


**CENTRE FOR
ECONOMIC
POLICY
RESEARCH**

CEPR PRESS



COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

ISSUE 12
1 MAY 2020

PRICE GOUGING

Luís Cabral and Lei Xu

SIR MODELS: A SURVEY

Christopher Avery, William Bossert,
Adam Thomas Clark, Glenn Ellison
and Sara Fisher Ellison

**THE BURDEN OF SOCIAL
DISTANCING**

Simon Mongey, Laura Pilossoph
and Alex Weinberg

**HEALTH RISK AND INCOME LOSS
WITHIN HOUSEHOLDS**

Warn N. Lekfuangfu, Suphanit
Piyapromdee, Ponpoje Porapakkarm
and Nada Wasi

**POLITICAL COMMUNICATION IN
BRAZIL**

Lucas Argentieri Mariani,
Jessica Gagete-Miranda and Paula Rettl

INEQUALITY IN PAST PANDEMICS

Davide Furceri, Prakash Loungani,
Jonathan D. Ostry and Pietro Pizzuto

FINTECH ADOPTION

Jonathan Fu and Mrinal Mishra

EFFICIENT TEST ALLOCATION

Christelle Baunez, Mickael Degoulet,
Stéphane Luchini, Patrick A. Pintus
and Miriam Teschl

Covid Economics

Vetted and Real-Time Papers

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

Founder: Beatrice Weder di Mauro, President of CEPR

Editor: Charles Wyplosz, Graduate Institute Geneva and CEPR

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

© CEPR Press, 2020

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

Guido Alfani, Bocconi University and
CEPR

Franklin Allen, Imperial College
Business School and CEPR

Oriana Bandiera, London School of
Economics and CEPR

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

Tito Boeri, Bocconi University and
CEPR

Markus K Brunnermeier, Princeton
University and CEPR

Michael C Burda, Humboldt
Universitaet zu Berlin 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

Barry Eichengreen, University of
California, Berkeley and CEPR

Simon J Evenett, University of St
Gallen 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

Tom Kompas, University of
Melbourne and CEBRA

Per Krusell, Stockholm University
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

Moritz Schularick, University of
Bonn and CEPR

Paul Seabright, Toulouse School of
Economics and CEPR

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

Thierry Verdier, Paris School of
Economics and CEPR

Jan C. van Ours, Erasmus University
Rotterdam and CEPR

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

Ethics

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

Covid Economics

Vetted and Real-Time Papers

Issue 12, 1 May 2020

Contents

Seller reputation and price gouging: Evidence from the COVID-19 pandemic <i>Luís Cabral and Lei Xu</i>	1
Policy implications of models of the spread of coronavirus: Perspectives and opportunities for economists <i>Christopher Avery, William Bossert, Adam Thomas Clark, Glenn Ellison and Sara Fisher Ellison</i>	21
Which workers bear the burden of social distancing policies? <i>Simon Mongey, Laura Pilossoph and Alex Weinberg</i>	69
On Covid-19: New implications of job task requirements and spouse's occupational sorting <i>Warn N. Lekfuangfu, Suphanit Piyapromdee, Ponpoje Porapakkarm and Nada Wasi</i>	87
Words can hurt: How political communication can change the pace of an epidemic <i>Lucas Argentieri Mariani, Jessica Gagete-Miranda and Paula Rettl</i>	104
Will Covid-19 affect inequality? Evidence from past pandemics <i>Davide Furceri, Prakash Loungani, Jonathan D. Ostry and Pietro Pizzuto</i>	138
The global impact of COVID-19 on fintech adoption <i>Jonathan Fu and Mrinal Mishra</i>	158
Sub-national allocation of COVID-19 tests: An efficiency criterion with an application to Italian regions <i>Christelle Baunez, Mickael Degoulet, Stéphane Luchini, Patrick A. Pintus and Miriam Teschl</i>	192

Seller reputation and price gouging: Evidence from the COVID-19 pandemic¹

Luís Cabral² and Lei Xu³

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

We test the theory that seller reputation moderates the effect of demand shocks on a seller's propensity to price gouge. From mid January to mid March 2020, 3M masks were priced 2.72 times higher than Amazon sold them in 2019. However, the difference (in price ratios) between a post-COVID-19 entrant and an established seller is estimated to be about 1.6 at times of maximum scarcity, that is, post-COVID-19 entrants price at approximately twice the level of established sellers. Similar results are obtained for Purell hand sanitizer. We also consider cumulative reviews as a measure of what a seller has to lose from damaging its reputation and, again, obtain similar results. Finally, we explore policy implications of our results.

1 Views presented in this paper are those of the authors, and do not necessarily represent the Bank of Canada's views.

2 Paganelli-Bull Professor of Economics, Stern School of Business, New York University; and CEPR Research Fellow.

3 Senior Economist, Bank of Canada.

1. Introduction

In the wake of unexpected negative supply shocks — or positive demand shocks — prices tend to surge. Many economists — most? — find such price hikes normal and in fact efficient: it's the law of supply and demand at work. By contrast, most of the world accuses sellers of price gouging and clamors for laws protecting buyers from unfair pricing. For example, in the aftermath of Hurricane Sandy, which struck the US in October 2012, New Jersey authorities filed civil suits accusing seven gas stations and one hotel of price gouging. Meanwhile, Libertarian TV personality John Stossel was inviting his readers to “hug a price gouger” today.

The tension between efficiency and fairness has existed for as long as economic relations have existed. It is not our goal to resolve this tension — nor would we have the ability to do so. Instead, we are interested in a third possibility, a third way between the efficiency-only and the fairness-only extremes, namely the possibility that the *seller reputation* may act as a moderator of the “price gouging” forces.

This is by no means a novel idea. For example, Akerlof (1980) addresses the issue of what we might call out-of-equilibrium customs, that is, customs from which there is a privately profitable deviation. If deviating from a custom is profitable and if the cost of deviation is lower when more agents deviate, then one might expect the custom eventually to disappear. Akerlof's (1980) point is that reputation — and the threat of its loss — provides the “binding force” that keeps in place a custom that's hard to keep, in the sense that it's costly to keep.

As an application, Akerlof (1980) considers the puzzle of involuntary unemployment. Suppose there is a social code barring employers from paying less than a certain threshold. Then the outcome may be excess demand in the labor market, unemployment. Clearly, there is a profitable deviation from this equilibrium, namely hiring unemployed workers at a lower wage, but such deviation does not take place due to the prevailing custom. One possible reason is that, by hiring workers at a wage below the prescribed minimum, the employer loses its reputation of being a “fair” employer, whereby existing workers “punish” the employer by not cooperating, for example, in the training of new workers.

Although Akerlof's (1980) analysis does not directly address the issue of price gauging, one might think of fair pricing as a generally accepted principle and custom (except possibly by a number of economists). As in Akerlof (1980), we observe that attitudes toward fair pricing remain remarkably stable even though there is a strong incentive to deviate, especially in times of large demand or supply shocks. One reason why sellers do not deviate is the law, as the Hurricane Sandy example illustrates. A second reason, proposed by Akerlof (1980) as a general principle and studied by us in the context of price gouging, is seller reputation. Similarly to Akerlof (1980), we posit that a seller who violates the fair-pricing custom and price gouges is likely to lose its reputation as a “fair” seller, which in turn may hurt it in the future. Specifically, a natural punishment for a badly behaved seller is subsequent boycott by buyers.

