single excluded instrument, $\hat{\beta} = \hat{\rho}/\hat{\gamma}$.

⁸ To be precise, we generate a grid and assign each county to the cell that contains its centroid.

⁹ We control for both cases per 100,000 and raw case counts at baseline because both are independently informative about social distancing and endline outcomes. That is likely because while the one measures the baseline rate of prevalence, the other drives initial local media coverage. It is also likely that a greater raw number of cases lowers the probability that the infection dies out because all initially infected self-isolate. The case growth rate, which we calculate as the average change in the inverse hyperbolic sine of case counts, is informative about the trajectory prior to the pre-shutdown weekend.

¹⁰ Total population; fraction of population in the bins 60-69, 70-79, and over 80; fraction African American; and median household income.

Table 1
Two-Stage Least Squares Estimates


Panel A: Interpreting the First-Stage						
	First-Stage		Activities Averted by Staying Home			
	(1)	(2)	(3)	(4)	(5)	(6)
	% Leaving Home	Exposure	Gatherings	Travel Near	Travel Far	Non-Locals
Rainfall	-0.432*** (0.126)	-0.876*** (0.327)	-1.670** (0.694)	-0.217* (0.131)	-0.336** (0.146)	-0.267** (0.122)
Counties	1946	1397	1757	1946	1946	1397
Clusters	139	113	124	139	139	113
Outcome Mean	64.77	37.35	38.34	41.41	21.42	9.13
F-stat: Rainfall	11.68	7.18	5.79	2.75	5.27	4.77
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X

Panel B: Reduced-Form						
	Endline Cases/100k			Endline Cases/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-6.776** (3.160)	-6.132*** (1.705)	-5.921*** (1.670)	-0.717 (0.463)	-0.581*** (0.216)	-0.537*** (0.176)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
Outcome Mean	58.12	58.12	58.12	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Panel C: Two-Stage Least Squares						
	Endline Cases/100k			Endline Cases/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
% Leaving Home	15.686 (9.653)	14.596*** (4.852)	14.824*** (5.130)	1.660 (1.274)	1.383** (0.556)	1.344** (0.517)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
First-Stage F	11.68	16.54	17.80	11.68	16.54	17.80
Outcome Mean	58.12	58.12	58.12	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: All standard errors are clustered using a 3°x 3° latitude-longitude grid to adjust for spatial correlation.
Panel A: “Exposure” refers to the Device Exposure Index, a measure of the number of devices (cell phones) visiting small indoor venues. “Gatherings” measures the number of times more than 5 devices ping in a single indoor venue within the span of an hour. Both of these measures are rescaled as a percentage of their level on the weekend 7–8 March. “Travel Near” and “Travel Far” give the percentage of devices that leave home and travel less than versus more than 16 kilometers. “Non-Locals” gives the percentage of devices in the county that were not present on any of the prior 14 days.
Panels B and C: “Baseline Case Controls” are the number of COVID-19 cases the day before the pre-shutdown weekend (both the raw count and the number per 100,000), and the average growth (change in the inverse hyperbolic sine) of cases in the week preceding the last weekend. “Demographic Controls” are total population; fraction of population in the bins 60-69, 70-79, and over 80; fraction African American; and median household income.
 *p=0.10 **p=0.05 ***p=0.01

Figure 2
First-Stage and Reduced Form



Note: Each panel shows a partial correlation plot of rainfall on the weekend before the statewide lockdown against either the percentage of people leaving home on that weekend (left-hand panel) or total cases per 100,000 as of 14 days after the lockdown. We calculate residuals from a regression of both X and Y variable on state fixed-effects, historical rainfall, current and historical temperature, and baseline case controls. We define bins based on residualized rainfall. Each dot shows the average residualized outcome within the bin, and the line shows the linear prediction. The histogram shows the number of observations that fall into each bin.

ture, a one-unit increase in our measure of rainfall causes a 0.4 percentage point decrease in the number of people who leave home. The F-statistic is 11.68, well above conventional measures of instrument strength.

Columns 2—6 explore what activities become less prevalent because of rainfall and because people are staying home. One concern might be that although some people stay home because of the rain, those who do leave will pack into bars and restaurants instead of visiting the outdoors. Column 2 shows that the average exposure, based on how many people visit small indoor venues, declines by 0.87 percentage points relative to its level the first weekend of March (prior to any lockdown). Column 3 shows that our measure of large gatherings declines by 1.7 percentage points relative to early March.

Is the impact of rainfall on the prevalence of COVID-19 driven more by reducing local transmission, or by reducing the spread of the virus over long distances and across counties? Columns

4 and 5 measure the impact on the percentage of people leaving home and traveling a short or long distance (based on whether they traveled more than 16 kilometers from home). The estimates suggest a larger impact on long distance travel (especially compared to the mean). Column 6 shows that a one-unit increase in rainfall causes a 0.27 percentage point decrease in the fraction of people in the county who had not been there in the previous two weeks, suggesting a sizable decline in cross-county travel.

Reduced-Form and Two-Stage Least Squares: Panel B of Table 1 shows estimates of the reduced-form impact of rainfall on COVID-19 cases and deaths per 100,000 at endline, which these regressions define as 14 days after the statewide lockdown. Column 1 shows that a 1 unit increase in rainfall on the weekend before lockdown lowers the number of cases at endline by 6.7 per 100,000. Columns 2 and 3 show that controlling for baseline prevalence and demographics tightens the standard errors without substantially changing the estimates. Columns 4–6 imply that the reduction in cases translates to a reduction in deaths, as well. A 1 unit increase in rainfall causes a 0.5 to 0.7 per 100,000 reduction in the death rate.

Figure 2 shows a partial correlation plot of the first-stage and reduced form of the regression in Column 2 (which includes baseline case controls). The plot illustrates how rainfall on the last weekend before the state-wide lockdown lowers both the percentage of people leaving home (left-hand panel) and the number of cases at endline (right-hand panel). The plot shows that our estimates are not driven by outliers, and that both relationships are approximately linear.

Under the assumption that rainfall only affects disease transmission through its impact on early social distancing, the two-stage least squares estimate—the ratio of the reduced-form and first-stage coefficients—gives the causal impact of early social distancing on COVID-19 cases and deaths. Panel C of Table 1 presents these estimates. All three specifications have a strong first-stage, with the F-statistic on the excluded instrument (weekend rainfall) varying from 11 to 18. The basic specification, which has no controls, is relatively noisy and statistically insignificant.

But after controlling for baseline case controls the standard errors become tight enough to make the estimates highly significant (Columns 2 and 5). The final specification additionally controls for county demographics, which makes little difference in size or significance of the estimates (Columns 3 and 6). Indeed, all three specifications produce near-identical estimates.

A 1 percentage point increase in the number of people leaving home on the weekend before the shutdown causes an additional 13 cases and 1.3 deaths per 100,000.

The size of these estimates relative to the mean of the outcome may seem surprising. As we discuss in Section 3.3, the average impact represented by these estimates is misleading because it is in large part driven by changes in the probability of large outbreaks. Clearly it is not the case that a 4 percentage point decrease in people leaving home would have eradicated the disease across the country. It is more accurate to say that it would marginally reduce the probability of a catastrophic outbreak.¹¹ Finally, it is not clear whether the proper benchmark is the population that is infected at endline or the population that is susceptible to infection. In the latter case the reference group is roughly the entire population, in which case our results imply that an additional 1 percentage point of people leaving home causes an additional 0.013 percentage points of the population to become infected.

3.2 Comparative Dynamics in Counties with Less Early Social Distancing


Table 2 gives a relatively limited picture of the trajectory of cases because all outcomes are measured at the fixed horizon of 14 days after the statewide lockdown. One advantage of our research design is that we can estimate the comparative dynamics of case rates between counties that quasi-randomly practiced different levels of early social distancing. Using the same specification as Column 2 of Table 1, we estimate the impact on cases per 100,000 2 days after the lockdown, 4 days after, and so on for every horizon $h = 2, 4, 6, \dots, 18$. Figure 3 plots each coefficient against h . The estimated impact appears to increase linearly over time with no sign of leveling off within the horizon available to us.¹² The figure suggests the impact of a one-time difference in early social distancing is surprisingly long-lived.

We find no evidence, however, that the growth *rate* of cases increases because of more people leaving home on the last weekend (see Appendix A.4). That is not surprising because the natural experiment induces some counties to begin early social distancing just before all counties go

¹¹ Another possibility is that the types of activities deterred by rainfall are the riskiest—for example, visits by family members to skilled nursing facilities. It is also possible that there are substantial spillovers across counties, and that each person staying home actually reduces the risk of spreading cases across several counties.

¹² At longer horizons we would start to lose states because our case count data ends 18 days after the last state in our sample to go on lockdown.

Figure 3
 The Excess Case Count in Counties with Less Early Distancing
 Continues to Increase Even 18 Days after Lockdown



Note: Using total cases per 100,000 at each horizon $h = 2, 4, 6, \dots, 18$ we estimate the two-stage least squares coefficient controlling for baseline case controls (analogous to Column 2 of Panel C, Table 1). Each coefficient is from a separate regression (and the regression at $h = 14$ is identical to that reported in Table 1).

uniformly into lockdown. The effect is analogous to quasi-randomly inducing some counties to begin lockdown with a larger infected population. As long as this difference in initial population does not affect how carefully the lockdown is observed, it will rescale the case count without affecting the transmission rate.¹³


3.3 Distributional Impact: Early Social Distancing Lowers the Chance of Right-Tail Outcomes

Given the nature of exponential growth, local COVID-19 outbreaks may quickly die down or rapidly spiral out of control. That feature of transmission dynamics suggests early social distancing may have extended rather than shifting the distribution. We test for the impact on the full distribution by defining dummies for whether the endline number of cases per 100,000 is

¹³ If endline case count is $Y_T = Y_0 \exp(gT)$, our natural experiment is analogous to increasing Y_0 .

Figure 4

Counties with Less Social Distancing are More Likely to Have Very Large (Right-Tail) Outbreaks



Note: We estimate the impact across the distribution of outcomes. Each point and confidence interval is the two-stage least squares estimate of the impact of early social distancing on the probability of having endline cases per 100,000 greater than the percentile or absolute number indicated on the horizontal axis. Each estimate controls for baseline case rate, count, and growth (analogous to Column 2 of Panel C, Table 1).

greater than each decile of the distribution. We estimate Equation 3 using these dummies as the outcomes (using the specification with baseline case controls). This procedure is analogous to testing how the inverse cumulative distribution function is shifted by a 1 percentage point reduction in early social distancing.

The left-hand panel of Figure 4 plots the estimates with their 90 and 95 percent confidence intervals. The figure suggests that although the estimated impact becomes positive around 0.4 (meaning less early distancing increases the probability of being above the 40th percentile), the effect only becomes significant at 0.7. That suggests early distancing is lowering the probability of a right-tail outbreak. The most precise estimate is the last. A 1 percentage point increase in the number of people leaving home on the weekend before lockdown causes a 2 percentage point increase in the probability of an outbreak that puts the county in the top 10 percent of the distribution. The right-hand panel clarifies just how large these right-tail events are. This panel is analogous to the first one, but it defines dummies based on having an endline case rate above

some absolute cutoff. The size and significance peaks at 100 cases per 100,000, a very large case count.

The results suggest early social distancing worked less by causing a moderate reduction in cases than by reducing the chance of a big outbreak. This result may be consistent with several recent studies that find that COVID-19 has a very low dispersion factor, meaning small groups of “superspreaders” are responsible for the vast majority of cases (Kupferschmidt, 19 May 2020). Endo et al. (2020) estimate using a mathematical model that as few as 10% of initially infected people may be responsible for as much as 80% of subsequent cases. Miller et al. (2020) find a similar result when they use genome sequencing to trace the virus’s spread across Israel. If early social distancing marginally reduces the probability a superspreader begins a transmission chain, it could explain why our estimates are driven by changes in the number of large outbreaks. Regardless of the cause, our estimates imply that most counties that began distancing sooner had little benefit, but those that did benefit did so tremendously.

3.4 Robustness and Threats to Validity

In the appendix we run several other tests:

Balance: One concern is that rainfall, even after controlling for state fixed effects, historical rainfall, and current and historical temperature, is not truly exogenous. We show in Appendix A.1 that rainfall is uncorrelated with baseline measures of COVID-19 prevalence and county demographic characteristics.

Heterogeneity: We show in Appendix A.2 that there is little evidence of heterogeneous impacts by baseline case levels, baseline case growth, the time between the last weekend and the start of the statewide lockdown, and a host of demographic characteristics. This seems largely a consequence of not having enough data to generate a strong first stage when splitting the sample or identifying an interaction as well as a direct effect. There is some slight evidence that early social distancing has less of an impact in counties with an older population, though the mechanism for that result is uncertain.

Outliers: Given that Section 3.3 shows the effect comes largely from changes in the likelihood of right-tail events, one may worry that the entire estimate is driven by a few outliers. Appendix

A.3 shows that Winsorizing the very largest outcomes still yields significant effects. Although the top of the distribution does drive the result, it is a genuine distributional impact rather than a handful of fluke outliers.

Other Outcomes: Though endline cases and deaths per 100,000 is the most logical outcome (see Section 2.1), we show in Appendix A.5 the results are qualitatively similar if we instead use raw counts and the log of endline cases and deaths per 100,000.¹⁴

Measurement Error in COVID-19 Prevalence: One inevitable challenge to any study of COVID-19 is that the true number of cases far exceeds reported cases. One strength of our design is that rainfall is unlikely to be correlated with local testing capacity, making it unlikely that our result is spuriously driven by non-classical measurement error. However, we cannot rule out that counties with larger outbreaks are more aggressive in testing. Then *any* variation that reduces COVID-19 cases rates, be it rainfall or a hypothetical randomized controlled trial, would find accentuated impacts. We acknowledge that this caveat applies to our study as it does to any other.

Local Policy Response: One concern is that even if rainfall is exogenous, local governments might respond to either social distancing or (more likely) rising numbers of cases by instituting their own emergency orders or lockdowns. Our estimates might reflect not just the initial shock to social distancing but the policy response triggered by that shock. Although such a response is possible, it is likely to be a countervailing response. Local officials would likely loosen restrictions wherever case counts are low and vice-versa.¹⁵ That would, if anything, bias our estimates towards zero. Nevertheless we show in Appendix A.7 that controlling for a dummy for whether the county has any policy restriction by the end of the 14 day horizon of our regressions does not change the results.

Direct Impact of Weather: Some news reports and health experts have observed that warmer countries (e.g. Singapore and South Korea) have been more successful in controlling outbreaks than more temperate ones (e.g. the U.S. and Western Europe). That has led to a theory that tem-

¹⁴ To be precise, we estimate a Poisson Maximum Likelihood estimator using Equation 2 as the link function. Unlike simply taking the log, the Poisson estimator is consistent even though endline cases and deaths equal zero in many counties (Silva and Tenreiro, 2006).

¹⁵ Brzezinski et al. (2020) find that states where people are already social distancing of their own accord are less likely to impose a lockdown.

perature may directly affect virus transmission (e.g Sajadi et al., 2020). If the weather directly affects transmission it could violate the single-channel assumption needed for a valid instrument.

We find no evidence for a link between transmission and temperature on the last weekend in our county-level results. Regardless, all of our specifications control for temperature, making it unlikely to be driving our results. Some reports have also suggested humidity may separately affect transmission.¹⁶ Though the evidence for this is limited, we test for whether humidity is driving the results. If the impact of rainfall on cases and deaths were through its correlation with humidity rather than its impact on social distancing, we would expect that the reduced-form impact of rainfall on cases and deaths would vanish after controlling for humidity. But we show in Appendix A.6 that the reduced-form coefficient is essentially unchanged.¹⁷ Other links are possible but not yet well substantiated. It is possible that sunlight, through ultraviolet radiation, reduces virus spread. If that is true it would bias our estimates towards rainfall *increasing* the number of COVID-19 cases.

That said, we cannot categorically rule out that rainfall has some unanticipated impact or interaction with the environment. Given what is currently known about the virus and the nature of our own results, we believe these effects to be second-order compared to the direct impact on human behavior.

4 Directions for Future Research

Our results suggest that a marginal increase in social distancing a few days before a statewide lockdown has persistent effects two to three weeks later. One interpretation is that policy makers wishing to (re)institute a lockdown would reap surprisingly large gains from moving more quickly.

Our results come with a few caveats. First, as noted above we cannot categorically rule out that rainfall directly affects COVID-19 transmission through some as-yet unknown mechanism. Second, the type of social distancing induced by rainfall may differ from that induced by a gov-

¹⁶ Luo et al. (2020) is one example, though they actually find that humidity predicts *lower* transmission.

¹⁷ Since we only have humidity data for 60% of the sample, controlling for it directly in all specifications (as we do with temperature) would be too costly for precision.

ernment order.¹⁸ Finally, the context of our natural experiment—the weekend before a statewide lockdown—was one in which many people were already voluntarily social distancing. Policy-makers may face a different context when deciding on whether to begin a future lockdown. We leave disentangling these mediating factors to future research.

¹⁸ For example, state and county lockdowns have sparked protests and political opposition, while a rainy weekend presumably would not.

References

- Bell, Michael**, “How do I calculate dew point when I know the temperature and the relative humidity?,” Accessed 17 May 2020. <https://iridl.ldeo.columbia.edu/dochelp/QA/Basic/dewpoint.html>.
- Brzezinski, Adam, Guido Deiana, Valentin Kecht, and David Van Dijcke**, “COVID Economics,” *Covid Economics*, 2020, 7, 115.
- Center for Systems Science and Engineering (Johns Hopkins University)**, “COVID-19 Data Repository,” 2020.
- Chin, Alex W H, Julie T S Chu, Mahen R A Perera, Kenrie P Y Hui, Hui-Ling Yen, Michael C W Chan, Malik Peiris, and Leo L M Poon**, “Stability of SARS-CoV-2 in different environmental conditions,” *The Lancet Microbe*, 2020, 1 (1), e10.
- Corona Data Scraper**, “Corona Data Scraper,” 2020.
- Courtemanche, Charles, Joseph Garuccio, Anh Le, Joshua Pinkston, and Aaron Yelowitz**, “Strong Social Distancing Measures In The United States Reduced The COVID-19 Growth Rate,” *Health Affairs*, 2020.
- Couture, Victor, Jonathan Dingel, Allison Green, Jessie Handbury, and Kevin Williams**, “Device Exposure Indices,” 2020. <https://github.com/COVIDExposureIndices/COVIDExposureIndices>.
- , —, —, —, and —, “Location exposure indices,” 2020. <https://github.com/COVIDExposureIndices/COVIDExposureIndices>.
- Endo, Akira, Sam Abbott, Adam J Kucharski, Sebastian Funk et al.**, “Estimating the overdispersion in COVID-19 transmission using outbreak sizes outside China,” *Wellcome Open Research*, 2020, 5 (67), 67.
- Fowler, James H., Seth J. Hill, Nick Obradovich, and Remy Levin**, “The Effect of Stay-at-Home Orders on COVID-19 Cases and Fatalities in the United States,” *medRxiv*, 2020.
- Institute for Health Metrics and Evaluation**, “COVID-19 Projections,” 2020. <https://covid19.healthdata.org/united-states-of-america>. Accessed 17 April 2020.
- Kupferschmidt, Kai**, “Why Do Some COVID-19 Patients Infect Many Others, Whereas Most Dont Spread the Virus at All?,” *Science*, 19 May 2020.

- Luo, Wei, Maimuna S Majumder, Dianbo Liu, Canelle Poirier, Kenneth D Mandl, Marc Lipsitch, and Mauricio Santillana**, “The Role of Absolute Humidity on Transmission Rates of the COVID-19 Outbreak,” *medRxiv*, 2020.
- Madestam, Andreas, Daniel Shoag, Stan Veuger, and David Yanagizawa-Drott**, “Do Political Protests Matter? Evidence from the Tea Party Movement*,” *The Quarterly Journal of Economics*, 09 2013, 128 (4), 1633–1685.
- Manson, Steven, Jonathan Schroeder, David Van Riper, and Steven Ruggles**, “IPUMS National Historical Geographic Information System: Version 14.0 [Database].,” 2019. Minneapolis, MN: IPUMS. <http://doi.org/10.18128/D050.V14.0>.
- Menne, Matthew J, Imke Durre, Russell S Vose, Byron E Gleason, and Tamara G Houston**, “An Overview of the Global Historical Climatology Network-Daily Database,” *Journal of Atmospheric and Oceanic Technology*, 2012, 29 (7), 897–910.
- Miller, Danielle, Michael A Martin, Noam Harel, Talia Kustin, Omer Tirosh, Moran Meir, Nadav Sorek, Shiraz Gefen-Halevi, Sharon Amit, Olesya Vorontsov, Dana Wolf, Avi Peretz, Yonat Shemer-Avni, Diana Roif-Kaminsky, Na’ama Kopelman, Amit Huppert, Katia Koelle, and Adi Stern**, “Full Genome Viral Sequences Inform Patterns of SARS-CoV-2 Spread Into and Within Israel,” *medRxiv*, 2020.
- Pei, Sen, Sasikiran Kandula, and Jeffrey Shaman**, “Differential Effects of Intervention Timing on COVID-19 Spread in the United States,” *medRxiv*, 2020.
- Rodriguez, Olga R.**, “Fast Decisions in Bay Area Helped Slow Virus Spread,” *Associated Press*, 19 April 2020. <https://apnews.com/10c4e38a0d2241daf29a6cd69d8d7b43>.
- SafeGraph**, “SafeGraph Social Distancing Metrics, Version 2,” 2020. <https://docs.safegraph.com/docs/social-distancing-metrics>.
- , “SafeGraph Weekly Patterns Metrics, Version 1,” 2020. <https://docs.safegraph.com/docs/weekly-patterns>.
- Sajadi, Mohammad M, Parham Habibzadeh, Augustin Vintzileos, Shervin Shokouhi, Fernando Miralles-Wilhelm, and Anthony Amoroso**, “Temperature and Latitude Analysis to Predict Potential Spread and Seasonality for COVID-19,” *Preprint, Available at SSRN 3550308*, 2020.

Silva, JMC Santos and Silvana Tenreyro, “The Log of Gravity,” *The Review of Economics and Statistics*, 2006, 88 (4), 641–658.

The National Association of Counties, “County Explorer,” Accessed 22 May 2020.
<https://ce.naco.org/?dset=COVID-19&ind=State%20Declaration%20Types>.

A Empirical Appendix

A.1 Balance Tests

Table 2
First Stage and Balance

Panel A				
	(1)	(2)	(3)	(4)
	% Leaving Home	Baseline Cases	Baseline Cases/100k	Baseline Case Growth
Rainfall	-0.432*** (0.126)	-3.135 (4.377)	-0.054 (0.324)	0.005 (0.004)
Counties	1946	1946	1946	1946
Clusters	139	139	139	139
F-stat: Rainfall	11.68	0.51	0.03	1.44
State FEs	X	X	X	X
Avg. Rain	X	X	X	X
Temperature	X	X	X	X
Baseline Case Controls				
Demographic Controls				

Panel B				
	(1)	(2)	(3)	(4)
	Baseline Deaths	Baseline Deaths/100k	Baseline Death Growth	Population
Rainfall	-0.174 (0.172)	-0.033 (0.032)	0.005 (0.004)	4258.716 (11014.234)
Counties	1946	1946	1946	1946
Clusters	139	139	139	139
F-stat: Rainfall	1.03	1.04	1.44	0.15
State FEs	X	X	X	X
Avg. Rain	X	X	X	X
Temperature	X	X	X	X
Baseline Case Controls				
Demographic Controls				

Panel C					
	(1)	(2)	(3)	(4)	(5)
	Median HH Income	Fraction 60-69	Fraction 70-79	Fraction over 80	Fraction Black
Rainfall	3148.256 (8802.852)	0.001* (0.001)	0.001 (0.001)	-0.000 (0.000)	-0.001 (0.002)
Counties	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139
F-stat: Rainfall	0.13	2.79	1.10	0.01	0.22
State FEs	X	X	X	X	X
Avg. Rain	X	X	X	X	X
Temperature	X	X	X	X	X
Baseline Case Controls					
Demographic Controls					

Note: We estimate Equation 1 using the basic specification on each outcome. Standard errors are clustered as in Table 1.
*p=0.10 **p=0.05 ***p=0.01

A.2 Heterogeneity

Covid Economics 24, 1 June 2020: 76-110

Table 3
Heterogeneity By Interaction Terms

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline Cases	Baseline Case Growth	Days Until Lockdown	Fraction over 80	Fraction Black	Median HH Income
Main Effect	16.482** (8.300)	16.981** (6.890)	8.614 (17.789)	15.350*** (5.102)	11.772** (5.133)	18.513* (9.741)
Interaction	-0.373 (0.812)	-19.173 (27.065)	1.587 (4.471)	-16.891*** (5.432)	27.187 (22.896)	-0.000 (0.000)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
K-P Stat.	0.14	6.49	1.73	9.64	5.77	1.60
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls	X	X	X	X	X	X

Note: Each regression adds $C_t \times D_t$ to Equation 3, instrumenting for it with $C_t \times R_t$. The column header gives the variable used for C_t . The outcome in all regressions is ending cases per 100,000.
*p=0.10 **p=0.05 ***p=0.01

Table 4
Heterogeneity by Splitting the Sample

Panel A						
	Baseline Cases		Baseline Case Growth		Days Until Lockdown	
	(1) Below	(2) Above	(3) Below	(4) Above	(5) Below	(6) Above
% Leaving Home	7.824 (5.830)	22.356* (12.145)	8.257* (4.655)	32.642 (21.240)	16.916** (8.328)	13.614** (6.646)
Counties	998	948	1450	496	1055	891
Clusters	123	105	133	85	79	93
First-Stage F	3.97	8.54	6.37	6.24	8.37	6.23
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls	X	X	X	X	X	X

Panel B						
	Fraction over 80		Fraction Black		Median HH Income	
	(1) Below	(2) Above	(3) Below	(4) Above	(5) Below	(6) Above
% Leaving Home	24.445** (12.197)	7.941* (4.688)	6.141 (7.875)	18.321*** (6.835)	11.647 (9.222)	62.631 (44.183)
Counties	973	973	973	973	973	973
Clusters	119	112	118	89	118	107
First-Stage F	6.28	9.65	3.46	21.11	3.25	2.29
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls	X	X	X	X	X	X

Note: The sample is split based on whether a county is above or below the median value of the variable given in the header. "Days Until Lockdown" is the difference between the date of statewide lockdown and the first day of the final pre-shutdown weekend.