The goal of this paper is to develop this idea in the context of the recent COVID-19 pandemic and the multiple cases of price gouging that followed. We propose a test of Akerlof's (1980) theory in the context of online sales of a specific set of products, namely masks and hand sanitizers. Theory predicts that, during a time of excess demand, sellers with a reputation at stake are less likely to hike their prices than sellers with no reputation

to defend. In the wake of the Hurricane Sandy crisis, Cowen (2017) argued that

The reluctance to raise prices is especially strong for nationally branded stores. A local merchant may not care much if people in Iowa are upset at his prices, but major companies will fear damage to their national reputations. The short-term return from selling the water at a higher price is dwarfed by the risk to their business prospects.

In the example considered by Cowen (2017), high reputation is associated with being a national brand (as opposed to a local brand). We consider two online equivalents of the distinction between sellers who have a lot to lose and sellers who don't. First, we would expect that, all else equal, incumbent sellers — that is, continuing sellers — have more to lose than entrants, where the latter are defined as sellers who only start selling after the demand shock takes place. Second, we would expect that large sellers, which we define as sellers with a greater number of consumer reviews, have more to lose than small sellers.

■ **The 2019-2010 COVID-19 pandemic.** The first known COVID-19 case can be traced back to December 1st, 2019, in Wuhan, Hubei, China. The local health authorities released a public notice announcing the new virus on December 31st. Throughout the next month, the virus spread to other Chinese provinces as well as to other countries. The first known US case was confirmed on January 20, 2020, a 35-year-old who had returned from Wuhan five days earlier. On January 29, the White House Coronavirus Task Force was established, and two days later the Trump administration declared a public health emergency. By March 26 the United States had overtaken China and Italy with the highest number of confirmed cases in the world.

As the number of cases increased and media coverage widened, American consumers rushed to stores and online sites in search of personal protective products such as face masks and hand sanitizer. Reports of price gouging soon appeared in the media. For example, a two liter bottle of Purell sold for as much as \$250.

As a result, on March 25th a coalition of 34 attorneys general sent letters to Amazon, Facebook, Ebay, Walmart, and Craigslist advising them that, “as the American community faces an unprecedented health crisis,” they need to do more to prevent price gouging on their platforms. Two days earlier Amazon had already itself published a statement to the effect that “price gouging has no place in our stores” and further reporting that

Amazon has already removed well over half a million of offers from our stores due to coronavirus-based price gouging. We have suspended more than 3,900 selling accounts in our U.S. store alone for violating our fair pricing policies. We began taking these enforcement actions promptly upon discovering this kind of misconduct, and we've been partnering directly with law enforcement agencies to combat price gougers and hold them accountable.¹

Notwithstanding Amazon's assurance that they operate “dynamic automated systems” that “locate and remove unfairly priced items,” we observe significant price hikes on Amazon from mid January till mid March.² One explanation is that the actual implementation of

1. <https://blog.aboutamazon.com/company-news/price-gouging-has-no-place-in-our-stores>, visited on March 26, 2020.

2. Throughout the paper, we refer to prices of items actually on sale, that is, we exclude prices of stocked-out items.

these mechanisms took place not long before the statement was published. Accordingly, our analysis focuses on the period from the US onset of the COVID-19 pandemic to middle March, when we believe Amazon changed its de facto policy.

■ **Summary of results.** We construct a dataset of prices, measures of seller reputation, and indicators of excess demand. To account for endogeneity issues we use time series of number of COVID-19 cases and number of COVID-19-related deaths as instruments of demand and number of competitors.

We then test that the price coefficient on the excess demand variable is decreasing on seller reputation, as implicitly predicted by Akerlof (1980). The results conform with theory: Higher-reputation firms are less likely to take advantage of shortages of supply. The effect is statistically and economically significant. For example, during our sample period the 2020 price of 3M masks was 2.72 times higher than the average 2019 price. However, the difference (in price ratios) between an entrant and an incumbent is estimated to be about 1.6 at times of maximum scarcity (or about 0.8 at times of average scarcity). In other words, at times of maximum scarcity entrants price-gouge at a level that is approximately twice as large as incumbent sellers.

We consider various measures of seller reputation — that is, estimates of the seller's stakes — and various measures of scarcity.³ We always find economically and statistically significant moderating effects of seller reputation. That said, our analysis also suggests that these reputation effects are insufficient to completely limit the extent of price hikes.

■ **Related literature.** In addition to Akerlof (1980), Kahneman, Knetsch, and Thaler (1986) is very germane to our paper. They study whether community standards of fairness matter when it comes to price setting. Based on telephone surveys, they argue in the affirmative. In particular, buyers would be less willing to repeat purchase from a seller who price gouged. A seller who places a large weight on future sales effectively corresponds to a seller for whom reputation is important. In this sense, our setting matches naturally that of Kahneman, Knetsch, and Thaler (1986). One important difference is that, whereas they use telephone surveys, our source of evidence is actual price setting.

Our paper is also related to some literature on Corporate Social Responsibility (CSR). In particular, Bénabou and Tirole (2010) present a vision of CSR which essentially corresponds to long-term corporate decision making. For example, they argue that

a firm could economize on safety or pollution control; this also increases short-run profits, but creates contingent liabilities down the road: risk of future lawsuits, consumer boycotts and environmental clean-up costs.

Along these lines, we may think of price gouging as a type of (lack of) CSR, where the corresponding dynamic link is subsequent consumer boycott.

■ **Roadmap.** The next section lays out the two basic econometric models we propose as a test of the Akerlof (1980) hypothesis. Next, Section 3 describes our data. Our results are included in Section 4. Section 5 concludes the paper.

3. The above results correspond to our second (of three) measures of scarcity. See Section 2 for details.

2. Model

We wish to test the basic Akerlof (1980) hypothesis that seller reputation moderates the immediate incentive to price gouge. Specifically, we posit that sellers with less to lose in terms of reputation are more likely to price gouge than sellers who have more to lose in terms of reputation. Consistent with this hypothesis, we propose two models, each corresponding to a different way of measuring seller reputation.

■ **Incumbents and entrants.** In the first model, we compare Amazon.com with third-party sellers, and further distinguish between incumbent and entrant sellers. Amazon is a platform where an item can be sold either by Amazon.com or by a third-party seller. Amazon.com is the largest seller on Amazon, accounting for roughly 40% of the total sales as of 2019. Though having no clear reputation measures, Amazon.com is widely known to be more reliable than third-party sellers. Also, for many buyers Amazon.com is effectively a “focal” or “default” seller, a reference point, namely in terms of price.

We separate third-party sellers into two categories: incumbents and entrants. Incumbent sellers of an item are those who were selling that item from January 1 to January 15, 2020 (that is, before the COVID-19 US outbreak). By contrast, entrant sellers are those who started selling the item in question after January 16, 2020. Our hypothesis suggests that continuing sellers — who have more at stake than incoming sellers — are less likely to increase prices, or do so by less. Accordingly, our first baseline model takes the form:

$$\text{PriceR}_{ijt} = \alpha_s + \alpha_j + \beta_A A_i \cdot \text{Scar}_{jt} + \beta_I I_i \cdot \text{Scar}_{jt} + \beta_E E_i \cdot \text{Scar}_{jt} + \text{Controls} + \epsilon \quad (1)$$

where i denotes the seller, j the item on sale, and t the day the item is on sale. α_s and α_j control for seller type and item fixed effects, respectively. The value of PriceR_{ijt} is defined as

$$\text{PriceR}_{ijt} = \text{Price}_{ijt} / \text{Amazon19}_j$$

where Amazon19_j is item j 's average price in 2019 when sold by Amazon.com. In other words, PriceR measures the price ratio between the price at time t and the 2019 price. In a steady state situation, we would expect PriceR to be approximately equal to 1 (except possibly for inflation, which during the period of analysis is minimal, or other exogenous factors). A_i , I_i and E_i are indicator variables equal to 1 if seller i is Amazon.com, an incumbent seller, or an entrant seller, respectively.

The main independent variable is Scar_{jt} , a measure of scarcity of item j in day t . We consider three different scarcity measures, that is, $\text{Scar} \in \{S1, S2, S3\}$. The first measure, $S1$, is a dummy variable that turns one if Amazon is out of stock of that item on that day:

$$S1 \equiv \begin{cases} 1 & \text{if Amazon is out of stock} \\ 0 & \text{otherwise} \end{cases}$$

The idea is that Amazon is generally believed and observed not to price gouge. Moreover, many American buyers use Amazon.com as their primary source. For these reasons, when there is a run on a given product Amazon.com is typically the first target and one of the first sellers to run out of stock. Note that the first measure applies to third-party sellers only, because once Amazon.com runs out of stock, there is no pricing data available.

The second measure of scarcity, S2, is defined as

$$S2 \equiv \text{percentage of incumbent sellers out of stock}_{jt}$$

In other words, S2 measures the percentage of item j 's sellers during the January 1-15 period who are out of stock at day t . If all sellers active at the beginning of 2020 — before the COVID-19 crisis developed — are stocked at time t , then we say there is no scarcity, that is, S2 equals 0. If none of them is stocked, then we say S2 equals 1, the maximum value. The idea underlying this measure is that, when a large supply or demand shock takes place that results in significant excess demand, speculators get in the scene and price gouge. This results in a high seller turnover rate — entry of speculators, exit of stocked-out sellers — which in turn results in a high value of S2.

A third possible measure of scarcity, S3, is the cumulative intensity of Google searches, as measured by Google Trends.

$$S3 \equiv \text{cumulative intensity of Google search for word } j$$

where word j relates to the product in question. To maintain consistency across three measures, our Google Trends measure is normalized to the $[0,1]$ interval.

In order to account for possible endogeneity of prices and scarcity measures, we consider a variety of instruments. Specifically, we consider the logarithm of the number of confirmed COVID-19 cases and deaths due to COVID-19 in three different regions: US, China and the rest of the world (ROW). One concern of using this set of instruments is the potential impact on prices through competition. To address this issue and to control for competition effects, we include the number of sellers in the regression, which is also instrumented using the same set of instruments.

■ **Seller history.** The distinction between entrant and incumbent provides a first approach at measuring what a seller has to lose from squandering its reputation. However, this is clearly not the only approach. Another possibility is the length of a seller's history. Specifically, we propose the following alternative model:

$$\text{PriceR}_{ijt} = \alpha_j + \beta_1 \log(\text{RC}_i) + \beta_2 \text{Scar}_{jt} + \beta_3 \log(\text{RC}_i) \cdot \text{Scar}_{jt} + \text{Controls} + \epsilon \quad (2)$$

where the main independent variable is $\log(\text{RC}_i)$, the logarithm of seller i 's review count. Cabral and Hortaçsu (2010) show that, on eBay, the length of a seller's customer feedback record is a good proxy for the seller's previous sales. In this sense, RC_i can be thought of as a proxy of the seller's past experience. The idea is that, the longer a seller's history and/or a seller's size is, the more the seller has to lose from a reputation breakdown, and thus the less the seller is likely to price gouge. Specifically, Akerlof's (1980) hypothesis implies a negative β_3 coefficient: the greater a seller's reputation, the less likely the seller is to increase price as the result of an increase in scarcity.⁴

3. Data

In this section we describe our data. First, we justify our choice of a sample period. Next we describe the product categories under consideration. Finally, we briefly review some basic descriptive statistics.

4. Note that Amazon.com is excluded in this analysis because the review count is only available for third-party sellers.

■ **Sample period.** As mentioned in the introduction, the first known COVID-19 case in the US can be traced back to January 15, 2020. By mid March, under pressure from both media and the law, Amazon implemented a series of measures to fight price gougers on its platform. Since we want to study the role of market reputations, we restrict to the period from January 15 to March 15. Our analysis is robust to changing the precise beginning and end dates.

■ **Product categories.** We focus on two categories of consumer goods that are heavily affected by the outbreak of COVID-19: face masks and hand sanitizers. There are countless items and brands being sold on the market in each of these categories. We focus on the two best known brands: 3M and Purell. We search for 3M face masks and Purell hand sanitizers on Amazon. To exclude illegitimate items and to calculate price margins, we further select items based on the following criteria: First, we restrict to items sold by Amazon.com for at least seven months during 2019. Second, we also restrict to items with more than 3 product reviews.⁵ Applying these filters, we obtain 31 items, including 14 3M different face masks and 17 different Purell hand sanitizers. Tables 1 and 2 list the various 3M and Purell products included in our sample. As can be seen, the differences across products (within the same category) correspond primarily to differences in package size.

We collect price history data and seller characteristics from Keepa, a website that tracks prices of products sold on Amazon. In addition to price, we obtain information on each seller's cumulative number of consumer reviews.

■ **Descriptive statistics.** Table 3 provides basic descriptive statistics of our data. An immediately striking feature is the enormous range in price ratios. For example, 3M masks price ratios vary from .74 (lower price in 2020 than in 2019) to a whopping 43.88, a range that would be unusual during normal times. Second, we observe considerable variation in seller size, with an average of about 10,000 reviews but a standard deviation of about 50,000. Our measures of scarcity also vary significantly, in fact S1 and S2 vary from 0 to 1 (S3 varies from 0 to 1 by construction).

In order to get a better idea of the evolution of scarcity, Figure 1 plots our three measures of scarcity over the first three months of 2020.⁶ As can be seen, S1 and S2 are highly correlated, with a sharp increase during the last weeks of January. The third measure, by contrast, is less correlated and also more volatile. For example, there is a significant spike in Google searches for masks in end of February, whereas searches for hand sanitizers peaked in mid-March.

4. Results

We finally come to our results. We structure this section as follows. First, we present some preliminary aggregate analysis of the change in prices by type of seller. This preliminary analysis suggests that there are significant differences between low-reputation and high-reputation sellers. We then turn to formal regression analysis, including IV regression (to

5. Other selection criteria are used as a robustness test. Our main results do not change significantly.

6. Our first measure, the dummy variable which turns 1 when Amazon is out of stock, only takes two values. The time series in Figure 1 corresponds to the fraction of products — within a given category — for which Amazon is out of stock.

account for endogeneity of scarcity as well as the number of sellers) and quantify the extent to which reputation acts as a moderator of the effect of scarcity on price. Finally, we consider various robustness tests and extensions of our basic results.

■ **Preliminary summary statistical analysis.** As a preliminary stage in our analysis, Table 4 shows some basic descriptive statistics of our price data. Price ratios are defined with respect to Amazon’s 2019 prices. For this reason, the values in the third are equal to 1 by construction. The remaining columns show that, in 2019, third-party sellers were selling for about 60% higher prices than Amazon. In 2020, Amazon set prices slightly higher than in the previous year (about 3 to 4%). By contrast, third-party sellers’ markup with respect to Amazon increased from about 60% to about 141% (masks) and 72% (hand sanitizer).

To further decompose the price changes, we divide 2020 into two time periods: before and after the US outbreak of COVID-19, which we place at January 16, 2020. Table 5 largely confirms our prior from Akerlof’s (1980) theory: incumbent sellers set a price close to twice the Amazon 2019 price, whereas entrants set a price that was on average between three to four times higher than Amazon’s 2019 price.