*p=0.10 **p=0.05 ***p=0.01

Table 5
Winsorized Outcomes

Panel A: Reduced-Form						
	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
	.01	.02	.04	.01	.02	.04
Rainfall	-4.088*** (1.143)	-3.027*** (0.891)	-2.199*** (0.712)	-0.160*** (0.059)	-0.130** (0.052)	-0.077*** (0.038)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
Outcome Mean	54.18	51.93	48.91	1.66	1.61	1.45
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls						

Panel B: Two-Stage Least Squares						
	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
	.01	.02	.04	.01	.02	.04
% Leaving Home	9.730*** (3.407)	7.204*** (2.677)	5.234** (2.155)	0.381** (0.157)	0.310** (0.136)	0.184* (0.097)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
First-Stage F	16.54	16.54	16.54	16.54	16.54	16.54
Outcome Mean	54.18	51.93	48.91	1.66	1.61	1.45
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls						

Note: Outcomes are Winsorized at at the percentiles shown in the column header.
*p=0.10 **p=0.05 ***p=0.01

A.3 Winsorized Outcomes

Table 6
Growth Rates

Panel A: Reduced-Form

	Average Growth Rate in Cases			Average Growth Rate in Deaths		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-0.000 (0.002)	-0.000 (0.002)	0.000 (0.002)	-0.000 (0.001)	-0.001 (0.001)	-0.001 (0.001)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
Outcome Mean	0.10	0.10	0.10	0.04	0.04	0.04
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Panel B: Two-Stage Least Squares

	Average Growth Rate in Cases			Average Growth Rate in Deaths		
	(1)	(2)	(3)	(4)	(5)	(6)
% Leaving Home	0.000 (0.004)	0.001 (0.004)	-0.000 (0.004)	0.000 (0.003)	0.001 (0.002)	0.001 (0.002)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
First-Stage F	11.68	16.54	17.80	11.68	16.54	17.80
Outcome Mean	0.10	0.10	0.10	0.04	0.04	0.04
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: We calculate the growth rate as the average of the day-to-day change in the inverse hyperbolic sine of cases and deaths from the pre-shutdown weekend through 14 days after the statewide lockdown.

*p=0.10 **p=0.05 ***p=0.01

A.4 Growth Rates

Table 7
Alternative Outcome: Raw Endline Counts of Cases and Deaths

Panel A: Reduced-Form

	Endline Cases			Endline Deaths		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-58.417 (53.157)	-31.696** (12.158)	-34.127*** (11.987)	-7.818 (7.393)	-4.886* (2.833)	-4.293* (2.175)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
Outcome Mean	164.71	164.71	164.71	7.21	7.21	7.21
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Panel B: Two-Stage Least Squares

	Endline Cases			Endline Deaths		
	(1)	(2)	(3)	(4)	(5)	(6)
% Leaving Home	135.222 (141.248)	75.448** (33.577)	85.438** (35.271)	18.097 (19.153)	11.631* (6.924)	10.749* (5.850)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
First-Stage F	11.68	16.54	17.80	11.68	16.54	17.80
Outcome Mean	164.71	164.71	164.71	7.21	7.21	7.21
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: The outcomes are endline cases and deaths without adjustment for county population.
*p=0.10 **p=0.05 ***p=0.01

A.5 Other Outcomes

Table 8
Alternative Outcome: “Log” of Cases and Deaths per 100,000

	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-0.110*** (0.033)	-0.078*** (0.023)	-0.076*** (0.023)	-0.306*** (0.090)	-0.233*** (0.084)	-0.164*** (0.047)
Counties	1946	1946	1946	1942	1942	1942
Clusters	139	139	139	139	139	139
Outcome Mean	58.12	58.12	58.12	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: We estimate a Poisson Maximum Likelihood model that assumes the outcome equals the exponential of the specifications in the main text. This is in concept similar to regressing the log of the outcome on each specification, but the Poisson estimate is consistent even though the outcome equals zero for many counties. We are unable to estimate second-stage IV coefficients because the GMM estimator is unable to converge to estimates of so many state fixed-effects.

*p=0.10 **p=0.05 ***p=0.01

Table 9
The Impact of Rainfall on Cases/Deaths Is Unchanged When We Control for Humidity

	Endline Cases/100k			Endline Deaths/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-6.132*** (1.705)	-7.307*** (2.069)	-6.870*** (2.040)	-0.707** (0.353)	-0.707** (0.353)	-0.703* (0.373)
Rel. Humidity			-10.416 (17.458)			-0.113 (0.993)
Counties	1946	1131	1131	1131	1131	1131
Clusters	139	135	135	135	135	135
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls	X	X	X	X	X	X
Demographic Controls						
Sample	Full	Humidity	Humidity	Full	Humidity	Humidity

Note: The “Full” sample is the sample used in the main text. The “Humidity” sample is the subsample of counties for which we have data on dew point.
*p=0.10 **p=0.05 ***p=0.01

A.6 Humidity

We use data from the Global Surface Summary of Day. The dataset does not record humidity but does record dew point temperature. We calculate relative humidity using an approximation of the Clausius-Clapeyron equation (Bell, Accessed 17 May 2020).¹⁹

$$\begin{aligned}
 E &= E_0 \exp \left\{ \frac{L}{R_v} \left(\frac{1}{T_0} - \frac{1}{T_d} \right) \right\} \\
 E_s &= E_0 \exp \left\{ \frac{L}{R_v} \left(\frac{1}{T_0} - \frac{1}{T} \right) \right\} \\
 H_R &= 100\% \times \frac{E}{E_s} = 100 \exp \left\{ \frac{L}{R_v} \left(\frac{1}{T} - \frac{1}{T_d} \right) \right\}
 \end{aligned} \tag{4}$$

where the terms in (4) are

- H_R : relative humidity
- T : Temperature (in Kelvin)
- T_d : Dew Point Temperature (in Kelvin)
- $\frac{L}{R_v} = 5423K$

¹⁹ In a few cases the calculation gives a number greater than 100%, likely because a measurement error in the

We average dew point for all stations within a county and calculate the inverse hyperbolic sine of the dew point on the last weekend before statewide lockdown.

We estimate the reduced-form of our specification

$$Y_i = \kappa_s + \omega R_i + \phi_1 \bar{R}_i + \phi_2 T_i + \phi_3 \bar{T}_i + X_i \vartheta + v_i$$

which gives the direct impact of rainfall on the last weekend on cases and deaths. We see if the reduced-form coefficient $\hat{\omega}$ changes when we add dewpoint to the set of controls X_i . The specifications in Table 9 first show the reduced form coefficient for the entire sample. Since we only have humidity data for a subset of this sample, the next specification estimates the same reduced-form coefficient using the restricted sample. The final specification adds relative humidity. The reduced-form coefficient is essentially unchanged when we control for humidity.

Table 10
Controlling for Local Policy Response Does not Change the Results

Panel A: Reduced-Form						
	Endline Cases/100k			Endline Cases/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall	-6.253** (2.506)	-6.004*** (1.705)	-5.884*** (1.682)	-0.761 (0.502)	-0.660** (0.293)	-0.626** (0.252)
Counties	1904	1904	1904	1904	1904	1904
Clusters	134	134	134	134	134	134
Outcome Mean	57.11	57.11	57.11	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Panel B: Two-Stage Least Squares						
	Endline Cases/100k			Endline Cases/100k		
	(1)	(2)	(3)	(4)	(5)	(6)
% Leaving Home	15.686 (9.653)	14.596*** (4.852)	14.824*** (5.130)	1.660 (1.274)	1.383** (0.556)	1.344** (0.517)
Counties	1946	1946	1946	1946	1946	1946
Clusters	139	139	139	139	139	139
First-Stage F	11.68	16.54	17.80	11.68	16.54	17.80
Outcome Mean	58.12	58.12	58.12	2.05	2.05	2.05
State FEs	X	X	X	X	X	X
Avg. Rain	X	X	X	X	X	X
Temperature	X	X	X	X	X	X
Baseline Case Controls		X	X		X	X
Demographic Controls			X			X

Note: We define a dummy equal to 1 if the county has adopted some measure (emergency declaration, safer-at-home instruction, shutting down businesses) by the end of the horizon for our outcome, 14 days after the statewide lockdown. All regressions control for this dummy (in addition to the controls discussed in the main text).
*p=0.10 **p=0.05 ***p=0.01

A.7 Policy Response

Using dates on county-level policy responses from The National Association of Counties (Accessed 22 May 2020), we define a dummy for whether the county has put any social distancing measure (emergency declaration, safer-at-home instruction, shutting down businesses) before the date at which we measure the outcome (14 days after the statewide lockdown). Table 10 reports our reduced-form and two-stage least squares estimates after controlling for the policy response.

B Data Appendix

B.1 Measures of Social Distancing

Device-exposure index (DEX): The index is computed using cellular data from PlaceIQ. Daily exposure of a device is defined as the number of distinct devices that visit the commercial venues that the particular device visits that day. DEX is then calculated by averaging the exposure values for all devices in the sample in the geographical unit (e.g. county) on a particular day. The set of devices included in the calculation of DEX are those that pinged on at least 11 days over any 14-day period from November 1, 2019 through the date in question. The venues covered are mainly commercial venues (with the largest category being restaurants). The set of venues is restricted to those “small enough such that visiting devices are indeed exposed to each other.” The set excludes Nature and Outdoor, Theme Parks, Airports, Universities, as well as any location whose category is unidentified by PlaceIQ.

Location-exposure index: The LEX dataset is a daily matrix of 2018 counties in which each cell $[i, j]$ reports, among devices that pinged on a particular day in county j and pinged anywhere in the previous 14 days, the share of devices that pinged in county i at least once during the previous 14 days. The dataset is restricted to counties with reasonably large device samples. We assume that diagonal elements of the matrix represent the fraction of cellphones pinged in a particular county that belong to that county itself, and hence $1 - lex[i, i]$ represents the total fraction of devices pinged in county i that had not been in i during the prior 14 days.

B.2 COVID-19 Cases and Deaths

The data for county-level COVID-19 cases and deaths was extracted from two sources: (i) CoronaDataScraper project, and (ii) JHU COVID-19 daily cases and deaths repository. Both the sources are updated daily. While the JHU dataset is more comprehensive of the two, we identified several county-date combinations for which:

- There were missing observations, or
- Data was discontinued for subsequent time periods

There were 136 such counties identified for confirmed cases and 63 for number of deaths.²⁰ For the counties we check, we also corrected two additional errors:

- The cumulative number of cases (or deaths) decrease after the particular date, implying a negative growth rate in cases (or deaths)
- The cases were reported with a lag of more than 1 day

We start with confirming the first reported case for the aforementioned 136 counties. This is important, since in some cases a presumed case was erroneously reported as the first confirmed case, or an administrative error assigned a case from another county or state to the county in question (or person was a temporary resident). What we observe is one (incorrect) entry in the number of cases on a particular day and then no observations for multiple days after that.

While some counties have regular press releases or a daily updated dashboard to check the numbers for a particular day, for the others we rely on multiple news reports. We follow the same procedure for other cases in the panel where the cumulative numbers on the subsequent dates mysteriously reduce only to increase again. Links to the rectification provided by the County Public Health Department as well as the news reports have been provided in the dataset. In cases where the county corrected the numbers but an associated press release was not found, we rely on multiple local news reports for the dates in question. We follow the same procedure for counties which did not have any confirmed cases but the dataset recorded one.

For randomly missing observation on particular dates, we look at the county public health department daily releases and dashboard charts, or the state public health department daily status updates for counties, and finally if there is a lack of information from both sources, we look at reports from the local media. Some state public health departments also provide a disclaimer attributing missing data for certain counties to lag in time between testing and reporting (e.g. Jeff Davis County, Georgia). For these county-date pairs, we rely solely on multiple local news reports that confirm the number of cases on that date.

We follow the above steps for correcting the cumulative number of deaths decreasing over time.

²⁰ A subsequent release of the JHU data corrected 33 of the case count errors.

Spatial distancing: air traffic, Covid-19 propagation, and the cost efficiency of air travel restrictions¹

Olivier Hubert²

Date submitted: 26 May 2020; Date accepted: 27 May 2020

We analyse how air traffic across countries contributed to the propagation of COVID-19 by fitting a Spatial Durbin-Watson model adapted to local projections. Such a model explicitly accounts for spatial dependence of observations and allows to track the effect of domestic and foreign new infections over time. Our estimates show that air travel-induced cases amount to 8-9% of all confirmed cases on average, and that these infections from abroad came in two waves: in mid-March and the fourth week of March. We also evaluate that air travel restrictions had a marked impact in reducing the progression of the pandemic from April onward. Closing all air traffic 4 weeks earlier could have prevented between 7,000 and 7,800 deaths. Based on standard values of a statistical life and on the latest estimates of GDP loss induced by air travel restrictions, we conclude that, just as social distancing, spatial distancing might be a cost-effective way to tackle COVID-19 in the short run.

1 The views expressed in this paper are those of the authors and do not necessarily reflect those of the OECD or the Banque de France.

2 Economist, UNamur.

Copyright: Olivier Hubert

1 Introduction

On December 31, 2019, Wuhan hospitals reported 27 cases of pneumonia of unknown origin. Less than two months later, COVID-19 had reached 5 continents and caused more than 85,000 confirmed cases worldwide, including 3,000 deaths. At the time of writing, more than 5.4 million cases of COVID-19 had been reported in more than 188 countries and territories, with a global death toll of about 350,000.¹ Such unprecedented speed and breadth of this propagation leave little doubt that air travel was key in spreading the virus.

In this paper, our objective is twofold. First, we estimate the contribution of air traffic to the spread of COVID-19 in a panel of 48 countries (OECD economies and economies that host the 50 largest airports worldwide). Second, based on our estimates, we provide a first-order cost-benefit analysis of air travel restrictions.

To achieve the first objective, we build a dynamic spatial autoregressive model that explicitly account for the air travel dependence of COVID-19 propagation in order to measure the effect of new cases across both time and space. Our study leverages two data sources. For the panel data on COVID-19 cases, we draw from the Johns Hopkins COVID-19 repository. To measure air travel intensity, we rely on the latest data available from *OpenFlights.org* on weekly bilateral commercial flight connections between all airport-country pairs.

As expected, our results indicate that air traffic with infected foreign regions is correlated with domestic infection. More interestingly, our model concludes that air traffic has a quantitatively large influence in explaining the prevalence of COVID-19 across countries. After two weeks, foreign new cases generate an additional 21 percentage points to the observed growth rate of infections, against 43 percentage points for new domestic infections. Over our estimation period, air travel accounts for 8-9% of reported cases, or about 108,000-121,000 cases. Based on a case fatality rate of 6.47%, air travel represents about 7,000-7,800 deaths.

Our assessment for the second objective relies on two types of data. On the one hand, we use on the OECD's evaluation of the initial impact of COVID-19 containment measures on economic activity (OECD, 2020a). On the other hand, we take standard estimates of the value of a statistical life (e.g. Viscusi's (2015) and Viscusi, and Masterman's (2017) meta-analyses). Our calculation shows that restricting air travel 4 weeks earlier would have largely passed a cost-benefit analysis, as the monetary value of saved lives would have amounted to USD 24-75 billion while the GDP cost associated with shutting down air transportation 4 weeks earlier would have been around USD 24 billion. Just like social distancing, *spatial distancing* might be a cost-efficient way to tackle COVID-19 in the short run.²

Beyond these quantification exercises, we propose a crucial methodological innovation by applying a Spatial Durbin-Watson model in a local projections framework. To the best of our knowledge, our paper is the only such occurrence. This modeling choice also allows us to account for the incubation of the virus.

There exists ample evidence that globalization and the steady growth in people's mobility contribute to spreading contagious diseases such as COVID-19 globally. Commercial air traffic is a major culprit, as it enables infected individuals to travel across regions in a matter of hours, rendering physical distance irrelevant to whether countries are spatially connected. For example, commercial flights contributed to spreading the H1N1 pandemic (Chaug-Ing Hsu and Hsien-Hung Shih, 2010). Recent papers discuss the role of air travel in the COVID-19 pandemic (Keitou, 2020; Krisztin, Piribauer and Wögerer, 2020; Lau et al., 2020; Craig et al., 2020; Zhuang et al., 2020). Lau et al. (2020) find a positive and significant correlation between passenger

¹See <https://ourworldindata.org/coronavirus>. Brazil reported its first confirmed case on February 26, making South America the 5th continent affected.

²Early assessments of the impact of social distancing measures suggest substantial economic benefits. See e.g. Greenstone and Nigam (2020) or Thunström and et al. (2020).

volume and domestic and international COVID-19 cases. Closely related to our study, Krisztin, Piribauer and Wögerer (2020) show that air travel is the main spatial driver of the propagation of COVID-19 at its onset, whereas other determinants such as geographical distance take over a few weeks into the pandemic. We depart from their analysis by providing a direct quantification of the burden of foreign new cases for the domestic epidemic. Finally, Keitou (2020) finds that more connected countries registered first infection cases significantly earlier, and that very early implementations of air travel restrictions were associated with a delayed onset of infections.

The paper is organized as follows. The following section presents the methodology and introduces the Spatial Durbin-Watson model. Section 3 provides more information about the data construction while Section 4 gathers the results. In particular, Sections 4.1 and 4.2 tackle the dynamic and time-varying properties of the spatio-temporal contagion. We provide two simulation exercises in Sections 4.3 and 4.4. The first displays how cases are distributed geographically after a new case is detected in a given country, while the second provides a back-of-the-envelope computation of the costs and benefits of restricting air travel. This simulation seems particularly well-suited at a time when governments contemplate re-opening aerial routes. Finally, Section 5 concludes.

2 Methodology

2.1 The SAR model

To model the spatial interdependencies across the countries considered, we use a Spatial Autoregressive (SAR) model. Spatial econometrics is particularly well-suited for the task at hand because it is designed to explicitly account for the dependence of geographical regions. As such, it expresses the observed data as a function of what can be observed in neighbouring countries. This feature is of particular interest when dealing with infectious diseases outbreaks.

In a globalized world where goods and people can cross borders in a matter of hours, the notion of geographic neighbour may be somewhat outdated. As a consequence, we prefer here the notion of 'strongly connected spatial units' to address the effect of spatial linkages. Tied to the use of spatial econometrics models is the description of the connectivity structure that summarizes the bilateral ties region i may have with region j .

The SAR model in its general dynamic panel form is expressed as:

$$cases_t^{total} = c + \rho W cases_t^{total} + X\kappa + \zeta cases_{t-1}^{total} + \epsilon_t \quad (1)$$

where $cases_t^{total}$ is a $N \times 1$ vector containing the (log) total number of confirmed cases at time in each country t , X_t may contain exogenous explanatory variables and W is the transmission matrix. The spatial weight matrix W is a $N \times N$ matrix that contains the number of flight routes between countries because we believe it best summarizes the contagion path followed by the SARS-CoV-2 virus. The first diagonal of W is filled with zeros as a country cannot be its own neighbour. The parameter ρ gives the strength of the spatial linkages. A value of $\rho = 0$ implies that neighbours do not affect the region considered. It is sufficient for spatial stability that $\rho \in (-1, 1)$ (LeSage and Pace, 2009). The matrix W is row-normalized such that $W cases_t^{total}$ is the weighted average of foreign cases where the individual weights are given by the share of inbound flight connections in the total number of flight connections.

Notice that Equation (1) cannot be estimated by OLS because of the presence of $cases_t^{total}$ on both sides

of the equal sign. We estimate this expression via well-known techniques that allow address this issue (see Elhorst, 2003 and LeSage, 2008). In the subsequent analysis, X is empty.

Such a model has been applied to the SARS-CoV-2 epidemic by Krisztin, Piribauer and Wögerer (2020) in a Bayesian framework where the spatial lag parameter and the spatial weight matrix are allowed to evolve over time. Their focus, however, does not lie on characterizing the value of the spatial parameter, but rather on determining which spatial diffusion process among flight intensity, bilateral trade, contiguity or free movement of people is the most relevant. They show that flight intensity explains spatial contagion the most until the third week of March. Past this date, the spatial weight parameter ρ is not statistically significant.

Other applications to Spanish and Chinese data include Orea and Álvarez (2020) and Guliyev (2020). These paper, however, focus on spatial transmission due to geographical proximity (contiguity) and do not address the dynamic feature of COVID-19 infections.

2.2 Spatial Local Projections

In addition to the static description of spatial spillovers in Equation (1), we provide a dynamic version that allows us to track the effect of new cases both across time but also across spatial units. To achieve this, we use the local projections framework (LP) first introduced by Jordà (2005) and popularized in macroeconomics by a host of authors due to its high tractability and robustness to model misspecification (see, for example Auerbach and Gorodnichenko ; 2012, 2017). Technically, such Impulse Response Functions (IRFs) consist in regressing the dependent variable at $t + h$, with h the horizon of the IRF, onto a set of lagged explanatory variables to prevent cyclicity and a *shock* variable at time t . The main advantage of LP is that the IRF and its inference can be read directly via the coefficient and standard error pertaining to the shock variable. In essence, local projections amount to a direct forecast of the dependent variable, as opposed to an iterative forecast that is common with Vector Autoregressions, for example.

The Spatial Local Projections take the form:³

$$cases_{t+h}^{total} - cases_{t-1}^{total} = \alpha + \beta cases_t^{new} + \beta_W W cases_t^{new} + \gamma cases_{t-1}^{total} + \gamma_W W cases_{t-1}^{total} + \rho W (cases_{t+h}^{total} - cases_{t-1}^{total}) + \eta_{t+h} \quad (2)$$

where $cases_t$ is the total number of confirmed cases at time t and the *shock* variable is the number of new cases $cases_t^{new}$. The lagged explanatory variables control for the level of the epidemic as it has become clear that the growth rate of infection is time-dependent. Both the prevalence ($cases_t^{total}$) and the incidence ($cases_t^{new}$) are expressed in log terms such that the dependent variable is a growth rate.

To the best of our knowledge, the only occurrence where local projections and spatial econometrics were merged in a consistent framework is Brady (2011). We extend his framework by taking full advantage of the spatial lag structure: current cases depend on foreign new cases but also on foreign lagged cases. This feature seemed important to us because of the incubation period of the virus. As such, this is the first time that a Spatial Durbin-Watson model is used in a local projections framework.

In terms of specification, we estimate the spatial impulse response functions as a Pooled panel over the whole sample period. Both total cases and new cases are expressed in log terms such that the impulse

³To favour readability, the h superscript is omitted. Each of the coefficients is allowed to vary with the horizon of the Impulse Response Function.

response functions should be read as the cumulative growth rate between $t-1$ and $t+h$.

3 Data

The infection data comes from the Johns Hopkins COVID-19 repository, which aggregates several data sources.⁴ The dataset contains daily information on total cases and new cases for more than 180 countries. We restrict our sample to OECD countries plus non-OECD countries with an airport among the 50 largest airports worldwide in terms of passengers traffic between 2016 and 2018. In total, we have a panel of 48 countries (see Annex A for the complete list).

The flight intensity matrix is constructed from the number of weekly direct inbound commercial flights between country pairs as referenced in *OpenFlights.org*. To better gauge the actual connection intensity between a country pair, we removed duplicate routes with a Code Share, i.e. flights operated by SN Brussels Airlines on a Lufthansa route, for example.

The time span for the analysis below covers March, 13 2020 to April, 8 2020. The starting point of our estimation is dictated by the heterogeneous start of the epidemic in each country. As such, we have to restrict our estimation to a period where every country in our sample has at least one case since we use logarithms of cases. For the ending date, we consider an extension of three weeks from the date at which at least 50% of the countries in the world have set travel bans from infected areas (March, 18 2020). We consider that such an extension prevents us from selecting too narrow a period (and thus cherry-picking the results) while at the same time ensuring that our spatial transmission matrix is sensible over the four weeks considered. This extended window also largely covers the incubation period of the virus.

4 Results

As a starting point, we provide results from Equation (1) in Table 1 estimated in its static form between March, 13 2020 and April, 8 2020. As the ρ coefficient is positive and statistically significant, we can already conclude to the presence of a positive spatial association due to flight intensity across country pairs: a higher number of cases in connected countries is associated with a higher number of cases domestically. However, we believe that a dynamic view brings additional insights on the evolution of the transmission of the Coronavirus across time. We address this feature in the subsequent sections.

Table 1: SAR model

Variable	Coefficient	t-stat	z-probability
c	0.18	8.72	<0.01
ζ	0.96	337.93	<0.01
ρ	0.03	6.41	<0.01
R ²	0.991		
Adj. R ²	0.991		

Note: the table reports the estimated coefficients of Equation (1). The first column indicates the value of the coefficient in the regression, the second the corresponding t-stat and the last column lists the corresponding p-value. The last two rows of the table report the R-squared and its adjusted counterpart.

⁴Dataset available at: <https://github.com/CSSEGISandData/COVID-19>

4.1 Dynamic spatial diffusion

Table 2 reports the coefficients of Equation (2) at horizon $h=15$.⁵ First of all, the goodness of fit of the regression is fairly high, with roughly 70% of the observed variance accounted for by our model. Second, all the variables are highly statistically significant. In particular, the spatial lag ρ is positive and therefore indicates a positive spatial association, confirming the static results. Third, the β_W coefficient is positive, which indicates that an increase in the number of cases abroad translates into an increase in cases domestically by 21%. This coefficient provides more information than the spatial lag because it allows us to quantify the total (direct and indirect) response of domestic cases to foreign new infections.

Table 2: Spatial Impulse Response Functions - horizon $h=15$

Variable	Coefficient	t-stat	z-probability
α	5.41	49.40	<0.01
β	0.43	20.17	<0.01
γ	-0.58	-30.72	<0.01
β_W	0.21	7.85	<0.01
γ_W	-0.41	-16.74	<0.01
ρ	0.18	6.24	<0.01
R ²	0.706		
Adj. R ²	0.670		

Note: the table reports the spatial impulse response of Equation (2) for $h=15$. The first column indicates the value of the coefficient in the regression, the second the corresponding t-stat and the last column lists the corresponding p-value. The last two rows of the table report the R-squared and its adjusted counterpart.


We present the results of the estimation of Equation (2) in Figure 1 and plot the response of the number of domestic confirmed cases following an increase in the domestic (in blue with triangle markers) and foreign cases (in red with circles). The shock takes the form of a one percent increase in the number of new cases and the response is expressed in terms of cumulative percentage change.

Several conclusions can be drawn from the exercise. First, both domestic and foreign new cases lead to a sizable increase in the number of domestic cases. As such this indicates that a high connectivity with infected regions strongly increases the domestic infection rate. Second, the relative influence of foreign cases is not marginal as the IRF for foreign cases is almost half of the one for domestic cases at horizon 15 and 25% at horizon $h=3$. The opening or closing of air traffic therefore has a tangible influence on the the speed of the epidemic.

It is possible to express the share of the variance of observed cases that is due to foreign new cases thanks to the forecast error variance decomposition (FEVD). Within the framework of local projections, we follow the method proposed by Gorodnichenko and Lee (2019). The procedure consists in two steps: first, we obtain the residuals from (2) as they represent the forecast error. Second, we regress the residuals onto a collection of shocks up to $t+h$. The R-squared of this regression gives the FEVD. As we have two shocks, we make sure that the sum of the R-squareds from each type of shocks (domestic or foreign) sums to the R-squared of the regression where both types of shocks explain the forecast error variance. The FEVDs for various horizons can be found in Table 3. At first, new cases, whether domestic or foreign, do not explain much of the observed growth rate of total cases. Their influence builds up with time in an unequal way. After two weeks, foreign new cases account for 5% of the observed growth rate of infections, against 12.5%

⁵Though not reported here, the spatial lag across other forecast horizons remains fairly constant. We see this feature as evidence of the robustness of the spatial association.

Figure 1: Impulse response function of domestic and foreign new cases



Note: the figure presents the impulse response functions of the number of domestic confirmed cases following an increase in the domestic cases (in blue with triangle markers) and foreign cases (in red with circles). The x-axis reports the horizon of the IRFs (in days). The IRFs are estimated in a log-log specification such that the scale (y-axis) should be read as the cumulative growth rate following a one percent increase in the number of new cases.

for domestic cases. The influence of foreign cases is therefore slightly below half of the influence of domestic infections.