Figures 2 and 3 further extend the analysis by plotting price ratios over time by type of seller. We only plot price data when an item is in stock. Since Amazon’s items moved frequently between in and out of stock, we plot Amazon’s data as a scatter plot. As for third-party sellers, we observe that there is a clearly defined date when they become out of stock. Accordingly, we plot a continuous price ratio line until that date. For example, 3M mask incumbent sellers (left panel in Figures 2) run out of stock on March 12, about the same time as Amazon, whereas entrant sellers run out of stock on March 25.

In Figure 2 we distinguish between incumbents and entrants. (Recall that entrants are sellers who begin to sell after the first US COVID-19 case.) Both for masks and for hand sanitizer the data clearly shows that (a) on average third-party sellers charge a higher price than Amazon; and, more important, (b) within third-party sellers, entrants are particularly prone to charge higher prices. For masks, our rough estimate is that entrants charge approximately twice as much as incumbents. For hand sanitizer, the difference varies a lot over time, but generally speaking — and with the exception of early February — entrants charge higher prices than incumbents.

While it makes sense to think of incumbents as having more to lose than entrants, presumably some of the entrants could be very large sellers who, not having a specific history in selling masks or purifiers, have a long history of selling on Amazon. To account for this possibility, we consider a second measure of seller reputation, customer review count. In Figure 3, we distinguish between high-RC sellers (RC greater than 1,000) and low-RC sellers (RC less than 1,000). Similar to the incumbent/entrant split, we see that, for masks, on average low-RC sellers set considerably higher prices than high-RC sellers. Regarding hand sanitizer, while the difference is not as clearcut as in the case of masks, the pattern of higher prices by low-RC sellers persists.

■ **Regression analysis results.** Our systematic regression analysis is guided by a series of choices. First, we consider two possible measures of seller “reputation”, that is, measures of what sellers have to lose from consumer boycott. These measures are (a) whether the seller is an incumbent or an entrant (all else equal, incumbents have more to lose); and (b) a seller’s cumulative number of reviews (all else equal, a seller with a larger number of

reviews — a proxy for seller size — has more to lose).

Second, we consider three measures of scarcity: (a) a dummy variable indicating whether Amazon is stocked out, (b) the fraction of incumbent sellers who are out of stock, and (c) the intensity of Google searches for the relevant term.

Finally, in order to account for possible endogeneity, we consider both OLS and IV regressions, where the instruments correspond to COVID-19-related cases and deaths in the US, China and the rest of the world.

All in all, this corresponds to a large number of regressions. In all of them, the dependent variable is a seller's price for a given product at a given date (measured as a price ratio with respect to Amazon's 2019 price). We organize the regressions as follows. Table 6 explores the contrast between Amazon.com, incumbent and entrant sellers, whereas Table 7 focuses on a seller's size as measured by its prior review count. In each case, we consider six different regression models, corresponding to three scarcity variables times two regression methods (OLS and IV).

Consider first the coefficient estimates displayed in Table 6, where we focus on the distinction between incumbent sellers and entrant sellers. As predicted by theory, the impact of scarcity on prices is greater for third-party sellers than for Amazon. And, more important, within third-party sellers, the effect is greater for entrants than for incumbents. Consider for example the IV regression using the S1 scarcity measure (Amazon is stocked out). When Amazon stocks out, entrants increase price by the equivalent of 100% of 2019's price, whereas the incumbents' increase corresponds to 60%. This is consistent with the idea that entrants have less to lose from ruining their reputation than incumbents have. Consider next the S2 IV regression (fourth model).

An alternative way of measuring the desired effect is to consider a one-standard deviation change in the independent variable. From Table 3, we see that the standard deviation of S1 is given by .47. This implies that a one-standard deviation in S1 is associated with a change in an entrant's price equal to 47% of Amazon's 2019 price, whereas for an incumbent the number is 28%. The standard deviation of S2 is given by .39. We thus estimate that a one-standard deviation increase in S2 leads entrants to increase in price by $2.479 \times 0.39 = .97$, or about 97% of the 2019 price. By contrast, incumbents increase price by $0.914 \times 0.39 = .36$, or 36% of the 2019 price. The corresponding values for S3 are 119% and 76% for entrants and incumbents, respectively.

In sum, we estimate that a one-standard deviation increase in scarcity leads incumbents to increase price by 47 to 119%, whereas entrants increase price by 28 to 76%. Although the numbers do vary a bit according to the scarcity measure we use, we note that they are broadly speaking of the same order of magnitude and, more important, the estimate is economically and statistically greater for entrants than it is for incumbents.

Consider now the coefficient estimates displayed in Table 7, where we focus on a seller's reputation count. Our hypothesis predicts that sellers with higher reputation counts have more to lose and are thus less likely to increase prices. In terms of regression coefficients, we predict a negative coefficient on the interaction of scarcity and review counts.

The regression results confirm this expectation: all estimation coefficients of the interaction review count times scarcity are negative. Moreover, the coefficients are significant from both a statistical and an economic perspective. From Table 3, the mean value of the scarcity measures is given by 0.64, 0.51 and 0.23. Moreover, the standard deviation of $\log(\text{RC})$ is 3.32. Restricting to IV regressions, we see that a one-standard deviation increase in $\log(\text{RC})$

leads to a decrease of 8.3% in the price increase with respect to 2019 when considering the S1 scarcity measure. In other words, reputation — as measured by $\log(\text{RC})$ — acts as a moderator of the effect of scarcity on prices. The corresponding values for S2 and S3 are 9.7 and 15.1%. Again, the fact these estimates fall within a reasonably tight range gives us confidence regarding our measurement of scarcity.

■ **Extensions and robustness tests.** We performed a series of robustness tests in terms of variable definitions and sample period. We found that our results are fairly robust, specifically concerning the sign and statistical significance of the key parameter estimates.

One particularly interesting extension consists in combining our two approaches to measuring seller reputation: the distinction between incumbent and entrant; and the cumulative number of buyer reviews. Specifically, we consider the following model, which basically combines (1) and (2)

$$\begin{aligned} \text{PriceR}_{ijt} = & \alpha_j + \alpha_{type} + \beta_1 \log(\text{RC}_i) + \beta_2 \log(\text{RC}_i) \cdot E_i + \beta_3 \text{Scar}_{jt} + \beta_4 \text{Scar}_{jt} \cdot E_i \\ & + \beta_5 \log(\text{RC}_i) \cdot \text{Scar}_{jt} \cdot E_i + \text{Controls} + \epsilon \end{aligned} \quad (3)$$

The question of interest now is whether β_5 positive or negative, that is, whether the moderating factor of seller history, $\log(\text{RC})$, differs from incumbent and entrant sellers.

Table 8 displays the estimates corresponding to (3). As before, we consider three different measures of scarcity (S1, S2 and S3) and two regression types (OLS and IV). Restricting to the set of IV regressions, we see that, except for the first scarcity variable (Amazon is stocked out), the coefficient estimate is statistically significant. In all three cases, the coefficient is negative, indicating that the moderating factor of $\log(\text{RC})$ is greater for entrants than for incumbents. This suggests that being an incumbent or being a large/old seller are substitute factors in terms of moderating the incentives for price gouging.

5. Conclusion

The COVID-19 pandemic has resulted in emergency declarations in most US states. These declarations have the effect of triggering anti-price gouging statutes. These statutes suffer from several limitations. First, as many economists point out, they may defeat the incentive effect of prices, in particular the incentive to increase supply. Second, they can be rather vague (to the extent that they refer, for example, to terms such as “reasonable” price increases or “necessary” goods and services).

Free-market proponents suggest that seller reputation may substitute for regulation when it comes to preventing price gouging, possibly to the point of rendering anti-price gouging unnecessary. Our evidence suggests that seller reputation does act as a limit on how much sellers price-gouge: Higher-reputation sellers are less likely to take advantage of shortages of supply. The effect is statistically significant. However, our analysis also suggests that these reputation effects are far from sufficient to limit the extent of price gouging.