Table 3: Forecast Error Variance Decomposition - Gorodnichenko and Lee (2019)

	Domestic cases	Foreign Cases
<i>Horizon: 0</i>	0.2	0.2
<i>1</i>	1.5	1
<i>7</i>	6	2.7
<i>15</i>	12.5	5.3


Note: the table reports the variance decomposition of Equation (2) for horizons 0, 1, 7 and 15 days. The computation of the FEVDs is based on Gorodnichenko and Lee (2019). The first column indicates the value of the variance of total domestic cases due to new domestic cases while the second column reports the share that can be attributed to new foreign cases. Inertia in the development of the epidemic captures the remainder of the variance.

4.2 Time-varying spatial diffusion

Until now, we focused our attention on a time-invariant version of the influence of air traffic on the propagation speed of COVID-19. In this section, we estimate Equation (2) for each time t and track the IRFs for each estimation point. We report the results in Figure 2. We can observe several important features. First, as in the time-invariant IRFs, domestic cases are more important than foreign cases, but imported cases are not marginal either. Time-varying impulse responses therefore follow the same general pattern and both methods are compatible. Second, while domestic cases experience a single peak in the first week of April (third week of March plus two weeks for the computation of the IRFs), foreign cases seem to have

experience two peaks nine days apart. Interestingly, these two peaks fit right before and after their domestic counterparts. Third, the influence of foreign cases dies out at the beginning of April. If one recalls that 50% of worldwide travel bans occurred in mid-March and that the incubation period of the SARS-CoV-2 virus is 2 to 15 days, we uncover evidence of the efficacy of such decisions. Our results corroborate Sekou (2020) who shows that countries that imposed a travel ban earlier experienced a slower pace of infection. This result also provides credence to the selection of air traffic as the main spatial diffusion process.

Figure 2: Time-dependent impulse response functions for domestic and foreign new cases




Note: the figure presents the impulse response functions of the number of domestic confirmed cases following an increase in the domestic cases (top panel) and foreign cases (bottom panel). The IRFs are estimated in a growth-log specification such that the scale (color bars) should be read as the cumulative growth rate (in %) following a one percent increase in the number of new cases.

As in the panel case, we can derive the FEVDs from the IRFs to express the share of the domestic cases

that is due to domestic of foreign factors at various horizons. On average, foreign cases amount to between 8% and 9% of the total observed cases in a given country.⁶ Due to the restriction that the γ and β coefficients (both domestic and foreign) be equal across spatial units, the difference in FEVDs across countries is limited. Spatial heterogeneity may nevertheless remain due to the country-specific spatial structure.

Figure 3: Time-dependent variance decomposition for foreign new cases



Note: the figure presents the variance of the number of domestic confirmed cases that is due to new foreign cases. The x-axis reports the horizon (in days) of the IRFs while the y-axis matches the time at which the IRFs are computed. The share is expressed in percent (right color bar).

The time-dependent variance decompositions closely match the IRFs and stress the existence of two waves of imported cases in mid-March and during the fourth week of March. From April onward, foreign cases lose their influence on the domestic epidemic.

4.3 Simulating the transmission of COVID-19 due to air travel

We provide a simulation of the spatial diffusion of COVID-19 due to air travel in Figure 4 derived from the panel estimation of Equation (2). We increase the number of new cases by one in China, Spain, Italy, Great-Britain and the United States (one at a time) and we map the response after 15 days in terms of confirmed cases in each of the 48 countries in our sample.


We can observe large heterogeneity in the responses depending on the country considered. First, an additional case in China usually remains in East Asia and South-East Asia, with the extent of the contagion limited to 0.4 new cases for each additional case. An additional case in the three European countries spreads to sensibly more countries, especially from the UK, though the size of the spillover is similar to the Chinese simulation. The US case is particularly worrisome as the strength of international transmission is much larger than the other countries. Indeed, one additional case in the United States translates into more than one additional case in Mexico, 0.8-1 supplementary case in Colombia and 0.6-0.8 case in Brazil and Canada.

We conclude from these simulations that exposure to well-connected travel hubs may prove particularly dangerous in the wake of an epidemic. The total size of the spillover is also of particular interest: each

⁶Anecdotal evidence suggests that this proportion is sensible: as of May, 22 2020, around 10% of Coronavirus cases in South Korea were imported.

additional case in China creates 3.6 cases worldwide, which is roughly the same as Italy (3.7). However, one additional case in the UK or the United States creates 6.2 and 6 new cases worldwide, respectively.

Figure 4: Country-specific IRFs for spatial diffusion of new cases



Note: the figure maps the spatial impulse response functions of the number of domestic confirmed cases at horizon=15 following a unit increase in the number of confirmed cases in a specific country. The numbers in gray correspond to the value of domestic response of the country considered to allow for an assessment of the relative strength of the transmission. The IRFs are based on the panel estimates and the estimation period is 03/13/20 to 04/8/20.

4.4 Cost-Benefit analysis of air travel restrictions

The OECD (2020a) evaluates the potential direct economic impact of widespread containment measures based on an output approach by looking at detailed categories of output and identifies the sectors most directly affected. With these sectors typically accounting for between 30-40% of total output, the overall direct initial hit to the level of GDP lies between 20-25% in most major advanced economies. The OECD's disaggregated estimates suggest that the hit in the sector of air transport, accommodation and food services amounts to about 2.5% of GDP.

Covid Economics 24, 1 June 2020: 111-125

Table 4: Cost-Benefit analysis of air travel restrictions

(In billions of USD)	48 countries	OECD
Cases (up to April, 8)	1,347,683	1,212,471
Deaths (up to April, 8)	87,158	82,161
Cases due to air travel (high)	121,291	109,122
(low)	107,815	96,998
Deaths due to air travel (high)	7,844	7,394
(low)	6,973	6,573
Aggregate GDP	96,000	63,000
Share of aviation in GDP	0.5%*	0.5%
Statistical value of life (high)	0.0096	0.0096
(low)	0.0034	0.0034
Benefit (high)	75.30	70.98
(low)	24.4	23.01
Cost	36.9	24.23

*: The share of aviation in the GDP of the 48 countries is assumed to be the same as for the OECD.

Note: the table reports the Cost-Benefit analysis of the imposition of a 4-week restriction on travel at the height of the epidemic (March, 13, 2020 to April, 8 2020). Case numbers are expressed in units while monetary aggregates are expressed in billions of US dollars. The *low* and *high* estimates correspond to the lower and upper bound of the variance decomposition.

For our purpose, we want to disaggregate air transport from the hotel and restaurant activities. In the OECD, air transportation activities (ISIC Rev4 industry 51) typically represent a small share of the air transport, accommodation and food services aggregate (ISIC Rev4 industry 51;55-56). However, the transportation support industry (ISIC Rev4 industry 52) is an important upstream activity (OECD, 2020b). For this exercise, we assume that air travel accounts for 20% of the air transport, accommodation and food services aggregate. Moreover, the impact on annual GDP growth depends on the duration of the containment measures. In this computation, we consider a 4-week extension of air travel restrictions. The aggregate GDP in 2019 amounted to about USD 63,000 billion for the OECD and 96,000 for our 48 countries. Thus, we evaluate the cost side of our analysis at about USD 30 billion and USD 46 billion, respectively ($0.025 \times 0.2 \times 4/52 \times \text{GDP}$). One needs to bear in mind that this estimation is likely to be an overestimation as the months of March and April are not the busiest period of the year for the aviation industry (Mao et al., 2015).

On the benefit side, our model enables us to quantify the death toll accounted for by spatial dependence. We showed above that air travel accounts for 8-9% of reported cases, or 108,000-120,000 cases. As of April 8 (the end of our estimation window), the global case fatality rate was 6.47%, or 7,000-7,800 deaths associated with air travel. Using a single value of a statistical life of USD 9.6 million (Viscusi, 2015; Viscusi and Masterman, 2017), we evaluate that imposing air travel restrictions would have brought a monetary benefit of around USD 66-75 billion. However, estimates of the statistical value of life typically depends on age profile. Therefore, we also perform our exercise based on an estimate of the statistical value of life for age class 55-62 years of USD 3.4 million (Aldy and Viscusi, 2007; Viscusi and Aldy, 2007). Table 4 summarizes the Cost-Benefit analysis for our panel of countries and the OECD.

We are well aware of the approximate nature of this exercise, and by no means do we take these figures for more than back-of-envelope calculations. Yet, we believe they provide useful points of reference for policymakers and are suggestive that, in the short-run, air travel restrictions might pass the cost-benefit test.

5 Conclusions

In this paper we have shown that commercial flights had a decisive contribution to the worldwide propagation of the SARS-CoV-2 virus during the height of the pandemic. To achieve this, we have presented a Spatial Durbin-Watson model adapted to local projections. This framework allowed us to track the influence of an increase in foreign new cases both temporally and spatially.

Our results can be summarized as follows. First, we estimate that foreign cases represent 8 to 9% of the observed cases. However, it is also relevant to consider the time-varying importance of air traffic for the transmission of COVID-19. Foreign contamination mainly came in two waves spaced by about 9 days (mid-March and third week of March). After the enactment of travel bans, epidemic developments abroad lose their importance for the domestic propagation of the virus. Second, we provide a counterfactual scenario in the form of a 4-week ban on commercial travel and show that such a decision would have saved 7,000 to 7,800 lives whose value might outweigh the cost for the aviation sector in the short-run. As the effect of air travel diminishes over time, however, the marginal benefit is decreasing while the output cost increases as the shutdown is extended. Moreover, finer calculation should take the age distribution of air travel-induced deaths into account.

Although these calculations can only be characterized as a first-order approximation, we nevertheless believe that they could serve as useful point of reference for policymakers, especially so at a time when governments contemplate re-opening national airspaces.

References

- Adam T. Craig, A. E. Heywood, J. H. (2020). Risk of covid-19 importation to the pacific islands through global air travel. *Epidemiology & Infection*.
- Aldy, J. E. and W. K. Viscusi (2008, August). Adjusting the Value of a Statistical Life for Age and Cohort Effects. *The Review of Economics and Statistics* 90(3), 573–581.
- Anselin, L. (1988). A test for spatial autocorrelation in seemingly unrelated regressions. *Economics Letters* 28(4), 335–341.
- Auerbach, A. J. and Y. Gorodnichenko (2012, November). Fiscal Multipliers in Recession and Expansion. In *Fiscal Policy after the Financial Crisis*, NBER Chapters, pp. 63–98. National Bureau of Economic Research, Inc.
- Auerbach, A. J. and Y. Gorodnichenko (2017, September). Fiscal stimulus and fiscal sustainability. Working Paper 23789, National Bureau of Economic Research.
- Brady, R. R. (2011). Measuring the diffusion of housing prices across space and over time. *Journal of Applied Econometrics*.
- Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020, March). The Macroeconomics of Epidemics. NBER Working Papers 26882, National Bureau of Economic Research, Inc.
- Elhorst, J. P. (2003, July). Specification and Estimation of Spatial Panel Data Models. *International Regional Science Review* 26(3), 244–268.
- Gorodnichenko, Y. and B. Lee (2019). Forecast error variance decompositions with local projections. *Journal of Business & Economic Statistics* 0(0), 1–24.
- Greenstone, M. and V. Nigam (2020, March). Does Social Distancing Matter? BFI Working Paper 2020-26, Becker Friedman Institute.
- Guliyev, H. (2020). Determining the spatial effects of covid-19 using the spatial panel data model. *Spatial Statistics*.
- Hsu, C.-I. and H.-H. Shih (2020). Transmission and control of an emerging influenza pandemic in a small-world airline network. *Accident Analysis and Prevention* 42(1), 93–100.
- Keita, S. (2020). Air passenger mobility, travel restrictions, and the transmission of the covid-19 pandemic between countries. *Covid Economics* 9, 77–96.
- Kip Viscusi, W. and J. E. Aldy (2007, May). Labor market estimates of the senior discount for the value of statistical life. *Journal of Environmental Economics and Management* 53(3), 377–392.
- Krisztin, T., P. Piribauer, and M. Wögerer (2020). The spatial econometrics of the coronavirus pandemic. mimeo.
- Lau, H., V. Khosrawipour, P. Kocbach, A. Mikolajczyk, H. Ichii, M. Zacharski, J. Bania, and T. Khosrawipour ((in press)). The association between international and domestic air traffic and the coronavirus (covid-19) outbreak. *Journal of Microbiology, Immunology and Infection*.
- Lesage, J. P. (2008). An Introduction to Spatial Econometrics. *Revue d'économie industrielle* 0(3), 19–44.
- Mao, L., X. Wu, Z. Huang, and A. J. Tatem (2015). Modeling monthly flows of global air travel passengers: An open-access data resource. *Journal of Transport Geography* 48, 52–60.

- OECD (2020a). Evaluating the initial impact of covid-19 containment measures on economic activity. OECD Policy Responses to Coronavirus (COVID-19), OECD.
- OECD (2020b). Structural analysis database.
- Orea, L. and I. C. Álvarez (2020). How effective has been the Spanish lockdown to battle COVID-19? A spatial analysis of the coronavirus propagation across provinces. Working Paper 2/2020, Economic Discussion Paper, Universidad de Oviedo.
- Òscar Jordà (2005, March). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review* 95(1), 161–182.
- Thunström, L., S. C. Newbold, D. Finnoff, M. Ashworth, and J. F. Shogren (2020). The Benefits and Costs of Using Social Distancing to Flatten the Curve for COVID-19. *Journal of Benefit-Cost Analysis* 11(2), 1–17.
- Viscusi, K. W. (2015). The role of publication selection bias in estimates of the value of a statistical life. *American Journal of Health Economics* 1(1), 27–52.
- Viscusi, K. W. and C. J. Masterman (2017). Anchoring biases in international estimates of the value of a statistical life. *Journal of Risk and Uncertainty* 54(2), 103–128.
- Zhuang, Z., S. Zhao, Q. Lin, P. Cao, Y. Lou, L. Yang, and D. He (2020). Preliminary estimation of the novel coronavirus disease (covid-19) cases in iran: a modelling analysis based on overseas cases and air travel data. *International Journal of Infectious Diseases* 94, 29–31.

A List of countries in the analysis

Table 5: List of countries in the analysis

Country	ISO code	Country	ISO code
Australia	AUS	Lithuania	LTU
Austria	AUT	Luxembourg	LUX
Belgium	BEL	Malaysia	MYS
Brazil	BRA	Mexico	MEX
Canada	CAN	Netherlands	NLD
Chile	CHL	New Zealand	NZL
China	CHN	Norway	NOR
Colombia	COL	Philippines	PHL
Czech Republic	CZE	Poland	POL
Denmark	DNK	Portugal	PRT
Estonia	EST	Russia	RUS
Finland	FIN	Singapore	SGP
France	FRA	Slovakia	SVK
Germany	DEU	Slovenia	SVN
Greece	GRC	South Korea	KOR
Hungary	HUN	Spain	ESP
Iceland	ISL	Sweden	SWE
India	IND	Switzerland	CHE
Indonesia	IDN	Taiwan	TWN
Ireland	IRL	Thailand	THA
Israel	ISR	Turkey	TUR
Italy	ITA	United Arab Emirates	ARE
Japan	JPN	United Kingdom	GBR
Latvia	LVA	United States	USA

Time series models based on growth curves with applications to forecasting coronavirus

Andrew Harvey¹ and Paul Kattuman²

Date submitted: 21 May 2020; Date accepted: 26 May 2020

Time series models are developed for predicting future values of a variable which when cumulated is subject to an unknown saturation level. Such models are relevant for many disciplines, but here attention is focussed on the spread of epidemics and the applications are for coronavirus. The time series models are relatively simple but are such that their specification can be assessed by standard statistical test procedures. In the generalized logistic class of models, the logarithm of the growth rate of the cumulative series depends on a time trend. Allowing this trend to be time-varying introduces further flexibility and enables the effects of changes in policy to be tracked and evaluated.

¹ Professor of Econometrics, Cambridge University.

² Reader in Economics, Cambridge Judge Business School, Cambridge University.

Copyright: Andrew Harvey and Paul Kattuman

1 Introduction

The progress of an epidemic typically starts off with the number of cases following an exponential growth path. Over time the growth rate falls and the total number of cases approaches a final level - the ‘leveling of the curve.’ Complex behavioural modeling of the progress of the disease, for example by using the ‘semi-mechanistic Bayesian hierarchical model’ implemented by the team at Imperial College London, depends on many assumptions and unknowns; see Flaxman et al. (2020). Simple and transparent time series models may offer an alternative way of making predictions of the trajectory of the epidemic; see, for example, Chowell et al. (2016, section 2) where such approaches are called ‘phenomenological’. Similar issues arise in economics where there is a contrast between calibrated and Bayesian models based on economic theory and data-based time series models¹. Avery et al (2020) give an excellent review of these issues.

The progress towards an upper bound or saturation level can be taken on board with a sigmoid curve, such as the logistic

$$\mu(t) = \bar{\mu}/(1 + \gamma_0 e^{-\gamma t}), \quad \gamma_0, \gamma, \bar{\mu} > 0, \quad -\infty < t < \infty, \quad (1)$$

where $\bar{\mu}$ is the final level, γ is a rate of progress parameter and γ_0 takes account of the initial conditions. Logistic curves can be shown to arise from a model of a simple epidemic; see, for example, Daley and Gani (2008, ch 2). More generally, sigmoid curves are used in many disciplines for a variety of applications, such as estimating the demand for new products and population growth of mammals subject to space and resource limitations. An early and influential discussion of growth curves was given by Gregg et al. (1964). More recently Panik (2014) describes growth curves and statistical methods for fitting them. Here we concentrate on an extension of the logistic curve called the generalized logistic (GL) or Richards curve; see Panik (2014, p 78-80). The Gompertz curve emerges as an important special case. Figure 1 shows a Gompertz growth curve together with the corresponding change, as given by the first derivative. A second change curve shows the effect of what turns out to be a key parameter, γ , on the peak.

By formulating a statistical model, parameters such as $\gamma_0, \bar{\mu}$ and γ can be estimated from observations, denoted Y_t , made at discrete times $t = 1, \dots, T$, on $\mu(t)$. One option is to work directly with the level, by basing a time series model on a deterministic trend, as in (1); for example Meade and Islam

¹See also Harris et al. (2018) in the context of energy forecasting.




Figure 1: Gompertz function with saturation level of 100 and point of inflexion equal to 30, with $\gamma_0 = 20$ and $\gamma = 0.1$. The bold line shows change in each period times 10. The dotted line has $\gamma = 0.15$.

(1995) report fitting a variety of growth curves to telecommunications data, while a recent article by Ciufolini and Paolozzi (2020) reports fitting a deterministic trend to total coronavirus cases in Italy. The methods discussed in Panik (2014, ch 4) are restricted to deterministic curves, possibly with a first-order autoregressive disturbance term². Our view is that, as with most economic and social time series, a deterministic trend fitted to the level is too inflexible. Instead we prefer to work with the change or the growth rate, with the specification of the model informed by the assumption that the total follows a growth curve. The saturation level can be continually updated as new observations become available. A second advantage of working with the difference or the growth rate is that when logarithms are taken³, estimation of the basic models derived from the GL class can be carried out by least squares regression, as proposed in Harvey (1984). Thus the method is viable even for a small number of observations. Finally the deterministic time trend in the estimating equations can be replaced by a stochastic one. This is particularly effective with a Gompertz function. The flexible trend, which can be estimated by the Kalman filter, as in the STAMP package of Koopman et al. (2020), allows parameters to evolve over time. As a result the model can adapt to significant events and changes in policy. Figure 2 shows predictions made for new cases of coronavirus in Germany. The predictions, which include a day of the week effect, were made using data up to the end of March and show a downward trend even though the series was only just reaching what subsequently turned out to be its peak. Moving beyond the peak signals that R_t , the effective reproduction number at time t , has gone below one; see Flaxman et al. (2020) and Aronson et al. (2020).

When numbers are small, as is the case with deaths at the beginning or end of an epidemic, there is a strong argument for adopting a negative binomial distribution. Models formulated under an assumption of Gaussianity need to be modified accordingly and we show how this may be done using the score-driven approach described in Harvey (2013) and implemented in the TSL computer package of Lit, Koopman and Harvey (2020).

Section 2 describes the GL class of growth curves, while Section 3 dis-

²Tests for a unit root are discussed but the implications are not followed up.

³One implication is that the forecasts of the change will tend asymptotically towards zero and never become negative. In some applications, Y_t can go down as well as up and in these situations the methods proposed here do not apply. Instead flexibility can be brought into the deterministic growth curve by allowing parameters, such as $\bar{\mu}$, to change over time. Young and Ord (1989) effectively do this by discounted least squares.




Figure 2: Predictions for new cases of coronavirus in Germany made on April 1st. NB. There was no observation for May 1st and so May 2nd can be regarded as the sum of 1st and 2nd.

cusses estimation. These methods are then applied in Section 4 to data on coronavirus in UK hospitals and in Germany. Forecasts are set out and evaluated. The effects of policy, primarily the lockdown imposed on 21st March in the UK, are assessed in Section 5. There is concern about a potential second wave of infections as restrictions start to be eased and ways in which this might be monitored are discussed in Section 6. Section 7 concludes.

2 Growth curves

Let $\mu(t) \geq 0$ be a monotonically increasing function defined over the real line. The rate of change or ‘incidence curve’ is $d\mu(t)/dt \geq 0$. The generalized logistic is

$$\mu(t) = \bar{\mu}/(1 + (\gamma_0/\kappa)e^{-\gamma t})^\kappa, \quad \gamma_0, \gamma, \kappa > 0, \quad (2)$$

where γ is a growth rate parameter. The parameter κ must be positive for there to be an upper asymptote; allowing κ to be negative gives the class of general modified exponential (GME) growth curves. The logistic is obtained by setting $\kappa = 1$, while letting $\kappa \rightarrow \infty$ yields the Gompertz curve. When γ_0

is determined by the value of the curve at $t = 0$, it is

$$\gamma_0 = \kappa [(\bar{\mu}/\mu(0))^{1/\kappa} - 1]. \tag{3}$$

Differentiation yields

$$\ln d\mu(t)/dt = \rho \ln \mu(t) + \delta - \gamma t, \tag{4}$$

where $\delta = \ln(\gamma_0 \bar{\mu}^{-1/\kappa} \gamma)$ and $\rho = (\kappa + 1)/\kappa$, so $0 < \kappa < \infty$ implies $1 < \rho < \infty$. Alternatively, because $d\mu(t)/dt = g(t)\mu(t)$,

$$\ln g(t) = (\rho - 1) \ln \mu(t) + \delta - \gamma t, \tag{5}$$

where $g(t)$ is the growth rate of $\mu(t)$. Note that $\rho - 1 = 1/\kappa$.

The generalized logistic differential equation implied by (2) is

$$\frac{d\mu(t)}{dt} = \gamma \kappa \left[1 - \left(\frac{\mu(t)}{\bar{\mu}} \right)^{1/\kappa} \right] \mu(t). \tag{6}$$

The term in square brackets is less than one and tends to zero as $\mu(t) \rightarrow \bar{\mu}$. The growth rate implied by (6) is

$$g(t) = \gamma \kappa \left[1 - \left(\frac{\mu(t)}{\bar{\mu}} \right)^{1/\kappa} \right]. \tag{7}$$

When $\kappa = 1$ (6) is a Riccati equation.

2.1 Where is the peak?

The point of inflexion on the growth curve is the point at which the number of new cases, $d\mu(t)/dt$, peaks. Taking the second derivative⁴ of $\mu(t)$ and setting to zero yields the condition

$$g(t) = -d \ln g(t)/dt = -g_g(t). \tag{8}$$

Where g_g is the growth rate of the growth rate. It follows from (5) that the point of inflexion in the GL is when $g(t) = \gamma/\rho$. The corresponding value of $\mu(t)$ is

$$\mu^* = \bar{\mu} \rho^{-\kappa} = \bar{\mu} (\kappa/(\kappa + 1))^\kappa. \tag{9}$$

⁴ $d^2\mu(t)/dt^2 = g(t)d\mu(t)/dt + \mu(t)dg(t)/dt$. Dividing by $g(t)\mu(t)$ gives the result.

The change declines more slowly than it ascends when $\kappa > 1$; see for example Figure 1.

When γ_0 is determined by $\mu(0)$, as in (3), the peak is at

$$t^* = \ln \gamma_0 / \gamma, \quad \gamma_0 > 1, \quad (10)$$

so it comes forward as γ increases.

2.2 The Gompertz curve

The Gompertz curve is

$$\mu(t) = \bar{\mu} \exp(-\gamma_0 e^{-\gamma t}), \quad \gamma_0, \gamma > 0, \quad -\infty < t < \infty. \quad (11)$$

When the GL is parametrized as in (2), letting $\kappa \rightarrow \infty$ gives (11). Taking the logarithm of $\mu(t)$ gives

$$\ln \mu(t) = \ln \bar{\mu} - \gamma_0 e^{-\gamma t} \quad (12)$$

leading to (5) with $\rho = 1$, that is

$$\ln g(t) = \delta - \gamma t, \quad (13)$$

where $\delta = \ln \gamma_0 \gamma$. The point of inflexion is in (10), which corresponds to $\mu^* = \bar{\mu}/e \approx 0.368\bar{\mu}$, where e is Euler's number. The maximum value of the change is $0.368\gamma\bar{\mu}$, while $g(t) = \gamma$. Letting $\kappa \rightarrow \infty$ in (9), so that $\rho \rightarrow 1$, also gives $\mu^* = \bar{\mu}/e$.

2.3 Statistical distributions and epidemics

Writing a GL growth curve as $\mu(t) = \bar{\mu}F(t)$ allows $F(t)$ to be interpreted as the cumulative distribution function (CDF) of the log of a Dagum distribution; see Kleiber and Kotz (2003, pp 212-13). Hence the probability distribution function (PDF), $f(t)$, is a special case of an Exponential Beta of the Second Kind (EGB2) distribution; see McDonald and Xu (1995).

The Gompertz distribution written in terms of the notation of (11) has $f(t) = \gamma\gamma_0 F(t) \exp(-\gamma t)$. Figure 1 shows a Gompertz PDF (labelled change) with $\gamma_0 = 20$ and $\gamma = 0.1$ as well as $\gamma = 0.15$. The effect of increasing γ is to raise the peak and bring it forward. Since $d\mu(t)/dt = \bar{\mu}f(t)$ the peak is brought down by a lower $\bar{\mu}$.

When estimating a deterministic growth curve, some researchers, for example Ciufolini and Paolozzi (2020), prefer to use a sigmoid defined in terms of the Gaussian error function, which is the CDF of a Gaussian variate. Similarly Murray et. al. (2020) use the Gaussian error function to model the logarithm of total coronavirus cases in US states. Given that the Gaussian error function has no closed form and is not straightforward to evaluate, it is difficult to see the appeal.

It follows from (6) that

$$d\mu(t)/dt = \bar{\mu}F(t)(1 - F(t)^{1/\kappa}), \quad (14)$$

where $\mu(t) = \bar{\mu}F(t)$. In a simple epidemic, $d\mu(t)/dt$ is proportional to a logistic growth curve, $F(t)(1 - F(t))$. Allowing κ to be other than one gives more flexibility and is a useful generalization if the model provides a good fit to the data. Indeed complex mechanistic models of epidemics, with the population assigned to compartments labeled Susceptible, Infected, and Recovered (SIR), often produce incidence curves that are positively skewed. An example is the model of Giordano et al. (2020) which is based on a system of eight differential equations.