In their letter to Amazon and other sellers, a group of 34 attorneys general (AG) recommended that the platform

Set policies and enforce restrictions on unconscionable price gouging during emergencies: Online retail platforms should prevent unconscionable price increases from occurring by creating and enforcing strong policies that prevent

sellers from deviating in any significant way from the product's price before an emergency. Such policies should examine historical seller prices, and the price offered by other sellers of the same or similar products, to identify and eliminate price gouging.

Our results suggest one specific way cross-seller price comparison can be made, namely with respect to incumbent sellers. In other words, a possible implementation of the AG recommendation is that, once an emergency period is triggered, new sellers be prevented from selling at a higher price than incumbent sellers. While this would not eliminate price increases, we suggest that it might significantly reduce their extent.

Admittedly, one weakness of this argument is that the price set by incumbent sellers is an equilibrium price, and running a counterfactual without a model is a dangerous exercise. It's possible that, absent the high prices set by entering firms, incumbents would themselves price higher than otherwise. This would defeat the purpose of the proposed rule.

Moreover, there is nothing in the AG argument or in our proposal that shows consumers would be better off by preventing price gouging. As mentioned in the introduction, it is not our purpose to provide a welfare analysis of price gouging or the efforts to prevent it. Such analysis would require considerably more and better data, at the very least data on sales and seller characteristics.

Figure 1

Scarcity measures over the first months of 2020

(S1: % items Amazon is out of stock, S2: % continuing sellers, S3: Google searches)





Figure 2
Price ratio by type of seller
(Amazon data as scatter plot to reveal stockouts)

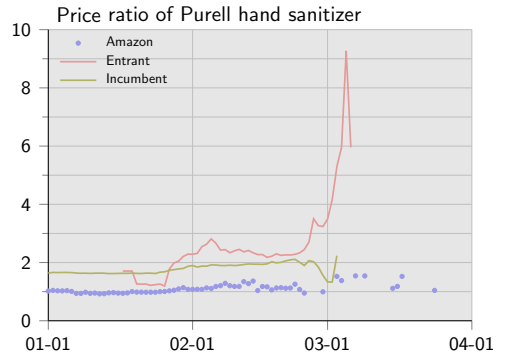





Figure 3
Price ratio by type of seller
(Amazon data as scatter plot to reveal stockouts)



Covid Economics 12, 1 May 2020: 1-20

Table 1

List of 3M products in sample

3M 50051138543438 Particulate Respirator 8511, N95 (Pack of 10)
3M 8210V Particulate Respirator with Cool Flow Valve, Grinding, Sanding, Sawing, Sweeping, Woodworking, Dust, 10/Box
3M 8233PA1-A Lead Paint Removal Valved Respirator
3M 8293 P100 Disposable Particulate Cup Respirator with Cool Flow Exhalation Valve, Standard
3M 8511 Respirator, N95, Cool Flow Valve (10-Pack)
3M 8577CA1-C-PS Chemical Odor Valved Respirator, 2-Pack
3M 8661PC1-A Home Dust Mask, 5-Pack
3M Air Pollution & Pollen Particulate Respirator, N95, Designed for Smoke and Smog Particles, 2 Pack, Adult
3M Particulate Respirator 8210PlusPro, N95, Smoke, Dust, Grinding, Sanding, Sawing, Sweeping, 10/Pack
3M Particulate Respirator 8211, N95
3M Particulate Respirator 8247, R95, with Nuisance Level Organic Vapor Relief, 20/Box
3M Particulate Respirator 8514, N95, with Nuisance Level Organic Vapor Relief
3M Particulate Respirator 8577, P95, with Nuisance Level Organic Vapor Relief
3M Particulate Respirator, 8110S, N95, Smoke, Dust, Grinding, Sanding, Sawing, Sweeping, Smaller Size, 20/Pack

Table 2

List of Purell products in sample

PURELL Advanced Green Certified Hand Sanitizer, Gentle & Free Foam, 535 mL EcoLogo Certified Sanitizer Table Top Pump Bottles (Case of 4)
PURELL Advanced Hand Sanitizer Green Certified Refreshing Gel, Fragrance Free, 12 fl oz Pump Bottle (Pack of 12)
PURELL Advanced Hand Sanitizer Naturals with Plant Based Alcohol, Citrus Scent, 12 fl oz Pump Bottle (Pack of 2)
PURELL Advanced Hand Sanitizer Naturals with Plant Based Alcohol, Citrus scent, 2 fl oz Pump Bottle (Pack of 24)
PURELL Advanced Hand Sanitizer Naturals with Plant Based Alcohol, Citrus scent, 2 fl oz pump bottle (Pack of 6)
PURELL Advanced Hand Sanitizer Naturals with Plant Based Alcohol, Citrus Scent, 28 fl oz Pump Bottle (Pack of 4)
PURELL Advanced Hand Sanitizer Naturals with Plant Based Alcohol, Citrus Scent, 8 fl oz Pump Bottle (Pack of 12)
Purell Advanced Hand Sanitizer Refreshing Gel 8 oz
PURELL Advanced Hand Sanitizer Refreshing Gel, Clean Scent, 2 fl oz Portable Flip Cap Bottle (Pack of 24)
PURELL Advanced Hand Sanitizer Soothing Gel for the workplace, Fresh scent, with Aloe and Vitamin E - 8 fl oz pump bottle (Pack of 4)
PURELL Advanced Hand Sanitizer Variety Pack, Naturals and Refreshing Gel, 1 fl oz Portable Flip Cap Bottle with JELLY WRAP Carrier (Pack of 36)
PURELL Advanced Hand Sanitizer, Refreshing Gel, 2 fl oz Portable, Travel sized Flip Cap Bottles (Pack of 24)
PURELL Advanced Hand Sanitizer, Refreshing Gel, 36 - 1 fl oz Portable, Travel Sized Flip Cap Bottles with Display Bowl
PURELL Advanced Hand Sanitizer, Refreshing Gel, 8 fl oz Sanitizer Table Top Pump Bottles (Pack of 2)
PURELL Healthcare Advanced Hand Sanitizer Foam, Clean Scent, 18 fl oz Pump Bottle (Pack of 4)
PURELL Naturals Advanced Hand Sanitizer Gel, with Skin Conditioners and Essential Oils, 12 fl oz Counter Top Pump Bottle (Case of 12)
PURELL SF607 Advanced Hand Sanitizer Foam, Fragrance Free, 535 mL Sanitizer Counter Top Pump Bottles (Pack of 4)

Table 3
Descriptive statistics

Variable	# obs.	Mean	St. Dev.	Min	Max
Price (3M)	13773	45.51	29.07	6.21	349.99
Price (Purell)	9239	51.45	35.27	2.99	199.99
PriceR (3M)	13773	2.72	1.77	0.74	43.88
PriceR (Purell)	9239	1.82	1.02	0.49	17.50
RC	961	10,600	51,558	0	954,539
log(RC)	961	5.42	3.22	0	13.76
S1	2697	0.64	0.47	0	1
S2	1956	0.51	0.39	0	1
S3	180	0.23	0.24	0	1
<u>Instruments</u>					
Cases (China)	88	44,785	36,282	0	81,999
Deaths (China)	88	1,544	1,382	0	3,299
Cases (US)	88	6,774	21,357	0	121,478
Deaths (US)	88	101	327	0	2,026
Cases (ROW)	88	42,138	95,986	0	457,229
Deaths (ROW)	88	1,996	5,068	0	25,327

Table 4
Average price ratio descriptive statistics

Brand	# Items	2019		2020	
		Amazon	Third-Party	Amazon	Third-Party
3M	14	1.00	1.60	1.04	2.41
Purell	17	1.00	1.57	1.03	1.72

Table 5

Price ratio by seller type

	Before outbreak			After outbreak		
	Mean	Min	Max	Mean	Min	Max
Amazon.com	0.99	0.54	1.6	1.08	0.54	1.92
Third-Party (incumbent sellers)	1.68	0.67	6.03	1.95	0.49	14.71
Third-Party (entrant sellers)	-	-	-	3.69	0.66	43.88

Table 6

Price gouging across different seller types: Amazon vs. incumbents vs. entrants
 Dependent variable: price ratio. p levels: 0.01, 0.05, 0.1

	Scarcity Measure 1		Scarcity Measure 2		Scarcity Measure 3	
	(1)	(2)	(3)	(4)	(5)	(6)
Scarcity \times Amazon			0.550*** (0.11)	0.432*** (0.12)	1.508*** (0.41)	1.468*** (0.42)
Scarcity \times Incumbent	0.541*** (0.02)	0.605*** (0.02)	0.902*** (0.03)	0.914*** (0.03)	3.185*** (0.10)	3.186*** (0.10)
Scarcity \times Entrant	0.780*** (0.05)	1.008*** (0.05)	2.665*** (0.07)	2.479*** (0.07)	4.939*** (0.10)	4.954*** (0.10)
N. of Sellers	0.002*** (0.00)	-0.007*** (0.00)	0.006*** (0.00)	-0.001 (0.00)	0.004*** (0.00)	0.003*** (0.00)
Constant	2.268*** (0.05)	2.384*** (0.05)	2.752*** (0.05)	2.873*** (0.06)	2.833*** (0.05)	2.860*** (0.05)
Item & Seller Type FE	x	x	x	x	x	x
Model	OLS	IV	OLS	IV	OLS	IV
N	22986	22986	22986	22986	22986	22986
R^2	0.560	0.546	0.584	0.577	0.595	0.595

Table 7

Effects of reputation on price gouging

Dependent variable: price ratio. p levels: 0.01, 0.05, 0.1

	Scarcity Measure 1		Scarcity Measure 2		Scarcity Measure 3	
	(1)	(2)	(3)	(4)	(5)	(6)
Review Count	-0.011*** (0.00)	0.004 (0.01)	0.006 (0.00)	0.027*** (0.01)	0.008** (0.00)	0.066*** (0.01)
Scarcity	2.023*** (0.04)	2.724*** (0.07)	2.875*** (0.05)	3.355*** (0.08)	8.235*** (0.12)	11.213*** (0.20)
Review Count \times Scarcity	-0.118*** (0.00)	-0.130*** (0.01)	-0.140*** (0.01)	-0.191*** (0.01)	-0.358*** (0.01)	-0.656*** (0.03)
N. of Sellers	0.003*** (0.00)	-0.014*** (0.00)	0.007*** (0.00)	-0.002** (0.00)	0.006*** (0.00)	0.006*** (0.00)
Constant	2.622*** (0.05)	2.243*** (0.06)	2.975*** (0.05)	2.945*** (0.06)	3.248*** (0.04)	2.722*** (0.06)
Item FE	x	x	x	x	x	x
Model	OLS	IV	OLS	IV	OLS	IV
N	22044	22044	22044	22044	22044	22044
R^2	0.524	0.462	0.564	0.551	0.576	0.565

Table 8

Effects of reputation on price gouging - incumbents vs. entrants

Dependent variable: price ratio. p levels: 0.01, 0.05, 0.1

	Scarcity Measure 1		Scarcity Measure 2		Scarcity Measure 3	
	(1)	(2)	(3)	(4)	(5)	(6)
Review Count	0.001 (0.00)	0.013*** (0.01)	-0.000 (0.00)	0.019*** (0.01)	0.013*** (0.00)	0.044*** (0.01)
Review Count \times Entrant	-0.044*** (0.01)	0.141* (0.08)	0.006 (0.02)	0.311*** (0.04)	0.015 (0.01)	0.212*** (0.02)
Scarcity	1.238*** (0.07)	2.095*** (0.15)	1.791*** (0.10)	2.603*** (0.19)	6.147*** (0.28)	9.189*** (0.66)
Scarcity \times Entrant	-0.261** (0.12)	2.052*** (0.59)	1.360*** (0.17)	3.046*** (0.40)	0.506 (0.33)	3.665*** (0.82)
Review Count \times Scarcity	-0.080*** (0.01)	-0.137*** (0.02)	-0.102*** (0.01)	-0.179*** (0.02)	-0.326*** (0.03)	-0.659*** (0.07)
Review Count \times Scarcity \times Entrant	0.044*** (0.02)	-0.114 (0.09)	0.015 (0.02)	-0.335*** (0.05)	0.030 (0.04)	-0.425*** (0.10)
Num of Sellers	0.002*** (0.00)	-0.008*** (0.00)	0.006*** (0.00)	0.001 (0.00)	0.004*** (0.00)	0.001** (0.00)
Constant	2.724*** (0.05)	1.131** (0.52)	3.144*** (0.05)	0.881*** (0.28)	3.119*** (0.05)	1.815*** (0.15)
Item & Seller Type FE	x	x	x	x	x	x
Model	OLS	IV	OLS	IV	OLS	IV
N	22044	22044	22044	22044	22044	22044
R^2	0.562	0.513	0.584	0.569	0.595	0.571

References

- Akerlof, George A. 1980. "A Theory of Social Custom, of which Unemployment may be One Consequence." *The Quarterly Journal of Economics* 94 (4):749–775.
- Bénabou, Roland and Jean Tirole. 2010. "Individual and Corporate Social Responsibility." *Economica* 77 (305):1–19.
- Cabral, Luís and Ali Hortagsu. 2010. "The Dynamics of Seller Reputation: Evidence from eBay." *Journal of Industrial Economics* 58:54–78.
- Cowen, Tyler. 2017. "Price Gouging Can Be a Type of Hurricane Aid." *Bloomberg Opinion* (September 5).
- Kahneman, Daniel, Jack L Knetsch, and Richard Thaler. 1986. "Fairness as a Constraint on Profit Seeking: Entitlements in the Market." *American Economic Review* 76 (4):728–741.

Policy implications of models of the spread of coronavirus: Perspectives and opportunities for economists¹

Christopher Avery,² William Bossert,³ Adam Thomas Clark,⁴ Glenn Ellison⁵ and Sara Fisher Ellison⁶

Date submitted: 27 April 2020; Date accepted: 28 April 2020

This paper provides a critical review of models of the spread of the coronavirus (SARS-CoV-2) that have been influential in recent policy discussions. It notes potentially important features of the real-world environment that the standard models do not incorporate and discusses reasons why estimating critical parameters is difficult. These limitations may bias forecasts and lead forecasters to overstate confidence in their predictions. They also provide social scientists with opportunities to advance the literature and enable improved policies. This paper also discusses how optimal policies might depend on what is learned from new data and models.