3 Statistical modeling

In the observational model, the cumulative total at time $t - 1$ replaces $\mu(t)$ and the (positive) change, $y_t = \Delta Y_t = Y_t - Y_{t-1}$, replaces $d\mu(t)/dt$. Equation (4) leads to

$$\ln y_t = \rho \ln Y_{t-1} + \delta - \gamma t + \varepsilon_t, \quad \rho \geq 1, \quad t = 2, \dots, T, \quad (15)$$

where the disturbance terms ε_t are assumed to be independently and identically distributed with mean zero and constant variance, σ_ε^2 , that is $\varepsilon_t \sim IID(0, \sigma_\varepsilon^2)$. Subtracting $\ln Y_{t-1}$ from both sides gives the form corresponding to (5) namely

$$\ln g_t = (\rho - 1) \ln Y_{t-1} + \delta - \gamma t + \varepsilon_t, \quad (16)$$

where $g_t = y_t/Y_{t-1}$, although it may also be defined as $\Delta \ln Y_t$. The parameters ρ, δ and γ can therefore be estimated by regression. If ρ takes a specified value, $\ln y_t - \rho \ln Y_{t-1}$ is simply regressed on a constant and time trend. The estimators of δ and γ are then efficient when the disturbance is Gaussian, that is $\varepsilon_t \sim NID(0, \sigma_\varepsilon^2)$.

The observational model for the Gompertz curve is

$$\ln y_t = \ln Y_{t-1} + \delta - \gamma t + \varepsilon_t, \quad t = 2, \dots, T, \tag{17}$$

or the simple time trend regression

$$\ln g_t = \delta - \gamma t + \varepsilon_t, \quad t = 2, \dots, T. \tag{18}$$

3.1 Forecasts

Forecasts of future observations and an estimate of the saturation level can be obtained from the predictive recursions

$$\hat{y}_{T+\ell|T} = \hat{\mu}_{T+\ell-1|T}^\rho \exp(\delta_T) \exp(-\gamma\ell) \tag{19}$$

$$\hat{\mu}_{T+\ell|T} = \hat{\mu}_{T+\ell-1|T} + \hat{y}_{T+\ell|T}, \quad \ell = 1, 2, \dots \tag{20}$$

where δ_T is the level at time T , that is $\delta_T = \delta - \gamma T$, and $\hat{\mu}_{T|T} = Y_T$. As $\ell \rightarrow \infty$, $\hat{\mu}_{T+\ell|T} \rightarrow \bar{\mu}$. Alternatively

$$\hat{g}_{T+\ell|T} = \hat{\mu}_{T+\ell-1|T}^{\rho-1} \exp(\delta - \gamma(T + \ell)), \quad \ell = 1, 2, \dots \tag{21}$$

$$\hat{\mu}_{T+\ell|T} = \hat{\mu}_{T+\ell-1|T}(1 + \hat{g}_{T+\ell|T}) \tag{22}$$

so that $\hat{y}_{T+\ell|T} = \hat{g}_{T+\ell|T} \hat{\mu}_{T+\ell-1|T}$ and $\hat{Y}_{T+\ell|T} = \hat{\mu}_{T+\ell|T}$.

For the Gompertz model, where $\rho = 1$, the forecasts for the growth rate are simply

$$\hat{g}_{T+\ell|T} = \exp(\delta_T) \exp(-\gamma\ell), \tag{23}$$

so (21) and (22) yield

$$\hat{\mu}_{T+\ell|T} = \hat{\mu}_{T|T} \prod_{j=1}^{\ell} (1 + \exp(\delta_T) \exp(-\gamma j)). \tag{24}$$

A future point of inflexion is given at $\ell^* = (\delta_T - \ln \gamma) / \gamma$. When⁵ $\ell \rightarrow \infty$, $\hat{\mu}_{T+\ell|T} \simeq Y_T \exp(\exp(\delta_T) / (\exp(\gamma) - 1))$.

Predictive distributions of future observations may be obtained by simulation.

⁵ $\ln \prod_{j=1}^{\ell} (1 + g_{T+\ell}) = \ln \sum_{j=1}^{\ell} (1 + g_{T+\ell}) \simeq \sum_{j=1}^{\ell} g_{T+\ell}$. Higher order terms can be neglected when $g_{T+\ell}$ is small. Then $\sum_{j=1}^{\ell} g_{T+\ell} = \delta_T \sum_{j=1}^{\ell} \exp(-\gamma j) \rightarrow \delta_T / (\exp(-\gamma) - 1)$ as $\ell \rightarrow \infty$

Remark 1 An assumption of normality for the disturbances in (15) implies that, conditional on information at time T , y_{T+1} is log-normal. Thus there may be a case for multiplying $\hat{y}_{T+1|T}$ by $\exp(0.5\sigma^2)$, where σ^2 is the prediction error variance. (When there is no stochastic trend $\sigma^2 = \sigma_\varepsilon^2$.)

3.2 Gompertz model with a time-varying trend (dynamic Gompertz)

A stochastic trend may be introduced into (18), that is

$$\ln g_t = \delta_t + \varepsilon_t, \quad \varepsilon_t \sim NID(0, \sigma_\varepsilon^2), \quad t = 2, \dots, T,$$

where

$$\begin{aligned} \delta_t &= \delta_{t-1} - \gamma_{t-1} + \eta_t, & \eta_t &\sim NID(0, \sigma_\eta^2), \\ \gamma_t &= \gamma_{t-1} + \zeta_t, & \zeta_t &\sim NID(0, \sigma_\zeta^2), \end{aligned} \quad (25)$$

and the normally distributed irregular, level and slope disturbances, ε_t , η_t and ζ_t , respectively, are mutually independent. When $\sigma_\eta^2 = \sigma_\zeta^2 = 0$, the trend is deterministic, that is $\delta_t = \delta - \gamma t$ with $\delta = \delta_0$. When only σ_ζ^2 is zero, the slope is fixed and the trend reduces to a random walk with drift. On the other hand, allowing σ_ζ^2 to be positive, but setting $\sigma_\eta^2 = 0$ gives an integrated random walk (IRW) trend, which when estimated tends to be relatively smooth; see, for example, Harvey (2008). Figure 3 shows smoothed estimates of the logarithm of the Italian growth rate using data for March 2020.

The STAMP package of Koopman et al. (2020) can estimate a stochastic trend using techniques based on state space models and the Kalman filter. The Kalman filter outputs the state at time T so the forecasts are given by (22) and (24). The forecasts respond to recent movements in the data. Thus an increase in γ_t , brought about by smaller than expected values of y_t , means a faster decline in the growth rate.

A stochastic trend can be introduced into the more general GL model. However, unless ρ is fixed, it may be hard to estimate in small samples.

The Kalman filter can be bypassed by adopting the reduced form, which comes from the innovations form of the Kalman filter; see Harvey (2008). For the GL curve

$$\ln y_t = \rho \ln Y_{t-1} + \delta_{t,t-1} + \varepsilon_t, \quad t = 3, \dots, T, \quad (26)$$




Figure 3: Smoothed estimates of logarithms of the growth rate of total cases in Italy in March.

or

$$\ln g_t = (\rho - 1) \ln Y_{t-1} + \delta_{tt-1} + \varepsilon_t, \quad t = 3, \dots, T, \quad (27)$$

where δ_{t+1t} is the filtered form of the IRW

$$\begin{aligned} \delta_{t+1t} &= \delta_{tt-1} - \gamma_{tt-1} \\ \gamma_{t+1t} &= \gamma_{tt-1} + \alpha \varepsilon_t. \end{aligned}$$

The parameter α , which is non-negative, plays a similar role to the signal-noise ratio, q_ζ . When it is zero the model reverts to (16).

3.3 Models for the growth rate ⁶

A stochastic model for the logistic curve⁷ can be based on (7), as in Levenbach and Reuter (1976), by adding a serially independent Gaussian disturbance

⁶This sub-section can be omitted without loss of continuity.

⁷The local logistic model described by Meade and Islam (1995, p 202) is somewhat similar.

term so that

$$g_t = \gamma - (\gamma/\bar{\mu})\widehat{\mu}_{t|t-1} + \varepsilon_t, \quad t = 1, \dots, T, \tag{28}$$

where $\widehat{\mu}_{t|t-1}$ is an estimate of μ_t , such as Y_{t-1} , based on information at time $t-1$ and ε_t is a serially independent Gaussian disturbance with mean zero and constant variance, σ_ε^2 , that is $\varepsilon_t \sim NID(0, \sigma_\varepsilon^2)$. Regressing g_t on Y_{t-1} gives estimates of the key parameters γ and $\bar{\mu}$. Generalization of this approach is based on (7) but leads to a nonlinear equation

$$g_t = \gamma\kappa - \left(\frac{Y_{t-1}}{\bar{\mu}}\right)^{1/\kappa} + \varepsilon_t, \quad t = 1, \dots, T, \tag{29}$$

which requires a search over the range of κ if it is unknown. In the Gompertz case, estimation is as straightforward as for the logistic model because

$$g_t = \gamma \ln \bar{\mu} - \gamma \ln Y_{t-1} + \varepsilon_t, \quad t = 1, \dots, T. \tag{30}$$

Models based on the logarithm of the growth rate are preferred because they seem to have better statistical properties. For example the disturbance term is less likely to be heteroscedastic.

3.4 Small numbers: the negative binomial distribution

When y_t is small, it may be better to specify its distribution, conditional on past values, as discrete. The usual choice is the negative binomial which, when parameterized in terms of a time-varying mean, ξ_{bt-1} , and a fixed positive shape parameter, v , has probability mass function (PMF)

$$p(y_t) = \frac{\Gamma(v + y_t)}{y_t! \Gamma(v)} \xi_{bt-1}^{y_t} (v + \xi_{bt-1})^{-y_t} (1 + \xi_{bt-1}/v)^{-v}, \quad y_t = 0, 1, 2, \dots,$$

with $Var_{t-1}(y_t) = \xi_{bt-1} + \xi_{bt-1}^2/v$. An exponential link function ensures that ξ_{bt-1} remains positive and at the same time yields an equation similar to (15):

$$\ln \xi_{bt-1} = \rho \ln Y_{t-1} + \delta - \gamma t, \quad \rho \geq 1, \quad t = 3, \dots, T. \tag{31}$$

A stochastic trend may be introduced into the model, as in sub-section 3.2, by developing a filter for the time-varying trend similar in structure to that of (26). Because the observations are not Gaussian, the dynamic

conditional score (DCS) framework described⁸ in Harvey (2013, pp 77-79) is used to give

$$\ln \xi_{it-1} = \rho \ln Y_{t-1} + \delta_{it-1}, \quad t = 3, \dots, T, \quad (32)$$

where δ_{t+1t} is the filtered form of the IRW

$$\begin{aligned} \delta_{t+1t} &= \delta_{it-1} - \gamma_{it-1} \\ \gamma_{t+1t} &= \gamma_{it-1} + \alpha u_t, \quad \alpha \geq 0, \end{aligned}$$

but with $u_t = y_t/\xi_{it-1} - 1$, which is the conditional score for $\ln \xi_{it-1}$, that is $v(y_t - \xi_{it-1})/(v + \xi_{it-1})$, divided by the information quantity. The dynamic Gompertz model has $\rho = 1$.

Predictions of future observations and the saturation level can be obtained from (19) and (20).

4 Forecasting Coronavirus in the UK and Germany

We began working on this project at the beginning of April. At that time coronavirus was not as far advanced in the UK as in Italy, and our initial exploration was focussed on Italy. Figure 4 shows new cases, as measured by hospital admissions on the European Centre for Disease Prevention and Control (ECDC) website⁹, in the UK and Italy. Italy is seen to be clearly ahead and it is apparent that the decline was slower than the rise. This asymmetry is not consistent with a simple logistic model and later analysis confirmed this to be the case for most other countries, including the UK.

Economic and social time series are typically subject to periodic variation, due to seasonal, day of the week and other effects. Preliminary analysis of data for hospital admissions and deaths in Italy indicated a day of the week pattern and this was confirmed for the UK; one reason is that laboratory confirmation for the virus tends to slow down at weekends. Day of the week effects were initially included in the equations by the introduction of a seven day cycle, which is more parsimonious than adding six dummy variables.

⁸A slightly more elaborate form corresponding to a cubic spline smoother may also be adopted.

⁹<https://www.ecdc.europa.eu/en/publications-data/download-todays-data-geographic-distribution-covid-19-cases-worldwide>




Figure 4: New cases of coronavirus in UK (bold) and Italy up to April 10th.

Although we estimated many models in early April, the first results reported here for the UK are those obtained on April 13th, that is with data up to the 12th, starting on March 5th. The models were updated on April 20th and 27th. Unfortunately, the data available to us after April 29th was not consistent with the earlier data as it is not just confined to tests carried out on hospital admissions.

4.1 Models fitted to new cases in the UK

Table 1 shows April 13th estimates for the GL, (16), the Gompertz, where ρ is set to one, and the dynamic Gompertz where the trend is an IRW. The maximized log-likelihood is $\ln L$ and PEV denotes the variance of the one-step ahead prediction errors. The results for the logistic model are not reported because it gave an inferior fit and was rejected by a ‘t-test’ on $\hat{\rho}$. The diagnostics presented are the Durbin-Watson (DW) statistic for residual serial correlation, a portmanteau Q -statistic for serial correlation based on P autocorrelations, the Bowman-Shenton normality test statistic and a heteroscedasticity statistic consisting the squares of the last third of the residuals over the first third.

Statistic			Dynamic	Daily	Dynamic
	GL	Gompertz	Gompertz	GL	Gompertz (April 20)
γ	0.081	0.033	0.054	0.091	0.061
δ_T	-5.62	-2.48	-2.63	-6.32	-3.15
ρ	0.264 (0.153)	1	1	0.323 (0.129)	1
q_ζ	0	0	0.001	0	0.001
$\ln L$	38.350	37.765	39.368	42.590	55.595
<i>PEV</i>	0.071	0.078	0.086 ¹⁰	0.049	0.058
<i>DW</i>	1.62	1.43	1.67	2.10	2.01
<i>Q</i> (6)	7.32	5.30	8.25	<i>Q</i> (9)=9.38	<i>Q</i> (10)=9.32
<i>Normality</i>	0.47	.46	1.09	1.29	1.64
<i>Hetero F</i>	0.37	.72	.45	.52	0.45

Table 1 Estimates for UK new cases as of April 13th.

Figure 5 shows the fit and residuals from the GL. Figure 6 shows the histogram and correlogram for the GL residuals but using data available on April 20th.

There was a substantive reduction in the prediction error variance for all models when daily effects were included. The peak day is Friday, the same day as was found for Italy. The results for the GL are shown in the last column. The likelihood ratio statistic is 8.18 which is significant at the 5% level of a χ^2_2 distribution.

4.2 Forecasts and forecast evaluation

The forecasts of new cases made on April 13th are shown in Figure 7, together with the actual values up to April 29th. The upper line is the dynamic Gompertz, while the lower lines are the GL with and without the daily effect. The actual outcome turns out to lie between the dynamic Gompertz and the GL. As with the German predictions shown in Figure 2, the important point is that for both models the forecasts are moving downwards even though the observations have barely peaked.

Predictions from the GL made on April 20th showed little change from those made on April 13th. The saturation level is estimated to be 186,000. By contrast the dynamic Gompertz predictions changed significantly. They are

¹⁰The PEV can be higher in a dynamic model even though the likelihood is bigger.




Figure 5: Fit of GL to logarithm of UK growth rates of total cases on April 13th




Figure 6: Residuals from GL fitted on April 20th




Figure 7: Forecasts for UK new cases made on April 13th

still higher than the GL predictions but the saturation level is now 237,790, whereas on April 13th it was 308,960. The dynamic Gompertz predictions made on April 27th showed little movement. The final predicted level was now 253,800. Figure 8 shows the dynamic Gompertz predictions without the daily component, but this was included when the models were estimated. The flexibility of the dynamic model is allowing it to adapt to a situation in which the observations are falling less rapidly than indicated by the static GL.

4.3 Germany

The data for new cases in Germany is from the Robert Koch Institut (RKI) and are the confirmed cases of COVID in all national hospitals and testing centres¹¹.

Fitting with data for all of March gave the results in Table 3 for a Dynamic Gompertz model with daily effects. Estimation of a GL gave $\tilde{\rho} = 1.21(0.33)$, so the t-statistic is only 0.64. The likelihood was smaller at 16.43. There is no evidence of residual serial correlation in either model and although the heteroscedasticity statistic indicates a diminishing variance, its value seems to

¹¹https://www.rki.de/DE/Home/homepage_node.htm




Figure 8: Forecasts for UK new cases made by dynamic Gompertz on April 20th and 27th.

be heavily influenced by just one observation (March 9th) near the beginning of the series. The fit is shown in Figure 9.

	Dynamic	Dynamic
Statistic	Gompertz (April 1)	Gompertz (May 7)
γ	0.077	0.062
δ_T	-2.41	-5.15
q_s	0.0015	0.0015
$\ln L$	18.43	65.43
PEV	0.182	0.101
DW	2.02	2.01
$Q(9) + Q(11)$	10.02	16.83
<i>Normality</i>	0.86	6.11
<i>Hetero F</i>	0.42	0.22

Table 3. Germany: Estimation results for Dynamic Gompertz models with daily effects

The forecasts shown earlier in Figure 2 are quite remarkable in their accuracy over the next 36 days, that is up to May 6th; see also Figure 10. The observations had not yet started to go down on April 1st, yet the sigmoid




Figure 9: New cases in Germany with smoothed estimates of level and slope.

nature of the underlying growth curve means that the forecasts pick up the subsequent downward movement.

Fitting the dynamic Gompertz using data up to and including May 6th shows the model to be stable; see Table 3 and Figure 11. In particular the signal-noise ratio is almost unchanged. The observations in April were less variable than in March so the residuals are much smaller. The slope appears to have fallen in early April but it then returns to a value similar to what it was at the end of March. In summary the dynamic Gompertz model works extremely well¹².

4.4 Deaths

Deaths in Germany behave in a similar way to new cases: Figure 12 shows the log growth rates. There were some zero values in the first three weeks of March so the dynamic Gompertz was estimated using data from March 22nd. Because the numbers are smaller it is perhaps not surprising that the signal-noise ratio is estimated to be zero. In other words we end up with the static Gompertz model.

¹²Even a deterministic trend gives a reasonably good fit for data from April 1st, although it is not as good as a model with a stochastic trend.




Figure 10: Forecasts for logarithm of growth rate of total cases in Germany made on April 1st.




Figure 11: Smoothed estimates of level and slope of logarithm of growth rates for model fitted on May 7th.

Covid Economics 24, 1 June 2020: 126-157




Figure 12: Logarithm of growth rate for total infected and deaths in Germany

The score-driven negative binomial model, introduced in sub-section 3.4, can be estimated with the TSL package of Lit et al. (2020). Observations from 11th March, which include some zero values, gave $\tilde{\alpha} = 0$ - corresponding to a deterministic trend - and $\tilde{\gamma} = 0.071$. The ACF of the scores had a relatively high value at lag seven indicating a daily effect. Modeling a fixed daily effect with dummy variables produced the fit in Figure 13, with the log-likelihood significantly increased. The parameter estimates were $\tilde{\gamma} = 0.070$, $\tilde{\delta}_T = -4.14$ and $\tilde{\nu} = 13.25$. It is reassuring that the values of $\tilde{\gamma}$ and $\tilde{\delta}_T$ are consistent with those reported for new cases. The final total is predicted to be 8714.

Up to April 27th, the UK data on deaths from coronavirus was restricted to deaths in hospitals. After April 28th deaths in the community were included and the earlier figured revised. Although the relationship between deaths and new cases before April 28th was similar to that in Germany, there is a disconnect after that and the data on new cases is of very little help in predicting deaths.

Figure 14 shows the fit of a Negbin model, with daily effects, to the series on UK deaths. The estimated α of 0.34 (0.05) indicates a changing trend, which is consistent with the models fitted to new cases. The slope at the end of the series, May 14th, was 0.042 while $\tilde{\nu} = 29.91$ (6.13). The




Figure 13: German deaths and Negbin filter with daily effects

day of the weeks multiplication factors are as follows starting with Monday: $-0.343, 0.271, 0.210, 0.106, 0.175, 0.038$ and -0.457 . Figure 15 shows the forecasts of the underlying trend up to one month ahead. The new observations after the forecasts were made are marked by dots.

4.5 Leading indicators

Figure 12 shows that the series for the logarithms of the growth rates of total infected and deaths in Germany move downwards together, maintaining a roughly constant distance apart. The fact that the series for deaths is higher than the one for infected is consistent with there being a lag. The Gompertz model is particularly adept at dealing with this situation. Let β^{\dagger} be the proportion of new cases who subsequently die and let $\beta_j, j = 1, \dots, J$, be the fraction of those who die after j days. The number of deaths at time t is then $y_t^D = \beta^{\dagger} \sum \beta_j y_{t-j}$, with $\sum \beta_j = 1$. When estimates of the lag structure are available, the weighted averages $y_t^w = \sum \beta_j y_{t-j}$ and $Y_t^w = \sum \beta_j Y_{t-j}$ may be used to constructed a growth rate $g_t^w = y_t^w / Y_t^w$; the β^{\dagger} parameter cancels out. The dynamic Gompertz model then yields

$$\ln g_t^D = \delta + \ln g_t^w + \varepsilon_t, \quad \varepsilon_t \sim NID(0, \sigma_\varepsilon^2), \tag{33}$$




Figure 14: Negbin model fitted to UK deaths




Figure 15: Forecasts of UK deaths made on May 15th for one month ahead.

with $\delta > 0$. A single lag of τ time periods in a static Gompertz model would have $\delta = \gamma\tau$. The number of new cases serves as a useful leading indicator for deaths¹³ if the lag structure is such that g_t^w can be constructed some time in advance.

5 The effect of policy interventions

The effect of significant interventions, which may be a policy change, such as the introduction of a lockdown, or an external event, such as the arrival of a cruise ship in a small port, may be modeled by intervention (dummy) variables. The difficulty is that the pattern of the response is rarely known and so it becomes difficult to obtain a meaningful estimate of the final effect. Nevertheless some notion of the response can be obtained by making forecasts at the time the policy is thought to have become effective¹⁴ and comparing these forecasts with the actual outcome. In order to investigate this possibility, the Dynamic Gompertz model was estimated on UK data up to and including March 31 (10 days after lockdown) using a fixed q_ζ of 0.001 (as estimated later with a larger sample and reported in Table 1). No day of the week effect was included, because the series is rather short. Figure 16 shows there is considerable overprediction. However, if q is increased to $q_\zeta = 0.01$ so the most recent observations receive more weight, the predictions are excellent. The saturation level is now approximately 290,000 as against 1.8 million. New cases are at their maximum, with a value of 5,400, on April 18th, whereas with $q = 0.01$, they do not peak until May 20th, when they reach 21,500. Although the variation in predictions is huge, the same is true of many of the large models where the output can be very sensitive to the assumptions made.

Information about the pattern of the response can also be obtained by fitting a dynamic Gompertz model and graphing the estimates of the slope. Figure 17 shows filtered and smoothed estimates of slope in such a model based on data up to April 29th. The filtered estimates are most informative as they show the evolving changes in the slope after the lockdown of 21st March. As can be seen, the big falls occur at the beginning of April, with little or no change after mid-April. Figure 9 shows a very similar movement in Germany, but taking place a few weeks earlier.

¹³We are grateful to Mark Salmon for suggesting that we investigate this issue.

¹⁴Flaxman et al. (2020) provide information on the incubation period.




Figure 16: Forecasts of UK new cases on April 1st with $q = 0.01$ and $q = 0.001$ (dashed).




Figure 17: Filtered and smoothed estimates of slope in dynamic Gompertz model with daily effects estimated on April 30th

Covid Economics 24, 1 June 2020: 126-157




Figure 18: Cumulative distribution for logistic (bold) and Gompertz with median set to 7. For logistic $\gamma = 0.5$ and for Gompertz $\gamma = 0.2$.

Growth curves can be used to parameterize a gradual response to an intervention. A permanent change is captured with the CDF and a temporary one with the PDF. A logistic CDF gives a response curve $W(t) = 1/(1 + \gamma_0^I \exp(-\gamma^I(t - t^I)))$, where t^I is the median. The I superscript distinguishes the parameters γ_0^I and γ^I from the ones used to model the time series itself. With t^L and t^U denoting the beginning and the end of the time span during which gradual response to the intervention occurs, the intervention dummies are defined by $w_t = 0$ for $t < t^L$, $w_t = W(t)$ for $t = t^L, t^L + 1, \dots, t^I, \dots, t^U$ and $w_t = 1$ for $t = t^U + 1, \dots, T$, or even just by $w_t = W(t)$ for $t = 1, \dots, T$. With the Gompertz CDF the response is $W(t) = \exp(-\gamma_0^I \exp(-\gamma^I(t - t^I)))$; in this case the point of inflexion, corresponding to the maximum change, comes before the median of the time span between t^L and t^U . Figure 18 shows the CDFs for logistic and Gompertz distributions with the median set to seven; for logistic $\gamma^I = 0.5$ and for Gompertz $\gamma^I = 0.2$.

If the effect of a policy is to change γ in the model, a slope intervention is needed in (16). Thus

$$\ln g_t = (\rho - 1) \ln Y_{t-1} + \delta - \gamma t - \beta t w_t + \varepsilon_t, \quad t = 1, \dots, T, \quad (34)$$




Figure 19: Estimates of logarithm of growth rate of total cases in UK with a logistic intervention and a daily effect

but unless the sample is moderately large, ρ will need to be fixed. When the full effect is realised, the slope on the time trend will have moved from γ to $\gamma + \beta$. A positive β will lower the growth rate, g_t , the peak of the incidence curve and the saturation level.

Fitting the static model in (34) to new cases in the UK with $\rho = 1$ and a logistic intervention, starting on March 26th and ending on April 12th, gave an estimate of β equal to 0.020 (0.004) and an estimate of γ also equal to 0.020. The picture in Figure 19 is not unreasonable but the estimate of the overall effect is 0.041 which may be a slight underestimate because minus one times the final slope in Figure 17 is close to 0.05. When the slope was allowed to be stochastic, $\beta_T + \gamma$ was estimated at 0.054. However, a stochastic slope risks some confounding with the intervention variable and in fact the estimate of β was reduced to 0.014 (0.006). Although neither model is completely satisfactory, both give a significant coefficient for the intervention variable, albeit after a degree of data mining.




Figure 20: New cases in Germany up to on May 7th together with filtered growth rate and its rate of change.

6 A second wave?

With the relaxation of the lockdown the unwelcome prospect of a second wave of infections arises. Dynamic GL models can monitor this possibility by tracking changes in the estimates of γ and the growth rate, g . It was shown in (8) that

$$g_y(t) = g(t) + g_g(t) \tag{35}$$

and for GL curves $g_y(t) = (\rho - 1)g(t) - \gamma$. In discrete time the negative of the growth rate of the growth rate, $g_g(t)$, is tracked by the filtered estimates of the slope, that is $\gamma_{t|t-1}$, while the growth rate itself is tracked by the exponent of the filtered level, that is $g_{t|t-1} = \exp(\mu_{t|t-1})$. Figure 20 shows $\gamma_{t|t-1}$ and $g_{t|t-1}$ for Germany from a dynamic Gompertz model, together with the daily number of new cases. The maximum is identified as April 3rd and after this date $g_{t|t-1} < \gamma_{t|t-1}$. The possible onset of a second wave is raised if at some point $\gamma_{t|t-1}$ starts to fall below $g_{t|t-1}$, or below $(\rho - 1)g_{t|t-1}$ in the general case. This would signal that the reproduction number, R_t , has moved up above one.