1 We are grateful to Marcy Alsan, Ben Bolker, Amitabh Chandra, Bill Clark, Jonathan Dushoff, Michael Kremer, Balachander Krishnamurthy, Nolan Miller, Ziad Obermeyer, Elizabeth Rourke, Bruce Sacerdote, Doug Staiger, Jim Stock, and Richard Zeckhauser for helpful comments and advice and to Louis Mitchell, Krista Moody, and JingKai Ong for editorial and research assistance

2 Professor, Harvard Kennedy School.

3 Professor, Harvard School of Engineering and Applied Sciences.

4 Postdoctoral Scholar, Helmholtz Centre for Environmental Research.

5 Professor, MIT.

6 Senior Lecturer, MIT.

I. Introduction

Predictive models have taken on newfound importance in response to the spread of the COVID-19 illness, and the SARS-CoV-2 virus that causes it. The fast pace of reports, predictions, and policy changes that have gripped much of the world since March have highlighted both the political relevance and the great uncertainty associated with real-time forecasting in complex systems. On Thursday March 12, Boris Johnson announced that the U.K. would not limit large gatherings or close schools because “The scientific advice is that this could do more harm than good at this time;”⁷ Patrick Vallance, the chief science advisor to the government, described a containment plan based on herd immunity, noting that this would likely require 60% of the population to contract the virus (Robinson and Blanchard (2020)). Four days later, Johnson reversed course and stated “now is the time for everyone to stop non-essential contact with others”⁸, coordinating his announcement with a related briefing on a new study from Imperial College to indicate that the results of that study had influenced the government’s change in policy (Booth (2020)). The United States also changed its guidelines that same day, suggesting that gatherings be limited to ten people (after the CDC had suggested a different limit of 50 people the day before)⁹, with Dr. Deborah Birx, Coronavirus Response Coordinator for the White House Coronavirus Task Force, also referencing the Imperial College report in support of the policy change: “What had the biggest impact in the model is social distancing, small groups, not going in public in large groups” (Fink (2020)). From a modeling perspective, Birx’s

⁷ <https://www.gov.uk/government/speeches/pm-statement-on-coronavirus-12-march-2020>

⁸ <https://www.gov.uk/government/speeches/pm-statement-on-coronavirus-16-march-2020>

⁹ There is some suggestion that the United States previously more relaxed policy had been influenced by a blog post by law professor Richard Epstein with back of the envelope calculations suggesting no more than 500 deaths due to the epidemic in the United States (Dawsey et al., (2020)). We discuss Epstein’s estimate in Section V.

comment is telling because those predictions about the relative impact of different distancing procedures are completely driven by *ad hoc* assumptions about how each affects contacts.

One perplexing element of these references to the Imperial College study was that government officials cited predictions that varied by surprising orders of magnitude. Donald Trump has emphasized the prediction of 500,000 deaths in the UK and 2.2 million deaths in the U.S. from uncontrolled spread of the epidemic, whereas Patrick Vallance focused on a prediction of 20,000 deaths in the UK from an extreme response of “suppression” (Ferguson et al., 2020) which requires substantive changes in behavior until the production of a vaccine (Johnston, 2020). On March 31, Deborah Birx reported that the White House consulted “five or six international or domestic modelers,” and concluded that full mitigation would reduce total deaths in the United States to 100,000 to 200,000. At the same time, Anthony Fauci, Director of the National Institute of Allergy and Infectious Diseases, cast doubt on these predictions: noting that “... I know my modeling colleagues are going to not be happy with me, but models are as good as the assumptions you put into them. And as we get more data, then you put it in and that might change.”¹⁰

At present, after several weeks of “social distancing,” deaths that have been officially linked to SARS-CoV-2 in the UK and US have only amounted to a fraction of those projected in the most pessimistic scenarios for the Imperial College model, but are also already above the predictions of other widely publicized projections. Since it is still an early phase of the pandemic, there should be a role for well-founded models to provide information and guide critical policy decisions. This paper provides a critical review of the literature from several

¹⁰ <https://www.whitehouse.gov/briefings-statements/remarks-president-trump-vice-president-pence-members-coronavirus-task-force-press-briefing-15/>

fields with the goal of providing a foundation for economists and other social scientists to conduct original research that may both save lives and bolster economic outcomes.

Economists have responded energetically and have produced a considerable amount of policy-oriented research related to the spread of the pandemic in a short period of time. Of most relevance to this paper, Atkeson (2020a) and Berger, Herkenhoff, and Mongey (2020) apply and propose extensions to the standard (“SIR”) epidemiological model, while Stock (2020) argues that policy analysis is limited for the moment given the limitations of existing data. We discuss several other recent papers by economists in the text and conclusion.

The paper proceeds as follows. Section 2 outlines the types of models generally used in the ecology/epidemiology literature to study the spread of disease. Section 3 describes a number of heterogeneities that are important in practice but are not incorporated in the baseline versions of these models. Section 4 describes how existing data is frequently used to parameterize existing models; we also include an appendix that lists standard and novel sources of data about the pandemic and hope it will be a useful resource for researchers. Section 5 discusses how this data guides the choice of parameters for the standard dynamic model of the spread of disease. Section 6 provides detail of five estimates of the spread reported to have influenced government policy and critiques the models used to provide those estimates. Section 7 discusses outstanding policy questions and suggests pressing research questions for economists to take on. Section 8 concludes.

II. Types of Models and Their General Properties

There are two primary approaches for modeling the spread of disease: (1) “Mechanistic” and (2) “Phenomenological”.¹¹ The distinction between these models is generally analogous to the distinction between structural and reduced form models in economics. Just as proponents of structural work tout the ability to extend their models to conduct counterfactual analysis, advocates for mechanistic approaches to disease modeling highlight the importance of out of sample predictions: “*All other things being equal, mechanistic models are more powerful since they tell you about the underlying processes driving patterns. They are more likely to work correctly when extrapolating beyond the observed conditions.*” (Bolker (2008), p. 7).

As a biological process, disease spread is usually much better characterized by models that focus on rates of change, rather than models that focus on the states of individual variables themselves. Thus, models usually end up with the same general components, e.g. representing the rates at which new individuals become infected, at which infected individuals recover or die, etc. The primary distinction between phenomenological and mechanistic models in epidemiology, therefore, tends to be more directly related to how models have been parameterized than on the functional forms themselves. Models fit based on a priori biological assumptions, or boots-on-the-ground efforts to identify infected individuals and trace their contacts and resulting infections, tend to be labeled as mechanistic. In this sense, mechanistic models in this literature may be seen as analogous to macroeconomic models that fit some parameters to data and then calibrate other parameters to match external evidence. Models that are parameterized through curve-fitting based on reported case or mortality data, tend to be labeled as phenomenological.

¹¹ See Hurford (2012) for further discussion.

The most common approach that has been used to model the spread of SARS-CoV-2 is the “Susceptible / Infectious / Recovered” (SIR) model.¹² In essence, SIR models can be viewed as continuous-time Markov chain models where only a limited number of transitions between states are possible. In its most basic form, the SIR model considers dynamics within a single homogeneous population (i.e. what physicists would call a “mean-field model”), with form

$$dS/dt = -\beta IS / N$$

$$dI/dt = \beta IS / N - \gamma I$$

$$dR/dt = \gamma I$$

where S , I , and R are the abundance of the three model states, β combines information about encounter rates and infectivity, and γ describes the combined rates of recovery and mortality from the disease.¹³

In this model, the disease will increase in incidence until $\beta IS / N = \gamma I$ (i.e. $S / N = \gamma / \beta$), at which point new infections can no longer keep up with the recovery rate (n.b. this point will always be reached eventually, since recovery results from a linear rate process that is directly proportional to I , whereas infection rates are non-linear and depend jointly on I and S).

Thereafter, the number of new infections will begin to drop, and the disease will die out. This threshold (often referred to as the “*herd immunity*” threshold) describes the minimum size of the susceptible pool required for the infection to spread through a well-mixed population.

Importantly, individuals continue to become infected (especially if the infected pool is large), but

¹² Kermack and McKendrick published seminal articles studying this model from 1927 to 1933. These articles were republished in 1991 (Kermack and McKendrick (1991a), (1991b), (1991c)). Sattenspiel (1990) provides a detailed discussion of the history of models of the spread of infectious disease.