The filtered estimates, $g_{t|t-1}$ and $\gamma_{t|t-1}$ are obtained by discounting of past observations, with the rate of discounting depending on the signal-noise

ratio, q_ζ . When a new policy is implemented, q_ζ may need to increase so that past observations are discounted at a faster rate, as illustrated in Figure 16. In these circumstances, the only viable tracking option may be to try a range of q_ζ values, bearing in mind the risk of triggering a false alarm. However, if the effects of a policy are spread over a period of time, as with a gradual relaxation of a lockdown, a value of q_ζ estimated with the complete sample may be perfectly satisfactory.

If $\gamma_{t|t-1}$ falls, some idea of sampling variability is needed to assess the implications. This may be obtained from the standard deviation of $\gamma_{t|t-1}$, as given by the Kalman filter. The variation in $g_{t|t-1}$ is harder to determine, but since it is typically much smoother than $\gamma_{t|t-1}$ it is likely to be small in comparison to the variation in $\gamma_{t|t-1}$.

Predictions made at a point after the change in policy may also be helpful in the same way as they were for assessing the impact of the lockdown.

7 Conclusions

A new class of time series models is developed for predicting future values of a variable which when cumulated is subject to an unknown saturation level. Such situations arise in a wide range of disciplines. The models provide a simple and viable alternative, or complement, to the forecasts produced by large scale mechanistic models.

Generalized logistic growth curves provide the basis for our models. Estimating equations for the logarithm of the growth rate of the cumulative variable provide a good fit to the data, as assessed by standard statistical tests. Such models feature a time trend which can be made time-varying using the Kalman filter. The Gompertz model is a special case which works particularly well with a stochastic trend and when this modification is made the fit is often better than with an unrestricted generalized logistic model with a deterministic trend. The dynamic Gompertz can adapt to changing conditions and tracking the slope, especially the filtered slope, can be informative about the effect of interventions.

Additional components can be included in the models. These include seasonal or day of the week effects. The latter turned out to be relevant for new cases and deaths in our application to coronavirus in the UK and Germany. The dynamic Gompertz model worked well for both countries but was particularly impressive for German new cases. Deaths were successfully modelled

by a Gompertz model with a negative binomial conditional distribution and a dynamic equation driven by the conditional score.

Estimating a model up to the point at which the effect of an intervention is likely to make itself felt and then making (unconditional) forecasts provides information about the effects of the intervention. The effectiveness of this approach is illustrated by fitting a dynamic Gompertz model to UK data and then investigating the effect of the coronavirus lockdown. Further insight into the effect of the lockdown is given by tracking the filtered estimate of the growth rate of the growth rate. The impact of the lockdown can be estimated *ex post* by including a growth curve as an explanatory variable, thereby allowing the intervention to have a gradual response.

Finally we suggest that the possibility of a second wave can be monitored by tracking the filtered estimates of new cases or deaths given by our model. Of course, if these methods are to be useful in practice they require reliable up to date observations on new cases, preferably at a disaggregated level. Implementing viable tracking procedures and relating them to current methods for tracking the reproduction number, R , is the next phase in our research program.

Acknowledgements

We would like to thank Robert Taylor, Dario Palumbo, Mark Salmon and Qingyuan Zhao for helpful comments and suggestions. Jonas Knecht supplied valuable research assistance.

8 References

Aronson, J.K., Brasset, J. & Mahtani, K.R. (2020) When will it be over? An introduction to viral reproduction numbers. <https://www.cebm.net/oxford-covid-19-evidence-service/>

Avery, C., Bossert, W., Clark, A., Ellison, G., Ellison, S.F. (2020) Policy Implications of models of the spread of coronavirus: perspectives and opportunities for economists. National Bureau of Economic Research Working Paper 27007 <http://www.nber.org/papers/w27007>

Chowell, G., Sattenspiel, L., Bansal, S. & Viboud, C. (2016) Mathematical models to characterize early epidemic growth: medRxiv, 18, 66-97. doi:10.1016/j.phrev.2016.07.005

Ciufolini, I. & Paolozzi, A. (2020) Mathematical prediction of the time evolution of the COVID-19 pandemic in Italy by a Gauss error function and

Monte Carlo simulations. *Eur. Phys. J. Plus* 135, 355.

<https://doi.org/10.1140/epjp/s13360-020-00383-y>

Daley, D. & Gani, J. (1999). *Epidemic Modelling: An Introduction*. 2nd Edition. Cambridge: Cambridge University Press.

Flaxman, S., Swapnil M., Gandy, A. et al. (2020). Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries (Report 13). Imperial College London. <https://www.imperial.ac.uk/mrc-global-infectious-disease-analysis/covid-19/report-13-europe-npi-impact/>

Giordano, G., Blanchini, F., Bruno, R. et al. (2020). Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy. *Nat Med.* <https://doi.org/10.1038/s41591-020-0883-7>

Gregg, J.V., Hossell, C.H. & Richardson, J.T. (1964). *Mathematical Trend curves: an aid to forecasting*. ICI monograph no. 1. Oliver and Boyd, Edinburgh.

Harris, T.M., Devkota, J.P., Khanna, V., Eranki, P.L. & Landis, A.E., (2018). Logistic growth curve modeling of US energy production and consumption, *Renewable and Sustainable Energy Reviews*, Elsevier, vol. 96, 46-57.

Harvey, A.C. (1984) Time Series Forecasting Based on the Logistic Curve. *Journal of the Operational Research Society*, 35, 641-46.

Harvey, A.C. (2006). Forecasting with Unobserved Components Time Series Models. *Handbook of Economic Forecasting*, edited by G. Elliot, C. W.J. Granger and A. Timmermann, pp 327-412. Amsterdam: North Holland.

Harvey, A.C. (2013). *Dynamic Models for Volatility and Heavy Tails: with applications to financial and economic time series*. Econometric Society Monograph, Cambridge University Press.

Kleiber, C. & Kotz, S. (2003). *Statistical Size Distributions in Economics and Actuarial Sciences*. New York: Wiley.

Koopman, S.J., Lit, R. & Harvey, A.C. (2020) STAMP 8.4 Structural Time Series Analyser, Modeller and Predictor. London: Timberlake Consultants Ltd.

Levenbach, H. & Reuter, B.E. (1976). Forecasting trending time series with relative growth rate models. *Technometrics*, 18, 26-72.

Lit, R., Koopman, S.J. & Harvey, A.C. (2020). Time Series Lab - Score Edition: <https://timeserieslab.com>

Meade, N. & Islam, T. (1995). Forecasting with growth curves: an empirical comparison. *International Journal of Forecasting*, 11, 199-215.

McDonald J. B. and Xu, Y. J. (1995). A generalization of the beta distribution with applications. *Journal of Econometrics* 66, 133-152.

Murray, C. J. L. (2020). Forecasting COVID-19 impact on hospital bed-days, ICU-days, ventilator-days and deaths by US state in the next 4 months. Medrxiv. March 30, 2020. <https://www.medrxiv.org/content/10.1101/2020.03.27.20043752v1>

Panik, M.J. (2014) *Growth Curve Modeling: Theory and Applications*. Hoboken, New Jersey: John Wiley & Sons, 2014.

Young, P. & Ord, J.K. (1989). Model selection and estimation for technological growth curves, *International Journal of Forecasting*, 5, 501-513.

The impact of Covid-19 on the US child care market: Evidence from stay-at-home orders¹

Umair Ali,² Chris M. Herbst,³ and Christos A. Makridis⁴

Date submitted: 21 May 2020; Date accepted: 27 May 2020

Stay-at-home orders (SAHOs) have been implemented in most U.S. states to mitigate the spread of COVID-19. This paper quantifies the short-run impact of these containment policies on search behavior and labor demand for child care. The child care market may be particularly vulnerable to a SAHO-type policy shock, given that many providers are liquidity-constrained. Using plausibly exogenous variation from the staggered adoption of SAHOs across states, we find that online job postings for early care and education teachers declined by 13% after enactment. This effect is driven exclusively by private-sector services. Indeed, hiring by public programs like Head Start and pre-kindergarten has not been influenced by SAHOs. In addition, we find little evidence that child care search behavior among households has been altered. Because forced supply-side changes appear to be at play, our results suggest that households may not be well-equipped to insure against the rapid transition to the production of child care. We discuss the implications of these results for child development and parental employment decisions.

1 The paper reflects the views of the authors and not of their affiliated institutions. The authors thank Rob Sentz and Kevin Kirchner at Emsi for their generous support of the data, as well as Angela Rachidi, Erdal Tekin, and Joe Vavra for comments.

2 School of Public Affairs, Arizona State University.

3 School of Public Affairs and IZA, Arizona State University.

4 W. P. Carey School of Business, Arizona State University & Sloan School of Management, MIT.

Copyright: Umair Ali, Chris M. Herbst, and Christos A. Makridis

I. Introduction

The COVID-19 pandemic, and the resulting “stay-at-home” orders in most states, are likely to have large—and potentially permanent—negative consequences for the U.S. child care market. Indeed, given that child care businesses generate most of their revenue from parent fees, it is not surprising that 17% of providers in a recent survey reported they would not survive a closure of any length, and an additional 30% said they would not survive a closure of more than two weeks (NAEYC, 2020). Evidence of large-scale program closures is already abundant. For example, 57% percent of Florida’s child care providers are temporarily shuttered, while 35% of Louisiana’s programs expect to close permanently (Bryson, 2020; Sonnier-Netto et al., 2020). Moreover, while recent evidence suggests that labor demand collapsed throughout the economy, Figure 1 shows that demand—as measured by online job postings—has fallen especially sharply for child care: these job postings are at roughly 60% of their January 2019 levels, compared to 80% for job postings overall (Kahn et al., 2020; Rojas et al., 2020).

These developments, together with the large literature showing the importance of early childhood education to later schooling and labor market outcomes, highlight the need for systematic evidence on how COVID-19 is influencing the child care market (Herbst, 2017; Ludwig & Miller, 2007; Havnes & Mogstad, 2011; Cunha, Heckman & Schennach, 2010). Moreover, as Dingel et al. (2020) point out, parental constraints in the market for child care may stifle the re-opening of the United States economy given that 32 percent of the workforce has someone in their household under the age of 14.

There are at least two mechanisms through which the pandemic—and the resulting stay-at-home orders—may influence the child care market. First, there could be a direct effect on labor demand, measured using job postings, arising from the mandatory closure of child care services for public health reasons associated with the national containment strategy.¹ Second, there may be an indirect effect in parent demand for child care, measured using search behavior. Given that most states have issued stay-at-

¹ In several states the mandatory closure of child care businesses is explicit. Indeed, businesses in 12 states were ordered to close as per governors’ executive action (Hunt Institute, 2020). In other states child care providers were not deemed “essential” or were given the option to close. Anecdotal evidence suggests that many have exercised this option, electing to furlough or layoff their employees.

home (and mandatory business closure) orders, working parents are now largely at home. In addition, the pandemic may change parent preferences for child care. The rapidly spreading virus could alter the health and safety perceptions of out-of-home child care settings, in which relatively large groups of children interact in close proximity. If a child's perceived risk of contracting the virus is sufficiently high, parents may elect to shift the locus of caregiving from out-of-home services to in-home, parent care so that the environment can be closely monitored. Disentangling these channels is important for understanding the welfare effects of the pandemic on parents: if the labor demand channel is operational, then parents may be unable to self-insure against the shock; otherwise the changes simply reflect evolving demand.

This paper provides the first quantitative evaluation of the impact of pandemic-related containment policy on the supply of and demand for child care in the U.S. To measure labor demand, we make use of unusually rich proprietary data on the number of online job postings for early care and education (ECE) teachers. The measure of search behavior is based on the intensity of internet searches for the topic “child care” extracted from Google Trends. We construct a state-by-day panel of observations over the period January 15 to April 14, 2020. Our identification strategy exploits plausibly exogenous variation in the staggered adoption of state-level “stay-at-home” orders (SAHOs) to estimate their causal effect on child care supply and demand. SAHOs are important policies because they are enforceable by law (Pearl et al., 2020); they have been shown to limit public gatherings while increasing the proportion of people remaining at home (Abouk & Heydari, 2020; Brzezinski et al., 2020); and they have substantially slowed the spread of COVID-19 (Dave et al., 2020). The nation's first SAHO was enacted in California on March 19, 2020; approximately four weeks later, 40 states and the District of Columbia had also issued one. We exploit this state-by-day variation using a difference-in-differences (DD) design along with an event study analysis to examine SAHO-driven changes in child care supply and demand.

Our results suggest that the enactment of a SAHO reduced the number of ECE job postings by over 13% *per day*—implying a reduction in labor demand—but the enactment has no discernible effect on Google searches for child care—implying little impact on demand. To put this in perspective, child care

job postings declined by roughly 40%, relative to the pre-pandemic trend in 2019, so we interpret the SAHO elasticity as economically meaningful. The reduction in overall labor demand is driven by preschool-age teacher positions as well as providers in the private-sector. Indeed, SAHOs do not appear to influence hiring behavior in the market for Head Start and pre-kindergarten teachers. Back-of-the-envelope calculations imply that as many as 1,000 fewer teachers are hired—resulting in a decreased capacity to care for 10,000 fewer children—every month that a SAHO is in effect.

Although Kahn et al. (2020), for example, document a secular decline in labor demand, we provide two pieces of evidence that the ECE sector has been particularly affected by the pandemic—not obvious *a-priori*. First, Figure 1 shows that the number of ECE job postings has fallen much faster over the past few months than those in all other sectors. Second, we provide regression-based evidence that the impact of SAHOs on ECE hiring is nearly twice as large as their impact on education hiring broadly. While there may exist general equilibrium considerations, we provide a static, reduced-form, and early approximation of the impact of SAHOs on the child care market.² Indeed, there are numerous avenues for future work on this topic, including whether these policies (and their subsequent loosening) are associated with long-run changes in the structure of the ECE market, parent preferences for ECE, and child development.

At the broadest level, our paper contributes to large literature on early childhood development and the allocation of time towards child care and home production activities. Dating back to at least Becker (1965), child care has been viewed as a form of household production with a degree of substitutability with market activities (Rogerson, 2007). However, parents value time allocated to child care, especially in comparison with other home production activities (Juster, 1985; Robinson & Godbey, 1990; Krueger et al., 2009). Moreover, the time allocated to child care is increasing in parents' wage rate and educational attainment (Guryan et al., 2008). Although the degree of substitutability between market- and home-based

² We have experimented with several diagnostic exercises to gauge the role of general equilibrium effects. Our main concern is that changes in job postings may reflect a collapse in the demand for child care. To investigate this possibility, we consider fixed effects regressions relating search activity for child care and job postings for early care and education, but we find no statistically or economically meaningful relationship. This is consistent with, although not fully causal evidence of, our interpretation that the costs of these SAHOs may have been borne by households with children who need care and supervision.

care varies based on its quality, there is a general recognition that investments in child development at early ages are particularly important because of the presence of “dynamic complementarities” (Cunha & Heckman, 2006; 2007). If the pandemic decreased the use of market-based child care more by force than by choice, it raises concerns that families may not be fully prepared to transition to home production quickly (or effectively) enough to avoid disruptions to the child development process. This could generate scarring effects on children, much like growing up (Giuliano & Spilimbergo, 2013; McGuire & Makridis, 2020) or graduating (Kahn, 2010; Oreopoulos et al., 2012) during a recession. Given that 13 million (or 60%) of preschool-aged children were regularly attending some form of non-parental child care prior to COVID-19, the scale of the child development effects may be large (Corcoran & Steinley, 2019).




Figure 1: Time Series Variation in Normalized Early Care and Education and Overall Job Postings, January 2019 to April 2020

Notes.—Source: Emsi. The figure plots the normalized monthly number of online job postings in early care and education (ECE) occupations (SOC 25-2011, 25-2012, 25-2021, 25-2052), as well as total number of job postings, between January 2019 and April 2020 (normalized to January 2019).

Our paper also contributes to the policy discussion on the role of family leave and child care policies in parental employment decisions. Although the U.S. has witnessed large increases in the share of employed mothers—thereby fueling the demand for parental leave and subsidized child care—there are persistent gaps in the coverage and generosity of such benefits (Herbst, 2018). For example, the federal Family and Medical Leave Act (FMLA) provides only 12 weeks of *unpaid* leave and covers only 60% of private sector workers (Rossin-Slater & Uniat, 2019). In addition, just six states and the District of Columbia have enacted *paid* leave programs (Bana et al., 2019; Bartel et al., 2018), and there is considerable variation at firm-level in the availability and quality of these benefits.³ As for child care, the largest subsidy programs—the Child Care and Development Fund (CCDF) and Head Start—are means-tested (both), contain strict work requirements (CCDF), or do not operate full-day, year-round programs (Head Start) (Herbst & Tekin, 2016). These constraints imply that such policies may be of limited value to many families for employment purposes. As a result, families either pay for child care services out-of-pocket or shift into unpaid, informal caregiving arrangements (Gathman & Sass, 2018; Herbst, 2018). This discussion suggests that, given the limitations of in-work family supports, coupled with the liquidity boost from the CARES Act, the recent transition to home-based employment may be beneficial to some parents, particularly those who have stronger preferences for spending time with their children, who face greater consequences from taking time off work, or who are highly sensitive to the price of child care.

II. Data Sources and Measurement

Data for this paper come from three sources. We begin by analyzing the *labor demand for child care* using the universe of online early care and education (ECE) job postings obtained from the labor market analytics company EMSI. EMSI's job posting data are advantageous because it combines information from multiple external sources, including Indeed and CareerBuilder, among many others.⁴ Our dataset includes

³ For example, Liu et al. (2019) use data from Glassdoor to show that firms offer higher quality maternity leave benefits in labor markets that contain fewer skilled female workers.

⁴ To our knowledge, these data have not been used previously, although they are comparable to those from Burning Glass Technology, which have been used in some recent papers (e.g., Hershbein & Kahn, 2018).

all job postings pertaining to the two-digit standard occupational classification (SOC) code for education (SOC-25). We then utilized a variety of keyword search methods to locate postings in three key sectors of the center-based ECE market: child care, Head Start (and Early Head Start), and pre-kindergarten. Although we were careful to limit the data to pedagogical (i.e., teaching) positions within each sector, we analyze the full spectrum of such positions, including lead and assistant teachers, teacher's aides, co-teachers, and floating classroom teachers. These positions were advertised by local and state government agencies, for- and non-profit centers (including national chains), places of worship, community-based organizations, and school-based before- and after-school programs. We further categorized the job postings according to the child-age of the classroom that the teacher would operate: infant, toddler, or preschool classrooms, as well as before- and after-school settings. Finally, these data were collapsed into state-by-day cells, giving us the number of ECE job postings over the period January 15 to April 14, 2020.

We then study a measure of *child care search behavior* based on the intensity of internet searches for child care, extracted from Google Trends.⁵ We examine Google search data for the *topic* “child care,” which includes search terms related to child care (i.e., day care), different spellings of the search terms (e.g., childcare), and varying languages. Google Trends creates a score representing the intensity of internet searches for a given term or topic. Specifically, it reflects the number of “child care” searches as a share of the total searches within a particular geographic area and time period. The share ranges between zero and 100 in each area, with zero representing the point in time with the lowest search intensity for “child care” and 100 representing the point in time with the highest intensity. This analysis relies Google search scores for all 50 states and the District of Columbia on each day over the period January 15 to April 14, 2020.

We also collect information on the enactment dates of several state-specific COVID-19 containment policies, importantly the enactment of statewide SAHOs. The SAHO data were collected from multiple sources, including Mervosh et al. (2020), Dave et al. (2020), and the National Governors

⁵ Google Trends have been used to study the demand for religion (Bentzen, 2020), unemployment insurance benefits (Goldsmith-Pinkham & Sojourner, 2020), and health care (Hanna & Hanna, 2019).

Association (NGA). California was the first state to implement a SAHO—on March 19—and 40 states plus the District of Columbia had enacted one as of April 14. See Appendix Table 1 for a list of states operating under a SAHO (and the enactment dates) during our study period.⁶ We also collect information on emergency declarations, mandatory non-essential business closures, and public school closures at the state-level. Enactment dates for these policies were collected from the NGA and Hunt Institute.^{7,8} We also found that, out of the 41 states that implemented a SAHO, 27 have some provision that designates child care as an essential service, further suggesting that our main effects are not driven by mechanical closures.⁹

Figure 2 provides descriptive evidence on the time series pattern of ECE job postings and child care internet search intensity over the period January 14 to April 14, 2020. Although there is substantial day-to-day variation, there is an unambiguous decline in both series starting in mid-March, which coincides with the timing of the national state of emergency declaration and social distancing recommendations.¹⁰ Although the decline is meaningful for both series, it is substantially larger for job postings. As we show in more detail later, this is consistent with our main result that the supply effects appear to dominate the demand side effects. Figure 3 shows the variation in our key independent variable—implementation of a SAHO—by displaying the fraction of states with a SAHO on each day between March 1 and April 14. We see that there is considerable variation in the timing that different states adopted these policies.¹¹

⁶ It is important to note that, in a few states, the stay-at-home order took effect at 11:59pm. Our coding establishes the following day as the first day of implementation.

⁷ The NGA data can be found here: <https://www.nga.org/coronavirus/>. The Hunt Institute data can be found here: <http://www.hunt-institute.org/covid-19-resources/k-12-state-specific-resources/>.

⁸ In some models, we control for the number of COVID-19 cases by day and state. Such data were extracted from the *New York Times* Coronavirus (Covid-19) Data in the United States Database. This database provides information on the cumulative number of cases and deaths by state (and county) and day, starting on January 21. A confirmed case is defined as a patient who tests positive for COVID-19 and is reported as such by a federal, state, or local government agency.

⁹ <https://www.sittercity.com/parents/child-care-essential-business-by-state>

¹⁰ The child care searches, in particular, display a clear within-week pattern, rising during the weekdays and then falling abruptly throughout weekend. This is consistent with the idea that the demand for child care is employment-driven.

¹¹ Out of an abundance of caution that these laws are introduced in response to deteriorating health conditions, we also test for the presence of reverse causality by estimating linear probability models of the timing of these laws on the number of infections; we find no statistically significant effects.




Figure 2: Time Series Variation in Early Care and Education Job Postings and Child Care Internet Search Intensity, January 15 to April 14, 2020

Notes.—Source: Google Trends and Emsi. The figure plots the national number of job postings in early care and education (ECE) and the Google Trends search intensity score for the topic of “child care” on each day between January 15 and April 14, 2020.

III. Identification Strategy

To quantify the impact of the COVID-19 pandemic on the supply of and demand for child care, we exploit plausibly exogenous variation in the staggered adoption of SAHOs across the states. That is, we begin with a variant of a difference-in-differences (DD) estimator by comparing supply and demand in states that adopted these orders sooner versus later than their counterparts:

$$Y_{st} = \gamma SAHO_{st} + \theta X'_{st} + \phi_s + \lambda_t + \epsilon_{st} \tag{1}$$

where Y denotes our measure of the supply or demand for child care in state s and day t , $SAHO$ denotes a binary indicator for whether a given state has implemented a stay-at-home order, and ϕ and λ denote

state and day-of-the-year fixed effects. Standard errors are clustered at the state- and month-levels (Bertrand et al., 2004). Our identifying assumption is that the timing of these SAHOs is plausibly exogenous—that is, that job postings or Google searches for child care in states that adopted a SAHO would have trended similarly to those that did not, conditional on observables and fixed effects. While we recognize that states with Republican versus Democrat governors behave differently, and that states with different demographic characteristics (e.g., population density) may be more at risk, our state fixed effects control for these potential threats to identification. In addition, there may be unobserved national shocks to child care supply and demand (e.g., federal mandates or recommendations as well as presidential announcements); these potential aggregate confounders are eliminated by the day fixed effects.




Figure 3: Time Series Variation in the Adoption of Stay-at-Home Orders

Notes.—Source: Mevosh et al. (2020), Dave et al. (2020), and National Governors Association. The figure plots the fraction of states that implemented a SAHO on each day between March 1 and April 30, 2020.

Nonetheless, one concern is that the supply or demand for child care might be influenced by factors that are also correlated with the timing of states' SAHO enactment. For example, governors

undertook a variety of symbolic and policy-oriented actions contemporaneously with the implementation of SAHOs. Therefore, we include in equation (1) a vector of state-level variables, denoted by \mathbf{X} , to control for the declaration of a state of emergency and a statewide school closure order.¹² We are also concerned that states may adopt a SAHO based on increasing concern among its businesses, which could be correlated with other economic fundamentals that shift the supply and demand for child care. Failure to account for these underlying economic conditions would also lead to biased estimates. To address this concern, in some specifications, we control for the overall number of education-related job postings and the number of COVID-19 infections, which addresses potential time-varying omitted variables that could be correlated with both the passage of state policies and the demand or supply of child care. By controlling for these two terms, we isolate variation that is uniquely affecting child care. As a final test of our identifying assumption, we implement an event study analysis to test our parallel trends assumption, which investigates whether child care demand and supply were already shifting in states prior to the enactment of SAHOs.

IV. Results

Main Results

We begin by estimating equation (1) under several different specifications in Table 1. Columns (1) and (5) show that there is a strong negative association between a SAHO and both ECE job postings (i.e., supply) and child care internet search behavior (i.e., demand). However, these effects are heavily influenced by cross-sectional and temporal confounders, especially those influencing the demand for child care. Once we introduce state and day fixed effects, we find a large decline in the coefficient magnitudes. We nonetheless find a statistically significant 13.1% decline in ECE job postings *per day* following the adoption of a SAHO, but no statistically significant effect on Google searches for child care. While the estimate is negative—consistent with the potential for a demand-side effect—it is indistinguishable from zero.

¹² We have also introduced child care closures as an additional control. While, not surprisingly, the direct effect is negatively correlated with job postings, it is not statistically significant (and lower in magnitude) with the inclusion of a SAHO indicator. Moreover, the interaction effect between child care closure and SAHO is not statistically significant (p -value = 0.33).

Table 1: DD Estimates for the Impact of SAHOs on ECE Job Postings and Google Searches for Child Care

	ln(number of postings for all ECE jobs)				ln(Google Trends search score "child care")			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1[e > SAHO]	-0.439*** (0.084)	-0.131** (0.062)	-0.132** (0.062)	-0.112* (0.059)	-0.501*** (0.067)	-0.058 (0.057)	-0.057 (0.057)	-0.054 (0.056)
1[e > state of emergency]			-0.048 (0.054)	-0.066 (0.051)			0.028 (0.041)	0.025 (0.041)
1[e > public school closures]			0.073 (0.096)	0.034 (0.097)			0.038 (0.066)	0.031 (0.066)
ln(no. all ed job postings)				0.283*** (0.021)				0.054* (0.029)
Observations	4,641	4,641	4,641	4,641	4,641	4,641	4,641	4,641
R-squared	0.03	0.66	0.66	0.68	0.04	0.37	0.37	0.37
State Fixed Effects	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Day Fixed Effects	No	Yes	Yes	Yes	No	Yes	Yes	Yes

Notes.—Sources: Emsi. The table reports the coefficients associated with regressions of logged number of all early care and education (ECE) job postings and the logged Google Trends search intensity score for the topic "child care" on an indicator for the passage of a stay-at-home order (SAHO), conditional on an indicator for a state of emergency declaration, school closure orders, the logged number of overall education job postings, and state and day-of-the-year fixed effects. Standard errors, adjusted for clustering in state and month cells, are in parentheses. * p < 0.10, ** p < 0.05, *** p < 0.01.

Table 2: DD Estimates for the Impact of SAHOs on ECE Job Postings, by Age Group and Sector

	ln(ECE jobs by age-group)				ln(ECE jobs by sector)		
	(1) All ECE	(2) Infant and Toddler	(3) Preschool- Age	(4) School- Age	(5) Child Care	(6) Head Start	(7) Pre- Kindergarten
1 t > SAHIO]	-0.132** (0.062)	-0.028 (0.049)	-0.150*** (0.050)	-0.072*** (0.023)	-0.139** (0.063)	-0.004 (0.018)	-0.030 (0.022)
Observations	4,641	4,641	4,641	4,641	4,641	4,641	4,641
R-squared	0.66	0.47	0.54	0.36	0.65	0.2	0.29
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes.—Sources: Emsi. The table reports the coefficients associated with regressions of logged number of early care and education (ECE) job postings by age-group and sector on an indicator for the passage of a stay-at-home order (SAHO), conditional on an indicator for a state of emergency declaration, school closure orders, and state and day-of-the-year fixed effects. Standard errors, adjusted for clustering in state and month cells, are in parentheses.
* p < 0.10, ** p < 0.05, *** p < 0.01

Table 3: DD Estimates for the Impact of SAHOs on ECE Job Postings, by Minimum Education Level Required and Work Hours Offered

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	All ECE	No Ed Listed	High School	AA Degree	BA Degree	Part-time Hours	Full-time Hours
1 t > SAHIO]	-0.132** (0.062)	-0.100* (0.052)	-0.089* (0.047)	-0.065** (0.031)	-0.070** (0.033)	-0.120*** (0.039)	-0.118* (0.062)
Observations	4,641	4,641	4,641	4,641	4,641	4,641	4,641
R-squared	0.66	0.53	0.47	0.32	0.39	0.44	0.63
State Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes.—Sources: Emsi. The table reports the coefficients associated with regressions of logged number of early care and education (ECE) job postings by minimum education level required and part-time/full-time hours offered on an indicator for the passage of a stay-at-home order (SAHO), conditional on an indicator for a state of emergency declaration, school closure orders, and state and day-of-the-year fixed effects. Standard errors, adjusted for clustering in state and month cells, are in parentheses.
* p < 0.10, ** p < 0.05, *** p < 0.01.

To address potential concerns associated with omitted variable bias, we introduce in columns (3) and (7) two indicators that control for the timing of state of emergency declarations and for school closure mandates. These indicators address the concern that the main effect is coming from other policies that were introduced around the same time as a SAHO. The coefficient on SAHO is robust to the inclusion of these controls in the model for child care supply, while that for demand remains statistically insignificant (although it remains negatively signed). Interestingly, the additional policy controls do not influence ECE supply or demand. Finally, in columns (4) and (8), we add a control for the total number of educational services job postings (excluding ECE job postings), which addresses the concern that there is a common state-specific shock that is correlated with both the enactment of a SAHO and the job posting outcomes. Inclusion of this control reduces somewhat the point estimate on SAHO in the model for ECE job postings, to 11.2%, but it remains statistically significant at the 10% level.

Given the robust association between SAHOs and ECE job postings, we now investigate whether these state policies differentially affect job postings by the age-group of children served and by sector, as shown in Table 2. Column 1 presents our baseline estimate for all ECE job postings, as in Table 1. We find that the adoption of a SAHO is not associated with job postings for infant/toddler teachers [column (2)], but is associated with a 15% reduction in postings for preschool-age teachers [column (3)] and a 7.2% reduction in school-age teacher postings [column (4)]. A potential explanation for this pattern is that parents may be particularly likely to take care of very young children regardless of the pandemic and presence of SAHOs, making the demand for infant/toddler child care more inelastic.

We uncover similar heterogeneity across different ECE sectors. In particular, we find striking evidence that the implementation of SAHOs reduced private-sector child care job postings by 14% [column (5)], but had no effect on job postings for the public-sector Head Start and pre-kindergarten programs [columns (6) and (7)]. That public-sector hiring is less sensitive to SAHO-driven shocks may be attributed to a few factors, including that these services are less sensitive to negative shocks. Indeed, they do comparatively little hiring even during periods of strong growth (in part because they are smaller

programs), and it is possible that they received stimulus funding (independent of the Paycheck Protection Program) to remain operational (e.g., serve the children of essential workers) during the pandemic.

Robustness and Heterogeneity

Our earlier results are identified off of within-state variation in the timing of SAHO adoptions. We have also shown that our main results are robust to the inclusion of overall education job postings, which isolates variation unique to child care. We now present additional robustness exercises that address identification challenges. First, we begin by examining the presence of pre-trends in our DD estimator. We adopt a state-by-day event study design, relying on the staggered adoption of SAHOs at different times. This methodology is useful for examining whether child care demand and supply were already shifting in states prior to the implementation of SAHOs. We estimate the event study model as follows:

$$Y_{st} = \sum_{j=-10}^{10} \gamma_{t+j} d_{st0+j} + \boldsymbol{\theta} \mathbf{X}'_{st} + \phi_s + \lambda_t + \epsilon_{st}$$

where $d_{s,t0+j}$ denotes a set of indicator variables centered around the day on which each state implemented its SAHO, t_0 . We construct an indicator variable for each of the 10 days prior to and after enactment of the policy, using as the benchmark period 15 to 11 days prior to enactment. The event study model includes the other state policy variables as well as the state and day fixed effects.

The event study results for ECE job postings are present in Figure 4.¹³ We estimate a separate version of the model for all ECE postings (Panel A), private-sector child care postings (Panel B), and the public-sectors programs Head Start (Panel C) and pre-kindergarten (Panel D). As shown in Panel A, while the pre-SAHO trends show a small relative decline in the three days prior to policy enactment, we uncover a sharper and larger reduction in overall ECE job postings immediately after enactment. Evidence of pre-trends is even less detectable in the sector-specific analyses shown in Panels B through D. Together, these findings provide support for our methodological approach. Moreover, consistent with our primary DD

¹³ We conduct a similar event study analysis of the Google Trends child care search intensity score. As shown in Appendix Figure 1, there is no evidence of pre-trends. Although in the 10 days following the enactment of a SAHO the impact is consistently negative—again implying a reduction in demand—the estimates are never statistically significant.

results, while we see little evidence of a decline in job postings for Head Start and pre-kindergarten teachers, we see an economically meaningful decline in the 10 days following a SAHO for overall ECE and private-sector child care job postings—roughly a 0.5% to 1.5% decline. As more data becomes available, the confidence intervals on these estimates will also become more precise.




Figure 4: Event Study Estimates for the Supply of Child Care

Notes.—Source: Emsi. The figure investigates the presence of pre-trends by regressing the logged number of job postings associated with overall early childhood education (ECE), child care, Head Start, and pre-kindergarten teachers on 10 daily lagged and 10 daily lead variables, conditional on state and day-of-the-year fixed effects. Standard errors are clustered at the state- and month-level.

Second, since SAHOs were introduced in part as a response to the number of local infections, our estimates might be biased if infections also had a direct effect on ECE job postings. While we view the decline in job postings largely a function of the national quarantine (Kahn et al., 2020), rather than the direct health effects, we nonetheless investigate the robustness of our results to the role that emerging health risks might have played. Specifically, we control for the logged number of cumulative COVID-19

cases. Doing so reduces our estimate on the SAHO indicator from -0.132 (p -value = 0.034) to -0.053 (p -value = 0.340). Although not statistically significant at conventional levels, we view this as an overly strong test given the correlation between confirmed cases and the timing of SAHO adoption.

We also explore potential heterogeneity in the treatment effects. When we estimate our baseline equation for ECE job postings, we find some evidence of heterogeneity by states: the estimated treatment effect on SAHOs is -0.116 in states with Republican governors and -0.053 in states with Democrat governors, but the standard errors are large and do not allow us to reject the null that they are the same. In addition, we find evidence that the reduction in ECE job postings are concentrated in states that are above the median in terms of COVID-19 cases per capita.

Finally, Table 3 explores heterogeneity across categories of job postings according to the minimum level of education required and work hours offered. For reference, column (1) begins with our baseline effect on all ECE job postings. We subsequently allow elasticities to vary by job postings without any educational requirements, those requiring a high school diploma, those requiring an AA degree, and those requiring a BA degree (or more).¹⁴ We find statistically significant declines in each category, although the magnitudes are somewhat smaller for job postings requiring an AA or BA degree. We also find that the elasticities are almost identical when we distinguish between part- and full-time jobs.

Is the SAHO-Driven Reduction in the Demand for ECE Teachers Really Distinctive?

Recall that our baseline DD estimate suggests that the number of ECE job postings fell by over 13% for each day that a SAHO was in effect. How does this compare to the demand for workers in all other education occupations? To investigate whether our ECE effects merely reflect across-the-board reductions in labor demand, we re-estimate equation (1), using the log number of *all* education-related job postings (except ECE) as the outcome variable. Results from this DD model are reported in Appendix Table 2. Column (3) shows that all other education job postings fell by a marginally significant 6.9%

¹⁴ As a point of reference: throughout the month of January—and prior to the implementation of SAHOs—41% of ECE job postings did not specify an education requirement; 30% required a high school diploma; 12% required an AA degree; and 17% required a BA degree or more.

following the enactment of a SAHO. This estimate is a little less than half the magnitude of that in the model for ECE job postings, suggesting that the ECE labor market—more so than other educational services—has been particularly affected by these containment policies.

V. Conclusion

Arguably the most robust state-level policy response to the COVID-19 pandemic has been the implementation of stay-at-home orders (SAHOs). As of April 14, 40 states and the District of Columbia had enacted such a policy in order to slow the spread of the virus and to alleviate any capacity constraints experienced by hospitals and other health care providers. Although SAHOs have been effective at mitigating the spread of COVID-19, in part because of widespread compliance with these orders (at least in the short-run), such policies may have caused substantial job loss and firm closure—perhaps in the short- *and* long-run. One particularly vulnerable, though essential, sector is the market for non-parental child care. Indeed, child care services—many of them small businesses—operate on thin profit margins, with some analyses suggesting that programs must keep enrollments close to maximum capacity in order to stay in business (Workman & Jensen-Howard, 2018). However, the labor demand for child care needs to remain intact for when states end their SAHOs and parents head back to work.

Although some previous work has surveyed child care providers about their plans to close or alter their hiring behavior in the wake of COVID-19, to date no study has quantified the impact of implementing containment policies like SAHOs on the U.S. child care market. The current paper attempts to fill this gap by estimating the short-run impact of SAHOs on labor demand (i.e., number of online job postings) and search behavior (i.e., through internet searches) for child care. Our results suggest a reduction of 13% in the number of ECE job postings *per day*—implying a reduction in supply—but no discernible change in internet searches for child care—implying little change in demand. The reduction in job postings is driven by preschool-age teacher positions as well as providers in the private-sector. Indeed, SAHOs do not appear to influence hiring behavior in the market for Head Start and pre-kindergarten teachers.

How do we interpret these results? While we cannot rule out a potentially negative effect of SAHOs on child care search behavior, it is statistically insignificant. In this sense, we interpret the effect of SAHOs on child care as largely a supply-side mechanism—that is, had the COVID-19 pandemic and the national quarantine not hit, parents would have continued searching for child care and job postings would have continued being posted. One way to assess the magnitude of our results is to compare the DD estimates to the number of ECE job postings prior to the full onset of the pandemic. For example, throughout the month of January, child care providers advertised for 7,723 ECE teacher positions—or an average of approximately 250 positions per day. Therefore, our DD estimates imply that as many as 1,000 fewer teachers are hired for every month that a SAHO is in effect. Given that states mandate a child-to-staff ratio of approximately 10-to-1, on average, in center-based settings, a reduction of 1,000 newly hired teachers means that 10,000 fewer preschool-age children can be cared for each month.

An important point to bear in mind is that we are estimating only the short-run impact of SAHOs. As of this writing, such policies remain in effect in some states and have been lifted in many others. It remains to be seen how the implementation and subsequent removal of SAHOs influences child care demand and supply. Therefore, additional work over the next few months will be needed to trace the full impact of these important containment policies.

References

- About, Rahi & Heydari, Babak. (2020). The Immediate Effect of COVID-19 Policies on Social Distancing Behavior in the United States. Available at SSRN: <https://ssrn.com/abstract=3571421>.
- Bana, S., Bedard, K., & Rossin-Slater, M. (2019). The Impacts of Paid Family Leave Benefits: Regression Kink Evidence from California Administrative Data. NBER Working Paper No. 24438. Cambridge, MA: National Bureau of Economic Research.
- Bartel, A.P., Rossin-Slater, M., Ruhm, C.J., Stearns, J. and Waldfogel, J. (2018), Paid Family Leave, Fathers' Leave-Taking, and Leave-Sharing in Dual-Earner Households. *Journal of Policy Analysis and Management*, 37, 10-37.
- Becker, Gary S. (1965). "A theory of the allocation of time." *Economic Journal* 75, 493-517.
- Bentzen, J.S. (2020). In crisis we pray: Religiosity and the COVID-19 pandemic. *Covid Economics*, 20. Available at: https://cepr.org/active/publications/discussion_papers/dp.php?dpno=14824.
- Bertrand, Marianne, Duflo, Esther, and Mullainathan, Sendhil. (2004). "How much should we trust differences-in-differences estimates?" *Quarterly Journal of Economics* 119(1), 249-275.
- Bryson, A. (2020). "Coronavirus Florida: over half of state's child care facilities are closed", April 13, 2020, The Palm Beach Post.
- Brzezinski, A., Kecht, V., Van Dijcke, D., & Wright, U. (2020). Belief in science influences physical distancing in response to COVID-19 lockdown policies. Working Paper No. 2020-56. Becker Friedman Institute for Economics at University of Chicago.
- Corcoran, L., and Steinley, K. (2019). *Early Childhood Program Participation, From the National Household Education Surveys Program of 2016* (NCES 2017-101.REV), National Center for Education Statistics, Institute of Education Sciences, U.S. Department of Education. Washington, DC.
- Cunha, Flavio and Heckman, James J. (2006). "Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation", *Journal of Human Resources* XLII, 4, 738-782.
- Cunha, Flavio and Heckman, James J. (2007). "The technology of skill formation", *American Economic Review* 97, 31-47.
- Cunha, Flavio, Heckman, James J., and Schennach, Susanne M. (2010). "Estimating the technology of cognitive and noncognitive skill formation", *Econometrica* 78(3), 883-931.
- Dave, D., Friedson, A., Matsuzawa, K., & Sabia, J. 2020. "When do shelter-in-place orders fight COVID-19 Best? Policy heterogeneity across states and adoption time." *IZA Discussion Paper* No. 13190. Bonn, Germany: Institute of Labor Economics.
- Dingel, Jonathan I., Patterson, Christina, and Vavra, Joseph. 2020. "Childcare obligations will constraint many workers when reopening the US economy." *BFI working paper*.
- Giuliano, Paola and Spilimbergo, Antonio. (2014). "Growing up in a recession", *Review of Economic Studies* 81(2), 787-817.
- Goldsmith-Pinkham, P & Sojourner, A (2020). Predicting initial unemployment insurance claims using

- Google Trends. Working Paper. Available at: https://paulgp.github.io/GoogleTrendsUINowcast/google_trends_UI.html.
- Guryan, Jonathan, Hurst, Erik, and Kearney, Melissa. (2008). "Parental education and parental time with children", *Journal of Economic Perspectives* 22(3), 23-46.
- Hanna, A. & Hanna, L.-A. (2019). What, where and when? Using Google Trends and Google to investigate patient needs and inform pharmacy practice. *Int J Pharm Pract*, 27, 80-87.
- Havnes, Tarjei and Mogstad, Magne. (2011). "No child left behind: Subsidized child care and children's long-run outcomes." *American Economic Journal: Economic Policy*, 3,97-129.
- Herbst, C.M. (2017). "Universal child care, maternal employment, and children's long-run outcomes: Evidence from the U.S. Lanham Act of 1940." *Journal of Labor Economics* 35, 519-564.
- Herbst, C.M. (2018). The rising cost of child care in the United States: A reassessment of the evidence. *Economics of Education Review*, 64, 13-30.
- Herbst, C.M. and Tekin, E. (2016), The Impact of Child-Care Subsidies on Child Development: Evidence from Geographic Variation in the Distance to Social Service Agencies. *Journal of Policy Analysis and Management*, 35, 94-116.
- Hershbein, Brad and Kahn, Lisa B. (2018). "Do recessions accelerate routine-biased technological change? Evidence from vacancy postings", *American Economic Review* 108(7), 1737-1772.
- Hunt Institute. (2020). COVID-19 resources. Available at: <http://www.hunt-institute.org/covid-19-resources>. Durham, North Carolina: The Hunt Institute.
- Juster, Thomas. 1985. "Preferences for Work and Leisure." In *Time, Goods, and Well-Being*, ed. F. T. Juster and F. P. Stafford, 333-51. Ann Arbor, MI: Institute for Social Research, University of Michigan.
- Kahn, Lisa. (2010). "The long-term consequences of graduating from college in a recession." *Labour Economics* 17, 2: 303-316.
- Kahn, Lisa, Fabian Lange, and David G. Wiczer. (2020). "Labor demand in the time of COVID-19: Evidence from vacancy postings and UI claims." NBER Working Paper 27061.
- Krueger, Alan, Daniel Kahneman, David Schkade, Norbert Schwarz, and Arthur Stone. (2009). "National Time Accounting: The Currency of Life." In *National Time Accounting and Subjective Well-Being*, ed. Alan Krueger, chap. 1. University of Chicago Press.
- Ludwig, Jens & Miller, Douglas, "Does Head Start improve children's life chances? Evidence from a regression discontinuity design", *Quarterly Journal of Economics* 122 (2007), pp. 159-208.
- Makridis, Christos A. & McGuire, Erin. (2020). Refined by Fire: The Great Depression and Entrepreneurship. SSRN Working Paper. Available at: <https://papers.ssrn.com/abstract=3371991>.
- Makridis, Christos A., Tim Liu, Paige Ouimet, and Elena Simintzi. (2020). The Distribution of Non-Wage Benefits: Maternity Benefits and Gender Diversity. *Review of Financial Studies*, R&R. Available at: <https://papers.ssrn.com/abstract=3088067>.

Mervosh, S., Lee, J., Gamino, L., & Popovich, N. (2020). See Which States and Cities Have Told Residents to Stay Home. *The New York Times*. Available at: <https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html?action=click&module=RelatedLinks&pgtype=Article>.

National Association for the Education of Young Children (NAEYC) (2020). Child care in crisis: Understanding the effects of the coronavirus pandemic. Washington, DC: National Association for the Education of Young Children.

Oreopoulos, Philip, von Wachter, Till, and Heisz, Andrew. (2012). "The short- and long-term career effects of graduating in a recession", *American Economic Journal: Applied Economics* 4(1), 1-29.

Pearl, B., Hunter, L., Lo, K., & Chung, E. (2020). The enforcement of COVID-19 stay-at-home orders. Washington, DC: Center for American Progress. Available at: <https://www.americanprogress.org/issues/criminal-justice/news/2020/04/02/482558/enforcement-covid-19-stay-home-orders/>.

Robinson, John, and Geoffrey Godbey. (1999). *Time for Life*. University Park, PA: Pennsylvania State University Press.

Rogerson, Richard. (2007). "Taxation and market work: Is Scandinavia an outlier?", *Economic Theory* 32, 59-85.

Rojas, F. et al. (2020). Is the Cure Worse than the Problem Itself? Immediate Labor Market Effects of COVID-19 Case Rates and School Closures in the U.S. NBER Working Paper No. 27127. Cambridge, MA: National Bureau of Economic Research.

Rossin-Slater, M. & Uniat, L. (2019). Paid Family Leave Policies And Population Health. Health Affairs. Health Policy Brief. Robert Wood Johnson Foundation.

Sonnier-Netto, L., Cope, H., Falgoust, T., Oakey-Frost, R., & Lewis, R, "The impact of COVID-19 on Louisiana child care providers", April 2020, Louisiana Policy Institute for Children.

Wallach, P. & Myers, J. (2020). The federal government's coronavirus response—public health timeline. Washington, DC: The Brookings Institution.

Workman, S. & Jessen-Howard, S. & (2018). Understanding the true cost of child care for infants and toddlers. Washington, DC: Center for American Progress.


Appendix Table 1: Implementation Dates of Statewide SAHOs

State	Date	State	Date
Alabama	April 04	Montana	March 28
Alaska	March 28	Nebraska	--
Arizona	March 31	Nevada	April 01
Arkansas	--	New Hampshire	March 28
California	March 19	New Jersey	March 21
Colorado	March 26	New Mexico	March 24
Connecticut	March 23	New York	March 22
Delaware	March 24	North Carolina	March 30
District of Columbia	April 01	North Dakota	--
Florida	April 03	Ohio	March 24
Georgia	April 03	Oklahoma	--
Hawaii	March 25	Oregon	March 23
Idaho	March 25	Pennsylvania	April 01
Illinois	March 21	Rhode Island	March 28
Indiana	March 25	South Carolina	April 07
Iowa	--	South Dakota	--
Kansas	March 30	Tennessee	April 01
Kentucky	--	Texas	April 02
Louisiana	March 23	Utah	--
Maine	April 02	Vermont	March 25
Maryland	March 30	Virginia	March 30
Massachusetts	--	Washington	March 23
Michigan	March 24	West Virginia	March 24
Minnesota	March 28	Wisconsin	March 25
Mississippi	April 03	Wyoming	--
Missouri	April 06		

**Appendix Table 2: DD Estimates for the Impact of SAHOs on
All Other Education Job Postings**

	ln(no. all education job postings)		
	(1)	(2)	(3)
1[t > SAHO]	-0.327** (0.131)	-0.072* (0.040)	-0.069* (0.039)
1[t > state of emergency]			0.063 (0.047)
1[t > public school closures]			0.137*** (0.050)
Observations	4,641	4,641	4,641
R-squared	0.01	0.83	0.83
State Fixed Effects	No	Yes	Yes
Day Fixed Effects	No	Yes	Yes

Notes: Standard errors, adjusted for clustering in state and month cells, are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.



Appendix Figure 1: Event Study Estimates for Google Child Care Searches

Notes.—Source: Emsi. The figure investigates the presence of pre-trends by regressing the logged Google Trends child care search intensity score on 10 daily lagged and 10 daily lead variables, conditional on state and day-of-the-year fixed effects. Standard errors are clustered at the state- and month-level.

The potential impact of the Covid-19 pandemic on the welfare of remittance-dependent households in the Philippines¹

Enerelt Murakami,² Satoshi Shimizutani³ and Eiji Yamada⁴

Date submitted: 26 May 2020; Date accepted: 26 May 2020

The outbreak of the Coronavirus Disease 2019 (COVID-19) is inevitably affecting remittance-dependent countries through economic downturns in the destination countries, and restrictions on travel and sending remittances to their home country. This paper explores the potential impacts of the COVID-19 pandemic on the welfare of remittance-dependent households using a dataset collected in heavily remittance-dependent regions in the Philippines prior to the outbreak. First, the empirical model pins down the relationship between the macroeconomic performance of the destination countries, the amount of remittances, and the welfare of households. Second, we use the difference in the IMF's forecasts for the 2020 GDP before and after the COVID-19 crisis to project potential impacts on households caused by the COVID-19 pandemic. Our projection shows that

- 1 This study was conducted as part of the project "Study on Remittances and Household Finances in the Philippines and Tajikistan" carried out by JICA Ogata Sadako Research Institute for Peace and Development. We would like to thank Alvin P. Ang, Jeremiaiah M. Opiniano, and Akira Murata for their leadership and technical contribution during the data collection in the Philippines, and the institute for providing us with the valuable dataset. We also thank Yasuyuki Sawada, Hiroyuki Yamada, Aiko Kikkawa Takenaka, Akio Hosono, Etsuko Masuko, Hiromichi Muraoka, Megumi Muto, Ryosuke Nakata, and Shimpei Taguchi for their constructive comments. The views expressed in this paper are our own and do not represent the official positions of either the JICA Ogata Sadako Research Institute for Peace and Development or JICA.
- 2 Research Fellow, JICA Ogata Sadako Research Institute for Peace and Development.
- 3 Executive Senior Research Fellow, JICA Ogata Sadako Research Institute for Peace and Development.
- 4 Research Fellow, JICA Ogata Sadako Research Institute for Peace and Development.

Copyright: Enerelt Murakami, Satoshi Shimizutani and Eiji Yamada

remittance inflow will decrease by 23-32% and household spending per capita will decline by 2.2-3.3% in one year as a result of the pandemic.

1. Introduction

The Coronavirus disease 19 (COVID-19) is a devastating pandemic with global effects. In only five months since the outbreak of the virus, there have been 4 million confirmed cases and 0.3 million people have died globally, with numbers are still increasing rapidly (WHO Situation-Report 119, 2020). The pandemic is extremely detrimental not only to health but also to economies around the world. The COVID-19 pandemic is undoubtedly one of the largest macro-level shocks to the world economy, as evidenced by the already ominous indicators in business cycles and financial markets. The GDP figures for the first quarter of 2020 have revealed the first view of the catastrophic impact of the COVID-19 crisis on economies around the world. According to the OECD's quarterly national accounts data,¹ quarter-on-quarter GDP growth in China dropped by 9.8% in the first quarter of 2020. The Eurozone countries experienced negative growth of 3.8% on average and the United States reported a 1.2% decline in the same quarter.

While the adverse effects on the economy are revealing at the macro-level, the impact of the COVID-19 pandemic is likely to be heterogenous across countries and individuals, depending on their condition prior to the outbreak. More importantly, the adverse effects may not be confined to the domestic markets but may be transmitted internationally, thus affecting individuals in other countries. This is likely to happen through trade and foreign direct investments of goods and services as well as through human migration, which is especially the case for developing countries.

This paper explores the potential impacts of the COVID-19 pandemic on the welfare of households in a remittance-dependent country, which is likely to be severely exposed to external shocks. The pandemic is expected to substantially reduce the amount of remittances that migrants from developing countries can send home. The World Bank reports that global remittances are estimated to decline sharply by about 20 percent in 2020, the sharpest in recent history, and that remittances to low and middle-income countries are projected to fall by 19.7%.² As the IOM (2020) points out, many migrants

¹ <https://stats.oecd.org/index.aspx?queryid=33940>

² <https://www.worldbank.org/en/news/press-release/2020/04/22/world-bank-predicts-sharpest-decline-of-remittances-in-recent-history>. The decline is projected to be 13% in East Asia and the Pacific region. The projection method is shown at <http://documents.worldbank.org/curated/en/989721587512418006/pdf/COVID-19-Crisis-Through-a->

may lose their jobs or be forced to accept lower wages due to lockdowns or oil price crashes that are currently hitting the economy of their destination countries. Moreover, many migrants may find it technically difficult to send remittances, as some countries enforce stringent movement restrictions and exclude money transfer service providers from the list of “essential services” (World Bank, 2020a). Furthermore, many intended migrants who had been preparing for their departure in the near future will be forced to change their livelihood plans for the coming years. In 2019, 80% of the world’s total remittances flowed to low-and-middle-income countries (World Bank, 2020b);³ therefore, the negative impacts of the COVID-19 outbreak may be more serious in developing countries whose citizens heavily depend on remittances from migrant family members. Indeed, the Filipino Government started to provide a series of direct assistance to the overseas migrant workers and their families in hardship. For example, cash relief is being delivered to the affected migrants and their families in Davao del Sur (one of the regions where we conducted our survey) under the initiative of Overseas Workers Welfare Administration (OWWA).⁴

In order to gauge the potential impact of the pandemic on the welfare of remittance-dependent households, we utilize a household-level dataset which was collected in 2016 and 2017, before the COVID-19 outbreak in the Philippines. The Philippines is a sensible case to study for several reasons. First, it is well known that the Philippines is one of the largest source countries for migrants and is one of the most remittance-dependent countries in the world (Yang, 2011). According to the World Bank (2020),⁵ the remittance inflow to the Philippines was 35,167 million US dollars in 2019 and the country was ranked fourth in the world with an amount close to that of Mexico, which was ranked third. Since the top three countries (India, China, and Mexico) have a larger GDP, the proportion of remittances relative to GDP was much higher in the Philippines

[Migration-Lens.pdf](#)

³ Figures are based on calculations made by World Bank staff using IMF Balance of Payments Statistics database and data releases from central banks, national statistical agencies, and World Bank country desks (<https://www.knomad.org/data/remittances>).

⁴ <https://www.owwa.gov.ph/index.php/news/regional/85-1-600-active-owwa-members-in-davao-del-sur-receive-cash-relief-assistance-from-owwa-xi>.

⁵ The amount of remittance inflow in 2019 was the largest in India (83,131 million US dollars or 2.8% of GDP), followed by China (68,398 million US dollars or 0.5% of GDP) and Mexico (38,520 million US dollars or 3.0% of GDP).

sitting at 9.9%, which shows the significant economic impact of remittance inflow on the economy. Moreover, the countries who host Filipino migrants are among some of the countries most seriously affected by lockdowns and oil price crashes. The number of overseas Filipino workers was estimated at 2.2 million in 2016 with the top destinations being Saudi Arabia (23.8%), the United Arab Emirates (15.9%), Kuwait (6.4%), Qatar (6.2%), Hong Kong (5.6%), and Singapore (5.6%), which combined accounts for two-thirds of total destinations (Philippine Statistics Authority, 2017).⁶ The diversity of destinations implies that the impact of COVID-19 may be heterogenous even among Filipino migrants and their households.

In this paper, we project the potential impact of the COVID-19 shock in destination countries on the welfare of remittance-dependent households in two rural municipalities in the Philippines. Household level data allows us to estimate the empirical relationship between remittance income and household expenditure by 2SLS (two-stage least squares), by instrumenting the remittance income by the weighted average of log GDP per capita of migrants' destination countries. We then impute the hypothetical remittances and household economic outcomes in 2020 under a "no-COVID" scenario using the GDP projection made by International Monetary Fund (IMF)'s "World Economic Outlook" in October 2019 and two "with-COVID" scenarios in April 2020, which the IMF revised taking the economic implications of the COVID-19 pandemic into account. Taking the difference between the predicted outcomes of with-COVID and no-COVID projections provides us with the potential shocks on the remittances and other economic welfare outcomes of remittance-receiving households. Our projections show that remittance inflow will decrease by 23-32% and household spending per capita will decline by 2.2-3.3% in one year, as a result of the pandemic.

The COVID-19 pandemic is an ongoing phenomenon and the situation in the Philippines and destination countries is rapidly changing. The total confirmed cases and deaths caused by COVID-19 has been growing in the Philippines (WHO, 2020). Thus, there is too much uncertainty at this point to fully gauge the effects of the pandemic but

⁶ According to the Stock Estimate of Overseas Filipinos (Commission on Filipinos Overseas, 2013), the top five destination countries for Filipinos were the U.S. (34.5%), followed by Saudi Arabia (10.0%), the UAE (8.0%), Malaysia (7.8%), and Canada (7.1%), all of which account for two thirds of total destinations in 2013. The proportions are the share of the sum of permanent migrants, temporary migrants, and irregular migrants.

we believe in the significance of exploring the potential impacts of COVID-19 on international remittances and remittance-dependent households as a way of informing academics and policy makers to acknowledge policy responses to them in advance. To our knowledge, there has been little research on how the COVID-19 pandemic will affect household expenditure and other economic outcomes using microdata.⁷

This paper proceeds as follows: Section 2 describes the dataset used in this study. Section 3 examines the effect of macroeconomic shocks on household living standards prior to the COVID-19 outbreak. Section 4 performs several projections to gauge the impact of the pandemic on household welfare. Section 5 concludes.

2. Data description

The dataset used in this study is the “Survey on Remittances and Household Finances in the Philippines,” which was conducted by the Japan International Cooperation Agency (JICA).⁸ Within the course of this study, two rounds of household surveys were conducted in two heavily remittance-dependent municipalities in the country: Dingras, Ilocos Norte located in the Northern Luzon Island and Bansalan, Davao del Sur located in the southern island of Mindanao (Figure 1).⁹ As will be explained below, the sampling used in the local household survey statistically represents each municipality. The sample size at baseline in the local household survey was 200 overseas migrant households and 200 non-overseas migrant households in each municipality, which were randomly selected in each area. Throughout this paper, we define a migrant household as a household which, on the date the survey data was collected, had at least

⁷ There are a couple of studies on the consumption response to the pandemic using high-frequency transaction-level microdata. Baker et al. (2020) use daily bank transaction data in the US and find that household spending, particularly on foods, increased due to stockpiling in the early onset of the COVID-19 crisis in March, followed by a sharp decline of 50%. Chen et al. (2020) explore daily transaction data from China for the 12 weeks after the outbreak in January. What they found was a sharp decline in consumption by 32% on average and as high as a 70% drop in Wuhan. The level of consumption bounced back to the pre-COVID level in March, but dropped again by 20% in April due to the elevated risk of a second wave of the outbreak.

⁸ The field survey was conducted by Orient Integrated Development Consultants Incorporated (OIDCI).

⁹ These municipalities were selected to oversample households with overseas migrants. Initial listing to construct a sampling frame of households with migrants requires cooperation from local administrative authorities and public service providers as they are the agencies who keep information on who in a barangay currently reside overseas. With the help of our Filipino research collaborators, these municipalities provided the necessary collaboration and the information for listing.

one member who permanently resides or used to reside in this household but is now currently working or living overseas. Given that the stock of overseas Filipino was about ten million in 2013 (Commission on Filipinos Overseas, 2013), which corresponds to one-tenth of total population, the migrant households were oversampled. Before the field survey, a list was created for both sites from October to December in 2015 with the help of local institutions.¹⁰ A total of 2,429 overseas migrant households and 5,172 non-overseas migrant households were listed in Dingras. For Bansalan, a total of 563 overseas migrant households and 19,797 non-overseas migrant households were listed. Next, stratified random sampling was carried out for each municipality. The barangays within each municipality served as strata and the sample households were randomly selected within each barangay.¹¹ The sample of 200 overseas migrant households was proportionately distributed among the barangays. Once the number of overseas migrant households was allocated among the barangays, the equal number of non-overseas migrant households was selected within each barangay.

The survey instrument contains information on household roster, household spending, budgets and assets, remittance-receiving behaviors, and household savings and loans. The eligible respondents were the primary financial decision-makers in each household. The data was collected using computer-assisted personal interviewing (CAPI) in English, Ilocano, and Cebuano.¹²

After a pilot survey and training of field interviewers, the first-round survey was conducted between 2 August and 22 September 2016 in 31 barangays in Dingras and 25 in Bansalan. The sample size for the first round was 834. The second-round survey was implemented between 15 June and 3 August 2017. The sample size in the second round was 668. The attrition rate was 19.9% (16.6% in Bansalan and 23.2% in Dingras).

¹⁰ For Dingras, information was provided from the Local Government Unit (LGU) of barangays (administrative unit under municipality) and puroks (divisions within a barangay). In Bansalan, the Overseas Filipino Worker (OFW) Family Circle, the Bansalan Cooperative Society (BCS), and Barangay Health Workers (BHWs) provided necessary information for listing.

¹¹ The barangay is the smallest political unit and a subdivision of a city or municipality in the Philippines.

¹² A major challenge was that the list has the names and puroks of the respondents but not their address, and the locations were very diverse with no transportation in some cases. On average, an interview lasted for 3.2 hours.

Figure 1 Location maps of two municipalities



Source: By authors based on the report by Orient Integrated Development Consultants Incorporated (OIDCI).

Table 1 reports the summary statistics of the variables used in this study for all households and households with migrants only.¹³ We use the data from the households that were surveyed during both the first and second rounds. We see that household expenditure is larger and income from domestic sources is smaller for households with migrants. We do not see a difference in saving deposits or loan repayments. “Destination per capita GDP (*ECON*)” refers to the weighted average of per capita GDP for all

¹³ We further investigated the sample selection bias caused by the attrition. When we compare the characteristics at baseline, we see that per capita expenditure is systematically larger and the ages of the heads of household is higher for the attrition households. Thus, households in the sample cover the lower side of the income distribution in the two village economies. The detail comparison is available upon request from the authors.

destination countries and for the Philippines (2016 for the first round and 2017 for the second round) in logarithm as explained in the next section. This variable proxies the household's exposure to macroeconomic performances in different countries. Turning to household characteristics, in households with migrants the heads of the household are slightly older and the household size, including overseas members, is larger. The educational attainment of the head of household is also higher; a quarter of the households are composed of college graduates or higher and the proportion of elementary school graduates is smaller. Looking at the occupations of the heads of migrant households, the proportion of managers is larger and agricultural taking is less prominent. Both the variables on education and those of occupation are binary, taking non-educated and non-working of heads of households as the reference.¹⁴

3. Empirical analysis

We use the data collected before the COVID-19 outbreak to examine the relationship between overseas remittance and a variety of outcomes on household welfare to gauge the impact of exogenous shocks on the living standards of households in the Philippines. To do so, we set up two types of specifications.

First, we directly examine the effect of remittances on a variety of outcomes to indicate household welfare. The specification is described as follows:

$$Y_{it} = \beta_0 + \beta(REMITTANCE_{it}) + \gamma X_{it} + barangay_i + \lambda_t + \epsilon_{it} \quad (1)$$

where i indexes households, and t refers to the survey round with 0 indicating 2016 and 1 indicating 2017. The dependent variables Y_{it} are a logarithm of (1) average monthly household expenditure per capita, (2) average monthly household incomes from domestic sources per capita, (3) average monthly household savings deposits per capita, and (4) average monthly household loan repayments per capita.¹⁵ The main explanatory variable " $REMITTANCE_{it}$ " is log average monthly overseas remittance income per capita. Y_{it}

¹⁴ It should be noted that while the job of a seamen makes up a large part of the migrant job market in the Philippines, our sample does not contain those migrants.

¹⁵ The denominator of all "per capita" variables from the household survey is the number of household members excluding migrating members.

Table 1 Summary statistics

VARIABLES	All households					Households with migrants				
	N	mean	sd	min	max	N	mean	sd	min	max
log per capita HH expenditure*	1,296	8.164	0.972	5.694	13.26	560	8.394	0.947	5.975	13.26
log per capita Income from domestic sources*	1,296	6.461	2.453	0	11.96	560	6.155	2.911	0	11.51
log per capita new saving deposit*	1,296	0.811	1.870	0	10.06	560	0.922	2.029	0	10.06
log per capita loan repayments*	1,296	1.319	2.726	0	13.19	560	1.238	2.657	0	13.19
log per capita remittance income*	1,296	3.473	3.779	0	11.41	560	7.006	2.169	0	11.41
Destination per capita GDP (ECON)**	1,296	8.522	0.678	7.981	10.62	560	9.167	0.517	7.981	10.62
Head's age	1,296	51.73	13.85	20	95	560	53.36	14.45	20	95
HH size including overseas members	1,296	4.945	2.143	1	15	560	5.495	2.283	2	15
Head's educational attainment										
Elementary	1,296	0.407	0.492	0	1	560	0.368	0.483	0	1
General High School	1,296	0.349	0.477	0	1	560	0.345	0.476	0	1
Technical Vocational	1,296	0.0409	0.198	0	1	560	0.0446	0.207	0	1
Post Secondary	1,296	0.00926	0.0958	0	1	560	0.00536	0.0731	0	1
College or more	1,296	0.185	0.389	0	1	560	0.234	0.424	0	1
Head's occupation										
Manager	1,296	0.0293	0.169	0	1	560	0.0429	0.203	0	1
Professional	1,296	0.0262	0.160	0	1	560	0.0214	0.145	0	1
Clerical	1,296	0.0147	0.120	0	1	560	0.0179	0.133	0	1
Service	1,296	0.0903	0.287	0	1	560	0.0750	0.264	0	1
Agriculture	1,296	0.271	0.445	0	1	560	0.246	0.431	0	1
Production	1,296	0.0231	0.150	0	1	560	0.0214	0.145	0	1
Municipality (=0 if Bansalan, = 1 if Dingras)	1,296	0.486	0.500	0	1	560	0.436	0.496	0	1

Note: The Authors' calculation.

* Original monetary value is in 1 + PhP (Philippines Peso)

** Refers to weighted average of per capita GDP of all destinations including the Philippines in logarithm.

and $REMITTANCE_{it}$ are calculated as the monthly average either over the past 12 months for the first round or for the period since the first round visit in the case of the second round.¹⁶ \mathbf{X} is a vector of household characteristics that includes the age of the household head, household size, the educational attainment level of the household head, and their occupation. We also include barangay fixed effect ($barangay_i$) and survey round fixed effect (λ_t). Lastly, ϵ_{it} is a well-behaved error term.

We first employ ordinary least squares (OLS) estimation to obtain the coefficients by pooling the observations for all households (two observations for each household). While the dataset is longitudinal, the interval is short (less than one year) and we see little change in the amount of remittances, the main variable, during the survey period. Thus, we use a level specification by pooling the observations at the first and second rounds, rather than a difference specification including a fixed effect model as it is difficult to obtain stable estimation results. There may be some concern about the endogeneity issue since household welfare outcomes are likely to be affected by remittances and vice versa. Thus, we use specification (1) as “preliminary” to obtain the pooled OLS estimates needed to construct the correlation between the remittances and the outcome variables of interest, and to address the issue of endogeneity in our second specification using an instrumental variable (IV) approach.

Table 2 shows the results of the estimations. Column (1) shows the coefficients when the dependent variable is total income from domestic sources and the coefficient on remittance income is negative and significant; this indicates that a 1% increase in remittance income decreases domestic income by 0.055%. In other words, above 5% change in remittances is compensated by a change in domestic income. Column (2) shows that the coefficients when the dependent variables are household spending per capita. The coefficient on the remittance income is positive and significant. The result means that a 1% increase in remittance income increases household expenditure per capita by 0.093%. Column (3) and (4) illustrate the coefficients when the dependent variable is new deposit savings and loan repayments, respectively. The coefficient is positive for new deposit savings and negative for loan repayments; however, neither is statistically significant.

¹⁶ Since the interval between the first and second round surveys is only 10 to 11 months, we use the value of the monthly average since the baseline visit. The qualitative results are not changed if we use the average over the past 12 months for the second round.

Table 2 Estimation results (Pooled OLS)

VARIABLES	(1)	(2)	(3)	(4)
	log per capita Income from domestic sources	log per capita HH expenditure	log per capita new saving deposit	log per capita loan repayments
log remittance income per capita	-0.0545* (0.0310)	0.0930*** (0.0118)	0.0384 (0.0323)	0.0175 (0.0354)
Head's age	0.0583** (0.0268)	0.0335 (0.0231)	0.0525 (0.0715)	0.165** (0.0831)
Square of Head's age	-0.000590** (0.000267)	-0.000261 (0.000215)	-0.000658 (0.000666)	-0.00178** (0.000780)
HH size including overseas members	-0.104*** (0.0327)	-0.142*** (0.0231)	-0.0417 (0.0577)	0.0404 (0.0657)
Head's educational attainment				
Elementary	2.767*** (0.928)	0.476* (0.289)	0.0678 (0.431)	-0.193 (0.739)
General High School	3.003*** (0.935)	0.529* (0.293)	0.0955 (0.470)	-0.164 (0.772)
Technical Vocational	3.182*** (0.998)	1.088*** (0.341)	0.664 (0.622)	-0.397 (0.866)
Post-Secondary	1.944 (1.193)	0.401 (0.411)	1.253 (0.821)	-0.473 (0.957)
College or more	3.368*** (0.946)	0.964*** (0.305)	1.444** (0.564)	0.493 (0.838)
Head's occupation				
Manager	0.265 (0.211)	0.215 (0.156)	1.306 (1.076)	0.493 (0.783)
Professional	0.453 (0.287)	-0.423 (0.259)	-1.526** (0.689)	-0.905 (1.289)
Clerical	0.175 (0.658)	0.556 (0.408)	0.759 (1.004)	0.476 (1.393)
Service	0.725*** (0.247)	0.134 (0.190)	-0.887*** (0.327)	0.307 (0.573)
Agriculture	0.523*** (0.176)	0.116 (0.0940)	0.384 (0.326)	0.458 (0.385)
Production	0.722*** (0.224)	0.506* (0.290)	-0.789 (0.663)	1.898 (1.439)
Constant	3.677*** (1.144)	7.613*** (0.643)	0.0370 (1.928)	-1.367 (2.193)
Barangay fixed effect	YES	YES	YES	YES
Survey round fixed effect	YES	YES	YES	YES
Observations	1,296	1,296	1,296	1,296
R-squared	0.228	0.307	0.234	0.361

Robust standard errors in parentheses

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

We now turn to the second specification. In this specification, we try to establish the impact of the macroeconomic conditions in the destination countries on the outcomes relating to household living standards through remittances. In order to achieve this goal, we employ a two-stage least squares (2SLS) estimation using an index of the macroeconomic performance of the destination countries as an instrumental variable.¹⁷ We construct the “economic performance (*ECON*)” variable by taking the weighted average per capita GDP of the country of residence of each household member, including overseas migrants. More specifically, the “*ECON*” variable is constructed as:

$$ECON_{it} = \ln \frac{\sum_{k \in \mathcal{K}(i)} g_{kt} \times n_{kit}}{\sum_{k \in \mathcal{K}(i)} n_{kit}}$$

Here, $\mathcal{K}(i)$ refers to the set of countries where the members of household i live, g_{kt} is the log GDP per capita in country k in t (2016 or 2017), and n_{kit} is the number of household i 's adult member who live in country k . Thus, by construction, we use per capita GDP in the Philippines for households without migrants.¹⁸

In the first stage, we regress the amount of remittances on the logarithm of the “*ECON*” variable and other covariates.¹⁹

$$REMITTANCE_{it} = \beta_0 + \beta(ECON_{it}) + \gamma X_{it} + \text{barangay}_i + \lambda_t + \epsilon_{it} \quad (2)$$

The notations are the same as in specification (1). This specification exploits cross-country variations of GDP per capita to explain variations in the amount of remittance

¹⁷ We use the sample of households with migrants and without migrants to maximize the observation so as to obtain stable coefficients; however, the results of estimation is largely comparable if we limit our sample to households with migrants only.

¹⁸ We assume that GDP per capita is exogenous to the amount of remittances in each household. We acknowledge a possibility that high endowment migrants are also likely to choose a high income destination country.

¹⁹ Dean and Martínez (2006) and Dean (2008) used the appreciation of the Philippine peso during the 1997 Asian financial crisis as an exogenous shock to international remittances and found the elasticity of Philippine-peso remittances with respect to the exchange rate was estimated to be 0.60 affecting positively on capital accumulation and entrepreneurship in origin households. Dean and Choi (2007) used regional rainfall in the Philippines as an instrumental variable for domestic income changes to explore whether remittances served as insurance for sending households. They found that consumption does not respond to domestic income shocks to households with migrant members but to those without.

across households, rather than exploiting within-household variations of remittances between the two survey rounds because the change in *ECON* is small between 2016 and 2017 and the remittance flows for each household is quite stable.

Column (1) of Table 3 shows the result of the first stage regression. The coefficient on “*ECON*” is positive and significant and indicates that a 1% increase in “*ECON*” leads to a 3.997% increase in income from remittances per capita. This implies that a significant economic recession in the destination countries as projected for the current COVID-19 pandemic, will lead to a substantial drop in remittances.

Next, we use the estimated dependent variable of remittances at the second stage regression.

$$Y_{it} = \beta_0 + \beta(\overline{REMITTANCE}_{it}) + \gamma X_{it} + \text{barangay}_i + \lambda_t + \epsilon_{it} \quad (3)$$

Again, the notations are the same as in (1) with the exception of the main independent variable, which is now the estimated value obtained from the first-stage regression (2).

Columns (2) – (5) of Table 3 convey the second stage of the 2SLS estimation results. Column (2) shows that the coefficient on remittance income is negative and significant in the case of income from domestic sources. The size of the coefficient is 0.107, showing that the degree to be compensated by domestic income is larger in this specification than in specification (1). Column (3) indicates that the coefficient on the remittance income is positive and significant for household spending per capita and the size is 0.099, which is comparable to specification (1). Columns (4) and (5) illustrate that the coefficient is positive for saving deposits and negative for loan repayments, but both are statistically insignificant, which is the same as in specification (1).

In summary, the estimation results using both specifications confirm that a decline in remittances discourages household spending per capita and is partly compensated for by domestic income.²⁰

²⁰ As a robustness check, we ran the regression separately for the first and second-round observations and obtained qualitatively the same results.

Table 3 Estimation results of 2SLS

VARIABLES	(1)	(2)	(3)	(4)	(5)
	First Stage Results	Second Stage Results			
	log per capita remittance income	log per capita domestic Income	log per capita HH expenditure	log per capita new saving deposit	log per capita loan repayments
ECON (Instrumental Variable)	3.997*** (0.131)				
log remittance income per capita		-0.107*** (0.0346)	0.0987*** (0.0126)	0.0362 (0.0321)	-0.0131 (0.0339)
Head's age	-0.0538** (0.0244)	0.0521** (0.0265)	0.0342 (0.0225)	0.0523 (0.0695)	0.162** (0.0807)
HH size including overseas members	0.000685*** (0.000254)	-0.000517* (0.000266)	-0.000269 (0.000210)	-0.000654 (0.000648)	-0.00174** (0.000758)
Head's educational attainment	0.104*** (0.0311)	-0.0957*** (0.0329)	-0.143*** (0.0228)	-0.0414 (0.0554)	0.0454 (0.0633)
Elementary	1.362*** (0.356)	2.871*** (0.902)	0.464* (0.279)	0.0723 (0.421)	-0.132 (0.715)
General High School	1.614*** (0.361)	3.128*** (0.908)	0.515* (0.283)	0.101 (0.462)	-0.0907 (0.748)
Technical Vocational	1.311*** (0.487)	3.308*** (0.969)	1.074*** (0.332)	0.669 (0.615)	-0.324 (0.841)
Post-Secondary	1.589 (1.079)	2.042* (1.131)	0.390 (0.397)	1.257 (0.799)	-0.416 (0.928)
College or more	1.860*** (0.383)	3.533*** (0.922)	0.946*** (0.297)	1.451*** (0.557)	0.589 (0.817)
Head's occupation					
Manager	-0.0287 (0.144)	0.280 (0.207)	0.213 (0.151)	1.306 (1.046)	0.501 (0.762)
Professional	0.156 (0.340)	0.441 (0.280)	-0.422* (0.252)	-1.526** (0.669)	-0.912 (1.255)
Clerical	-0.661* (0.355)	0.166 (0.639)	0.557 (0.397)	0.758 (0.975)	0.471 (1.352)
Service	-0.409** (0.164)	0.701*** (0.241)	0.137 (0.185)	-0.889*** (0.319)	0.293 (0.559)
Agriculture	-0.0154 (0.174)	0.520*** (0.173)	0.116 (0.0912)	0.384 (0.316)	0.456 (0.374)
Production	0.138 (0.397)	0.735*** (0.221)	0.505* (0.281)	-0.788 (0.643)	1.906 (1.401)
Constant	-30.77*** (1.571)	3.850*** (1.119)	7.594*** (0.625)	0.0444 (1.883)	-1.265 (2.133)
Barangay fixed effect	YES	YES	YES	YES	YES
Survey round fixed effect	YES	YES	YES	YES	YES
Observations	1,296	1,296	1,296	1,296	1,296
R-squared	0.569				
F-test for weak IV (p-value)	990.06 (0.00)				

Robust standard errors in parentheses

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

4. Projections

Using the coefficients obtained in the 2SLS estimation we now project the impacts caused by the COVID-19 pandemic. To quantify the scale of the economic shocks caused by the COVID-19 pandemic on the relevant countries, we use the per capita GDP predictions available for each country in 2020 from the International Monetary Fund (IMF)'s "World Economic Outlook" published in October 2019 and April 2020 (Appendix Table 1). The outlook from October 2019 can be seen as a "no-COVID" forecast which helps us to construct the hypothetical "*ECON*" variable in the case where a global COVID-19 pandemic had not taken place (column (1) of Appendix Table 1). Conversely, the revised outlook from April 2020 can be used to construct the "with-COVID" economic scenarios that will affect remittances from migrant workers. There are two forecasts in the April 2020 outlook: the "with-COVID scenario one" is based on the assumption that the COVID-19 pandemic will fade out in the second half of 2020, followed by a gradual lifting of containment measures, which corresponds to the "baseline" of the 2020 April outlook. The "with-COVID scenario two" assumes that the pandemic lasts roughly 50% longer and global output is 3% lower than scenario 1 (IMF, 2020).²¹ While some countries might have a downward revision of the GDP prediction in 2020 without the COVID-19 outbreak, we implicitly assume that the change in the prediction of GDP in 2020 in the two different timings (October 2019 and April 2020) is entirely attributed to the pandemic.

We compute the predicted values by plugging the hypothetical *ECON* variables constructed using each of the three different GDP per capita forecasts for remittance-receiving households into our 2SLS estimates. We then compare the mean predicted values for the various outcome variables in each projection scenario. The difference between using the "with-COVID" and the "no-COVID" scenario captures the potential impact of the COVID economic shock on the economic welfare of households. Our

²¹ The "with COVID" scenario 1 (column (2) of Appendix Table 1) corresponds to the IMF's baseline projection and depends on the assumption that the negative effects are concentrated mostly in the second quarter of 2020 (except for in the case of China) and 8% of the working days in 2020 are lost in heavily affected countries (5% in other countries), which is exacerbated by tighter financial conditions, weaker external demand, and poorer terms of trade caused by stagnant commodity prices (IMF, 2020). The "with COVID" scenario 2 corresponds to the IMF's alternative scenario (column (3) of Appendix Table 1) in which the GDP growth rate is uniformly reduced by 3 percent from scenario 1 as the GDP estimates in the alternative scenario are not available by country.

prediction depends on several assumptions. First, we assume that the adverse effects caused by the COVID-19 pandemic through a variety of factors including restricting migrants by way of a ban on international movement, limiting remittance transactions by labelling it a non-essential activity, and a stagnant economy of host countries, is summarized in a negative change in per capita GDP; this is then used to construct the *ECON* variable at the first stage regression. Second, in the projection of the potential impact on expenditure, savings, and loan repayments, we do not consider the compensating effect of domestic income on the decline of remittances sufficient to smooth changes in overall income levels, which was shown by Column (1), Table 2 and Column (2), Table 3. This assumption is reasonable given that the Philippine economy is also seriously affected by the pandemic; the GDP growth rate for 2020 is projected to be negative 1.7% (IMF, 2020).

Table 4 shows the predicted potential impacts of the COVID-19 pandemic. Column (1) shows the effect of using the “with-COVID scenario 1”. We observe that the negative impact of the pandemic on remittances is serious, with a decline of 23.2%. In monetary value, the average monthly remittance of 2,670 PhP projected under a non-COVID scenario, may decline to 2,051 PhP. This result, which is based on our microdata, is quite comparable with the World Bank’s forecast of a 19.7% global decline in remittance for the year 2020 based on macro data. In this case, household spending per capita would decline by 2.2%. Column (2) reports the impacts of the pandemic using the “with-COVID scenario 2”. We see that the adverse effects are more pronounced; the negative effect on remittances is 32.4% for all households. Household spending per capita would reduce by 3.3%, which would likely put significant negative pressure on household living standards.

In sum, our predictions show that remittance inflow will decrease by 23-32% and household spending per capita will decline by 2.2-3.3% in the space of one year during the COVID-19 pandemic. There is a growing uncertainty about how the COVID-19 pandemic is affecting the world economy and it is difficult to forecast the consequences of the disaster. The potential impacts we presented, must therefore be understood in conjunction with a number of reservations, which we set out below.

First, we use household data from heavily remittance-dependent regions in the Philippines prior to the outbreak. Thus, our estimates do not necessarily conform to the average in the Philippines. As we discussed, the households which dropped from the

Table 4 Potential impacts on household welfare

	2020 forecast, pesos per month			Percent changes	
	No-COVID	With-COVID Scenario 1	With-COVID Scenario 2	With-COVID Scenario 1	With-COVID Scenario 2
Remittances	2,670.2	2,051.3	1,804.9	-23.2	-32.4
Per-capita household expenditure	5,066.3	4,956.0	4,897.0	-2.2	-3.3
Savings Deposit (new saving)	3.87	3.84	3.83	-0.6	-0.9
Loan repayment	1.57	1.58	1.59	0.9	1.4

Note: This table summarizes the difference of projected outcomes between the non-COVID scenario and the two with-COVID scenarios (scenario 1 and scenario 2).

Scenario 1 is based on the IMF baseline projections for per-capita GDP growth for 2020, assuming that the pandemic fades in the second half of 2020.

Scenario 2 is based on the IMF alternative projections for per-capita GDP growth for 2020, assuming that the pandemic lasts longer in 2020.

sample during the survey period enjoyed larger household spending and the heads of households were older. Our projection therefore depends on a sample of households in lower income distribution in remittance-dependent municipalities.

Second, our projection is based on the 2020 IMF outlook. Thus, we capture a short-run effect of the COVID-19 pandemic on household welfare. The negative impact would be more serious if we consider it over a longer horizon. Indeed, IMF (2020) provides more severe scenarios assuming that there is a second outbreak in 2021 that is roughly two-thirds as severe as in the baseline, or that it takes longer to contain the outbreak in 2020 and that there is a second outbreak in 2021. Under those more devastating scenarios, the effect of the COVID-19 pandemic would be non-linear and more detrimental to the economy.

Third, we summarized all aspects of the virus outbreak into a change in per capita GDP. There are several channels for international transmission of the COVID-19 pandemic through a ban on international movement, limiting remittance transactions by deeming them a non-essential activity, and a stagnant economy in host countries. While we implicitly assume that the former two components are also captured by change in per capita GDP, we may need to take a more nuanced approach using data on international restrictions on travels and remittance transactions.

Fourth, migration and remittances depend on a serial decision-making process within a household, which involves several selection steps; whether to migrate or not, who migrates where to get which job, whether to stay abroad permanently or seasonally, and whether or not to remit money home and if so, how much. In this paper, we boldly sum up all of these complex processes into one variable - the amount of remittance. Disentangling the effect of the COVID-19 pandemic on these different steps in the migration process, is no doubt an important agenda for future research.

5. Conclusion

This paper explores potential impacts of the COVID-19 pandemic on the welfare of households in the Philippines, a remittance-dependent country that is likely to be exposed to external shocks. We utilize a household-level dataset in heavy migrant-dependent regions before the outbreak in the Philippines. We estimate the effect of adverse macroeconomic shocks in destination countries on remittance levels and household living standards and use the estimates to gauge the potential impact of the COVID-19 pandemic using a revision of the 2020 GDP projection made by the IMF's World Economic Outlook. Our projection shows that remittance inflow will decrease by 23-32% and household spending per capita will decline by 2.2-3.3% in one year as a result of the pandemic.

The pandemic is still ongoing and the situation in the Philippines and destination countries is rapidly changing. We acknowledge the difficulty of forecasting the adverse effects it may have but we believe that our projection could be useful for academics and policy makers. Future research should use the actual data in migrant-sending countries after the COVID-19 outbreak to quantify the adverse effects on household living standards. While it is not easy to conduct a survey during the pandemic, together with our projection, this line of research will be very informative for future policy responses.

References

- Baker, Scott R., Farrokhnia, R.A., Meyer, Steffen, Pagel, Michaela, and Yannelis, Constantine. (2020). “How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic,” *NBER Working Paper Series* 26949.
- Chen, Haiqiang, Qian, Wenlan, and Wen, Qiang. (2000). “The Impact of the COVID-19 Pandemic on Consumption: Learning from High Frequency Transaction Data.” Available at SSRN: <https://ssrn.com/abstract=3568574> or <http://dx.doi.org/10.2139/ssrn.3568574>.
- Commission on Filipinos Overseas, Department of Foreign Affairs, and Philippine Overseas Employment Administration. (2013). *Stock Estimate of Overseas Filipinos*. Retrieved from <https://cfo.gov.ph/yearly-stock-estimation-of-overseas-filipinos>.
- International Monetary Fund (IMF). (2019). *World Economic Outlook, October 2019 Global Manufacturing Downturn, Rising Trade Barriers*. Retrieved from <https://www.imf.org/en/Publications/WEO/Issues/2019/10/01/world-economic-outlook-october-2019>.
- International Monetary Fund (IMF). (2020). *World Economic Outlook, April 2020: The Great Lockdown*. Retrieved from <https://www.imf.org/en/Publications/WEO/Issues/2020/04/14/weo-april-2020>.
- International Organization for Migration (IOM). (2020). *COVID-19 Analytical Snapshot #16: International Remittances*. Retrieved from https://www.iom.int/sites/default/files/documents/covid-19_analytical_snapshot_16_-_international_remittances.pdf.
- Philippine Statistics Authority. (2017). *Survey on Overseas Filipinos 2016: A Report on the Overseas Filipino Workers*. Retrieved from <https://psa.gov.ph/sites/default/files/2016%20Survey%20on%20Overseas%20Filipinos.pdf>.
- World Health Organization (WHO). (2020). *Coronavirus Disease (COVID-19) Situation-Report 119*. Retrieved from <https://www.who.int/docs/default-source/coronaviruse/situation->

[reports/20200518-covid-19-sitrep-119.pdf?sfvrsn=4bd9de25_4](https://blogs.worldbank.org/psd/remittances-times-coronavirus-keep-them-flowing).

- World Bank. (2020a). *Remittances in times of the coronavirus-keep them flowing*. Retrieved from <https://blogs.worldbank.org/psd/remittances-times-coronavirus-keep-them-flowing>.
- World Bank. (2020b). *Remittances Data*. Retrieved from <https://www.knomad.org/data/remittances>.
- Yang, Dean. (2008). “International Migrations, Remittances and Household Investment: Evidence from Philippine Migrants’ Exchange Rate Shocks.” *Economic Journal* vol.118 (April): 591–630.
- Yang, Dean. (2011). “Migrant Remittances.” *Journal of Economic Perspectives* vol.25(3): 129–152.
- Yang, Dean and Hwajung Choi. (2007). “Are Remittances Insurance? Evidence from Rainfall Shocks in the Philippines.” *World Bank Economic Review* vol.21(2): 219–248.
- Yang, Dean and Claudia Martínez. (2006). “Remittances and Poverty in Migrants’ Home Areas: Evidence from the Philippines.” Ç. Özden and M. Schiff eds. *International Migration, Remittances, and the Brain Drain* Chapter 3: 81–123. Washington, DC: World Bank.

Appendix Table 1 GDP per capita in the IMF Scenarios for 2020 in no-COVID and with-COVID cases

	(Per-capita GDP, constant 2010 USD)		
	(1)	(2)	(3)
	2020 Scenarios		with-COVID
	no-COVID	with-COVID 1	2
Australia	56,768	52,449	50,766
Austria	49,751	45,528	44,054
Bahamas	27,385	25,650	24,836
Bahrain	22,348	20,932	20,268
Belgium	47,287	43,273	41,872
Brazil	11,111	10,342	10,012
Canada	51,560	47,314	45,780
China	7,710	7,374	7,155
Cyprus	31,110	28,995	28,076
Denmark	63,674	58,154	56,272
Egypt	2,845	2,665	2,581
France	43,432	39,819	38,529
Germany	47,480	43,633	42,225
Hong Kong	39,206	37,042	35,904
Israel	34,585	32,394	31,367
Italy	35,344	31,975	30,922
Japan	48,826	46,114	44,661
Jordan	3,295	3,086	2,988
Kuwait	33,693	31,558	30,557
Lebanon	6,334	5,933	5,744
Macao	58,285	55,069	53,376
Malaysia	12,175	11,530	11,178
Norway	92,922	84,866	82,119
Oman	16,196	15,199	14,717
Palau	12,255	11,467	11,103
Phillipines	2,994	2,835	2,749
Qatar	64,269	60,196	58,287
Russia	11,688	10,850	10,506
Saudi Arabia	20,735	19,825	19,204
Singapore	59,635	55,379	53,677
South Africa	7,446	6,938	6,714
South Korea	27,015	25,524	24,740
Spain	32,735	29,571	28,603
Sweden	58,228	53,180	51,459
Taiwan	25,213	23,822	23,090
United Arab Emirates	41,105	38,500	37,279
United Kingdom	43,398	40,000	38,710
United States	54,157	49,941	48,341

Note: Scenario "no-COVID" is based on the IMF's projection of GDP in 2020 as of October 2019.

Scenario "with-COVID 1" is based on the IMF projections for per-capita GDP growth for 2020 as of April 2020, assuming that the pandemic fades in the second half of 2020.

Scenario 2 is based on the IMF alternative projections for per-capita GDP growth for 2020 as of April 2020, assuming that the pandemic will last longer in 2020.

Source: World Development Indicators and World Economic Outlook

Which measures flattened the curve in Germany?¹

Enzo Weber²

Date submitted: 26 May 2020; Date accepted: 28 May 2020

We evaluate the effects of containment measures on flattening the COVID-19 infection curve in Germany. Constructing a regional daily panel dataset, we make use of the fact that different containment measures were implemented by the German state governments at different times and not uniformly nationwide. The results show that the cancellation of mass events, school and childcare closures and curfews played an important role, just as further unobserved factors beyond government interventions. In contrast, we find only limited evidence for additional effects of the closures of service sectors in public life.

- 1 I am grateful to Tobias Hartl, Maximilian Studtrucker and Anja Bauer for support. Of course, any errors are mine.
- 2 Head of Research Department, Institute for Employment Research (IAB), and Full Professor, University of Regensburg.

Copyright: Enzo Weber

1. Introduction

In spring, the corona virus spread exponentially in many countries. During the second half of March, in Germany, as elsewhere, comprehensive containment measures were taken. In the public sector, for example, companies and institutions in sectors such as hotels and gastronomy, retail or culture and leisure were closed. Furthermore, childcare and schools were discontinued and mass events prohibited. The borders with most neighbouring countries were closed, followed by contact restrictions and curfews.

Since then, the growth rate of confirmed COVID-19 cases in Germany has declined sharply (Hartl et al. 2020a). In the meantime, worldwide there is intensive discussion about relaxing the measures. The key task is to develop a strategy for reactivating the economy without risking a medical overload (Baldwin 2020). The effects of such an easing will only become apparent over time. However, it is already of great importance to determine which measures played what part in containing the virus. That is the aim of this article.

Methodologically, we make use of the fact that different containment measures were implemented by the German state governments at different times and not uniformly nationwide. The resulting temporal and regional variation in the introduction of the measures allows us to estimate the effects on slowing down the spread of the virus. This approach complements recent studies on epidemiological modelling such as Ferguson et al. (2020) or Dehning et al. (2020) and on estimating containment effects in an international cross section (e.g., Banholzer et al. 2020, Jüni et al. 2020, Deb et al. 2020). While the latter studies may benefit from larger data variation, a regional approach on a national basis has advantages such as comparable epidemiological conditions and institutional regulations. For instance, due to homogenous school systems and social infrastructure, similar cultural events, and a uniform standard of living across the German federal states one can expect the implemented policy measures to exhibit a comparatively similar impact on the cross-sectional observations.


The paper is structured as follows. The next section introduces the data on infections and containment measures as well as the methodology. Section 3 presents the result. The last section concludes.

2. Data and Methods

We use as dependent variables in a panel estimation the confirmed COVID-19 infection cases of the Robert Koch Institute by German federal states. Figure 1 shows the daily time series from the 15th of February to the 4th of May for Germany on a logarithmic scale; for the analysis at the federal state level we use available data from the 11th of March onwards. The 16 time series are presented in Figure 2 in the Appendix.

Hartl et al. (2020a,b) find in their trend break analyses a reduction in the growth rate of confirmed cases of infection in Germany by 13 percentage points on the 20th of March and by a further eight percentage points on the 30th of March. Such a flattening is also clearly visible in the data for the federal states. We seek to determine which concrete measures led to this flattening.

Figure 1: Confirmed cases of COVID-19 infections (Robert Koch Institute), logarithmic scale



The time series are modelled with loglinear difference equations, similar to Hartl et al. (2020b). As explanatory variables we use the different containment measures. Specifically, we create indicators that assume the value 0 and 1 as soon as the respective measure comes into effect in the respective state. The regression coefficients then indicate by how many percentage points the daily infection growth rate has changed as a result of the respective measure. The containment measures per se were usually very similar in the different states, so that we focus on differences in the timing.

We look at closures of service sectors in public life, schools and childcare, mass events (over 1.000 participants), the closure of external borders (with Austria, Belgium, Denmark, France, Luxembourg, and Switzerland) and curfews¹. The former cover retail, accommodation, gastronomy, cinemas, trade fairs/events, other education (such as music schools etc.), art/entertainment/recreation and hairdressers/cosmetics. The data by federal state were determined in the course of a comprehensive search and compiled in a data set (Bauer/Weber 2020). Table 1 shows for the measures the average date of entry into force (with one decimal place) as well as the standard deviation in days across the federal states.

¹ Depending on the federal state, for example, contact with more than one person outside one's own household was prohibited in public spaces, or one was only allowed to leave one's own home for a valid reason.

Table 1: Time of occurrence of the containment measures in the federal states

Measure	Entry into force	Standard deviation
	average date [march]	
Mass events	11.2	0.91
Schools/Childcare	16.4	0.73
External borders	16.0	*
Curfews	22.8	1.17
Arts/Entertainment/Recreation	16.9	1.63
Cinemas	17.3	2.35
Other education	17.5	1.03
Trade fairs/Events	17.6	1.97
Retail	18.0	0.82
Accommodation	18.6	2.22
Gastronomy	20.8	1.43
Hairdressers/Cosmetics	22.6	1.37

* The borders to Belgium, Denmark, France, Luxembourg, Austria and Switzerland were closed from the German side on the 16th of March. The indicator applies to the neighbouring federal states, where the main effects would be expected – also in view of the fact that the long-distance travel had already declined significantly during the data period. We also applied a scheme weighting the indicator for all federal states according to the distance of their capital to the next affected border. However, such weighting deteriorated the model fit. The borders with the Netherlands, Poland and the Czech Republic were not closed from the German side. Notwithstanding, it does not lead to relevant differences in the results if the indicator is additionally applied to the neighbouring federal states.

The effects of the measures can only appear in the data of confirmed cases of infection with a delay. This results, amongst others, from the incubation period (Lauer et al. 2020, Linton et al. 2020) and the duration for the test and data collection. We allow the measures to have a lasting effect from q days after entry into force, which can build up over r further days, e.g.,

due to varying incubation periods. We determine the optimal time span within a range for $q + r$ from 5 to 16 days (following information from the Robert Koch Institute). For this purpose, we refer to the Bayesian information criterion that weighs the fit to the data against the model complexity. This leads to $q = 7$ and $r = 8$ additional lags. While this choice is optimal from a statistical perspective, in robustness checks we found no substantial impact of varying the delay parameters within a sensible range. We allow for separate lag structures for the packages of containment measures that became effective within the same week.

In general, it is likely that factors other than the indicators considered influence the infections, such as changed behaviour at the workplace or decreasing mobility. Similarly, changing testing behaviour and capacity may be relevant – whereas a general degree of underreporting would cancel out since we use growth rates. To take into account such effects, we include linear federal-state-specific trends and three autoregressive lags of the endogenous variables. In addition, fixed effects for the federal states are considered. These controls account for any differences between the federal states as well as permanent and temporary developments over time. These terms have the advantage of controlling for a multitude of potentially relevant factors that are unobservable. The modelling therefore does not force the flattening of the curve to be explained by the containment measures.

The resulting panel equation for state i reads

$$y_{i,t} = \mu_{1,i}t + \mu_{2,i} + \mu_{3,i}d_t + \sum_{j=1}^p \alpha_j y_{i,t-j} + (1 + \sum_{k=1}^r \gamma_k L^k) X_{i,t-q} \beta + u_{i,t}, \quad (1)$$

where $y_{i,t}$ holds the growth rate of confirmed corona infection cases for state $i = 1, \dots, 16$ at time t , $\mu_{1,i}$ accounts for a state-specific linear time trend, while α_j are autoregressive coefficients. Note that since our number of observations in the time domain is considerably high and autocorrelation will be rather low, one can expect the Nickell bias not to be an issue here. The row vector $X_{i,t}$ holds the policy measures from Table 1 in its columns, while β is a row vector holding the coefficients. L is the lag operator and γ_k are the weights for the r further lags of the policy measures. d_t holds day-of-the-week effects that typically occur in reported infection data due to institutional reasons. $\mu_{2,i}$ represents the fixed effect and $u_{i,t}$ is a state-specific disturbance term.

3. Results

We estimate equation (1) by feasible generalised least squares correcting for heteroscedasticity and contemporaneous correlation of the residuals and use White standard errors. We report total effects, i.e. $(1 + \sum_{k=1}^r \gamma_k)\beta / (1 - \sum_{j=1}^p \alpha_j)$, which mirror the full impact of the containment measures reached once adjustment is completed.

As shown in Table 1, the category of closures of service sectors is divided into eight variables. Due to the temporal proximity of various overlapping measures, naturally there is high multicollinearity in the data. In such situations, the possibilities of clearly distinguishing all effects can be limited, and small sample variation can lead to mutually offsetting estimates. Particularly, any positive values would have to be offset by other negative effects that could be unduly inflated. In order to avoid these issues, we constrain the sector closure effects to be non-positive. This seems sensible since we can safely assume that the closures do not increase infections.

The results with total effects and standard errors (computed using the delta method) are shown in Table 2.

Table 2: Panel regression of the growth rate of confirmed corona infection cases on containment measures

Measure	Total effect	Standard error
Mass events	-0.067	0.005
Schools/Childcare	-0.082	0.007
Curfews	-0.035	0.008
Retail	0	*
Accommodation	-0.004	0.005
Gastronomy	-0.019	0.006
Other education	0	*
Cinemas	0	*
Trade fairs/Events	0	*
Arts/Entertainment/Recreation	0	*
Hairdressers/Cosmetics	-0.002	0.008
External borders	0.018	0.031

* Coefficients hitting the non-positivity constraint.

According to the estimates, school and childcare closures, the prohibition of mass events and curfews played a major role in containing the spread of the virus. These effects have a high statistical significance (with p -values < 0.01). School and childcare closures reduced the growth rate of confirmed infections by 8.2 percentage points, while stopping mass events stands for 6.7 percentage points. The curfews account for 3.5 percentage points. This impact might be limited by the fact that that public life was already significantly restricted when the curfews were introduced.

By contrast, the closure of most of the service sectors reveals no dampening effect on the growth rate of infections. Exceptions are gastronomy² and potentially accommodation. While the former reach statistical significance, the total effect is lower than for the above-

² A further even finer differentiation into restaurants and bars/clubs would not alter the other estimation results.

mentioned measures.³ For the closures, naturally also the size of the sectors plays a role, e.g. for gastronomy compared to a rather minor subsector such as cinemas. An insignificant effect ($p\text{-value} > 0.1$) can be seen in the case of border closures. Naturally, the situation could have been different, had the borders been closed earlier, e.g. before the winter holidays (returners exerted a special effect on the confirmed infections visible in Figure 1 in early March before the start of our regional sample).

The effect of school and childcare closures can work via avoided infections in these facilities themselves, including chains of infection. By the same token, it could be reinforced by parents staying away from their workplace for childcare reasons. Estimates of the amount of work lost in Germany due to the closures can be found in Wanger/Weber (2020). Childcare and the different types of schools were normally closed on the same day. Therefore, in this regard no further differentiation of the effects is possible within the underlying framework.

The model presented here considering the concrete measures quite accurately mirrors the overall flattening of the infection curve. Thereby, also the federal-state-specific linear trends play a role: on average, they have reduced the growth rate in the model by 0.1 percentage points per day, or 5.2 percentage points over the whole sample. So, while other factors also apparently played a role in the flattening, in sum the estimated effects of the measures considered here are predominant. We also checked still more flexible specifications, e.g. allowing for a trend break on the 20th or 30th of March, the break points found by Hartl et al. (2020a,b). This led to only slightly weaker effects of the containment measures.

The results can be compared to some recent or parallel studies (without claim of completeness). Jüni et al. (2020) also find a major impact of school closures, restricting mass gatherings, and social distancing rules on the spread of the corona virus, where social distancing rules were defined as any measure that prevented small gatherings of ten or fewer individuals. In a further cross-country study, Banholzer et al. (2020) find similar effects for curfews (in their study: gathering bans and lockdowns), but a higher impact of venue closures and border closures (where no closure was specified for Germany) and a smaller impact of school closures. Deb et al. (2020) state that containment measures have been, on

³ In an estimation without the non-positivity constraints, the sum of all sector closure effects is insignificant at the one percent level.

average, very effective, including gathering restrictions and school closures. Ferguson et al. (2020) conclude that while the effectiveness of any one intervention in isolation is likely to be limited, multiple combined interventions would have a substantial impact on transmission. The findings of Dehning et al. (2020) mirror our results in that three intervention packages including the cancellation of large events, the closing of schools and childcare and the contact ban reduced the spreading rate.

4. Conclusion

Based on empirical data, we estimate the effects of different containment measures on COVID-19 infection growth in Germany. The results indicate that the flattening of infections is associated with the prohibition of mass events, school and childcare closures as well as curfews. In addition, further unobserved factors played a role. Regarding the economic shutdown, we find no effects of the closures of most service sectors in public life measurable on an aggregate level. Exceptions are gastronomy and potentially accommodation, but with comparatively limited effects.

The future of the containment strategies is currently subject to debate. Thereby, opening schools and childcare is of immense importance from a social, educational and labour market perspective. In this regard, our results indicate that effective concepts to prevent infection are key. The same is true in view of the relaxation of curfews for the public space. Mass events, for example in sports, must still be viewed critically. According to our results, however, there could be good chances that a controlled opening of service sectors in public life is possible without a considerable increase in the number of infections.

This is to be seen in relation to the drastic downturn of labour markets in many countries. The lockdown effects on the economy and the labour market (e.g. Bauer/Weber 2020, Coibion et al. 2020) show that all possibilities must be explored in order to reconcile adequate control of the spread of the virus with economic activity. Bauer/Weber (2020), for example, find that via more layoffs and fewer new hires, the shutdown measures in Germany caused more than half of the short-run unemployment effect of the crisis.

With regard to the discussed results, it should be noted that all measured effects of empirical investigations apply under the conditions that prevailed during the observation period. This includes the sequence of containment measures. For example, the effect of curfews might have been stronger if they had been introduced before other measures in public life. Or closures of public venues might have shown a larger effect if it had happened before people adjusted their behaviour. Nussbaumer-Streit et al. (2020) indicate that the combination of different measures might increase their effectiveness. Therefore, the effects of any opening steps should be monitored and empirically examined in order to collect further evidence.

References

Baldwin, Richard (2020): COVID, remobilisation and the 'stringency possibility corridor': Creating wealth while protecting health. VoxEU.org, 10 April.

Banholzer, Nicolas; van Weenen, Eva; Kratzwald, Bernhard; Seelinger, Arne; Tschernutter, Daniel; Bottrighi, Pierluigi; Cenedese, Alberto; Salles, Joan Puig; Vach, Werner; Feuerriegel, Stefan (2020): Impact of non-pharmaceutical interventions on documented cases of COVID-19. medRxiv 2020.04.16.20062141.

Bauer, Anja; Weber, Enzo (2020): The Unemployment Impact of the COVID-19 Shutdown Measures in Germany. IAB-Discussion Paper 16.

Coibion, Olivier; Gorodnichenko, Yuriy; Weber, Michael (2020): The cost of the COVID-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. VoxEU.org, 12 May.

Deb, Pragyan; Furceri, Davide; Ostry, Jonathan D.; Tawk, Nour (2020): The effect of containment measures on the COVID-19 pandemic. Covid Economics: Vetted and Real-Time Papers, 19, 53-86.

Dehning, J.; Zierenberg, J.; Spitzner, F. P.; Wibral, M.; Neto, J. P.; Wilczek, M.; Priesemann, V. (2020): Inferring change points in the spread of COVID-19 reveals the effectiveness of interventions. *Science*, American Association for the Advancement of Science.

Ferguson, N. M. et al. (2020), Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand, Imperial College COVID-19 Response Team.

Hartl, Tobias; Wälde, Klaus; Weber, Enzo (2020a): Measuring the impact of the German public shutdown on the spread of COVID-19. *VoxEU.org*, 14 April.

Hartl, Tobias; Wälde, Klaus; Weber, Enzo; (2020b): Measuring the impact of the German public shutdown on the spread of COVID-19. *Covid Economics: Vetted and Real-Time Papers*, 1, 25-32.

Jüni, Peter; Rothenbühler, Martina; Bobos, Pavlos; Thorpe, Kevin E.; da Costa, Bruno R.; Fisman, David N.; Slutsky, Arthur S.; Gesink Dionne (2020): Impact of climate and public health interventions on the COVID-19 pandemic: a prospective cohort study. *CMAJ*.

Lauer, S. A. et al. (2020): The incubation period of Coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: Estimation and application. *Annals of Internal Medicine*.


Linton, N. M. et al. (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.

Nussbaumer-Streit, B.; Mayr, V.; Dobrescu A.I. et al. (2020): Quarantine alone or in combination with other public health measures to control COVID-19: a rapid review. *Cochrane Database Syst Rev*, 4(4):CD013574.

Wanger, Susanne; Weber, Enzo (2020): Wegen der Corona-Krise können viele Beschäftigte nicht zur Arbeit kommen. *IAB-Forum*, 6 May.

Appendix

Figure 2: Confirmed cases of COVID-19 infections (Robert Koch Institute) by federal states, logarithmic scale



Covid Economics 24, 1 June 2020: 205-271