¹³ Chapter 15 of Lehman, Loberg, and Clark (2019) provides a textbook description of the SIR and other models.

on average, the number of new infections per time step will begin to decrease once the herd immunity threshold has been reached.

One potentially confusing aspect of the basic SIR model is that commonly used methods for identifying and analyzing point equilibria are not necessarily indicative of states that will be realized as time goes to infinity. This is because of the "inertia" in the model caused by existing infections. Although the rate at which new individuals become infected drops after the herd immunity threshold is reached (and, recall that this threshold represents the "equilibrium" where $dI/dt = 0$ and $dS/dt = -dR/dt$), the fraction of the population that will ultimately be infected depends on a number of other factors, including the current number of infections, the rate at which new individuals are infected, and the recovery rate. Under some specialized conditions, equilibria in the SIR model can be derived in closed form. (Harko, Lob, and Mak (2014)).

Note that individuals in the SIR model are assumed to be infinitely divisible (i.e. there is no "minimum" step size for transitions among states), and that all state transitions are modeled as exponential processes which can be characterized by their half-lives. Additional compartments can be used to incorporate time-lags into the model. For example, since the incubation period for SARS-CoV-2 is estimated to be about five to six days, the extension of SIR to SEIR to allow for a fourth state ("Exposed") has been especially common for studying SARS-CoV-2 (Hethcote & Driessche (1991) and Li & Muldowney (1995)), although many other variants of the model exist (e.g. "SAIR" models, which include an additional "Asymptomatic" infected state). Collectively, these models are known among epidemiologists as "compartmental models".¹⁴

¹⁴ See Brauer (2008) and in Blackwood & Childs (2018) for general reviews of this term and associated concepts.

The structure of transfer rates and states, and of resulting model predictions can vary enormously even within a specific type of compartmental model. Some factors that differ across SIR models include separation of β into terms describing contacts vs. infection events (Blackwood & Childs (2018)), separation of γ and the “Recovered” bin to distinguish between recoveries and deaths (Gallos & Fefferman (2015)), inclusion of geographic (Bolker (1999)), and demographic (Hethcote (2000)) structure, the application of reproduction and death rates across population bins (Harko et al. (2014)), and loss of immunity (Hethcote & Driessche (1991)). Although these many variants of the model allow it to capture different kinds of biological nuance, it can also inhibit comparisons of predictions across models, e.g. if parameters describe fundamentally different processes or relate to very different types of individuals.

Bolker has compiled a list of more than 40 open source models of the SARS-CoV-2 epidemic.¹⁵ These models are mostly SIR models with considerable overlap in mathematical structure. Some primary differences among them involve: (i) the number and identity of compartments in the model; (ii) whether time-steps are continuous or discrete; (iii) whether dynamics are entirely deterministic or include stochastic aspects; (iv) how any incorporated stochasticity is introduced into the model (e.g. discrete time-steps, Gillespie’s method, tau-leaping)¹⁶; (v) whether methods are employed to quantify or reduce observation error; (vi) the inclusion of demographic and/or spatial structure in the model; and (vii) the types of scenarios and outcomes that they are able to consider (e.g. hospital capacity, number of ICU beds, social distancing strategies, etc.). Jointly, these models and methods provide a summary of the “state of the art” methods currently available for simulating and parameterizing SIR-like models.

¹⁵ https://docs.google.com/spreadsheets/d/1hUZIVDPfa5C8KgURoP_3dAiUQgI6rdb7A5e_g8NcPaY/edit#gid=0

¹⁶ See Wilkinson (2011) and Gillespie’s notes (<https://www.slideshare.net/csgillespie/the-27257701>) for details.

There is a distinct literature that models the spread of disease in a network context. These models typically take bilateral links between people as exogenous primitives and study the probabilistic spread of infection from an initial set of infected people through the rest of the network. It is possible to solve for steady state levels of disease in some versions of a network model with SIR structure, where this steady state level varies with the geometry of the network and the underlying properties of the infection.¹⁷ The epidemiological models that we discuss in Section 5 do not seem to make use of network techniques.

A notable shortcoming of the basic SIR model is that it does not allow for heterogeneity in state frequencies and rate constants.¹⁸ We discuss several different sources of heterogeneity in more detail in Section 2.

The most important and challenging heterogeneity in practice is that individual behavior varies over time. In particular, the spread of disease likely induces individuals to make private decisions to limit contacts with other people. Thus, estimates from scenarios that assume unchecked exponential spread of disease, such as the reported figures from the Imperial College model of 500,000 deaths in the UK and 2.2 million in the United States, do not correspond to the behavioral responses one expects in practice. Further, these gradual increases in “*social-distancing*” that can be expected over the courses of an epidemic change dynamics in a continuous fashion and thus blur the distinctions between mechanistic and phenomenological models.¹⁹ Each type of model can be reasonably well calibrated to an initial period of spread of

¹⁷ See Acemoglu and Ozdaglar (2009) and Easley and Kleinberg (ch. 21, 2010) for class and textbook summaries of the SIR model in a network context.

¹⁸ Murray et al. (2020) specifically argues that a phenomenological approach has an advantage over SEIR models that base transition probabilities on an assumption of “random mixing”.

¹⁹ See Wood (2001) for an example of a “semi-mechanistic” model that incorporates phenomenological components in an otherwise mechanistic model.

disease, but further assumptions, often necessarily ad hoc in nature, are needed to extend either type of model to later phases of an epidemic.

In addition to problems related to heterogeneity, there are three additional challenging aspects of SIR models that make their dynamics especially difficult to predict. First, because the model is nonlinear, small changes in parameter values and initial states can have large effects on dynamics. For example, given the reported doubling-times of SARS-CoV-2 early in the epidemic, comparing two otherwise identical regions that begin with 1 vs. 100 infected individuals can lead to more than a three-week lag-time in case numbers. This challenge is especially pernicious because of the high uncertainty that is usually associated with all of these values, meaning that predictions that fully incorporate uncertainty often span many orders of magnitude. Second, because dynamics in these models tends to be both complex and non-monotonic, classic model diagnostics and fitting tools may not be good indicators of whether a model will produce good extrapolations. For example, many models have high predictive ability when fit to the early stages of an epidemic, where growth in the number of infected individuals is approximately exponential. However, comparatively few models are able to accurately predict the saturation point at which the number of new infections begins to decline or the expected number of infected individuals at the peak of the epidemic. Third, these problems are compounded by the fact that disease transmission involves substantial time lags. Not only can these lags confound models that fail to incorporate them correctly, but they also make it difficult to identify the effects of interventions on disease spread, since changes in observed case numbers lag infection events by at least several days.

Sadly, all three of these challenges seem to be particularly acute for SARS-CoV-2. Uncertainty in rates and states are especially high for the current pandemic, both because

forecasts must be made in real-time, and because there are unprecedented stresses at a global scale on the institutions that are typically responsible for collecting and reporting data. Due to the unusually high asymptomatic rate of infection, it is also difficult to distinguish differences in reported case numbers that are driven by differences in actual infections vs. differences in reporting processes, so nonlinearities and saturation points are relatively difficult to identify. Because SARS-CoV-2 appears to have a particularly long latent stage in comparison to similar diseases, lags between policy interventions and observable changes in dynamics have often taken a week or longer.

III. Heterogeneities and Modeling

Several sources of heterogeneity complicate any modeling approach to SARS-CoV-2; it would be natural to attempt to incorporate one or more of them in future models.

Heterogeneous Exposure

The standard SIR model assumes a single transmission rate for the entire population, but systematic differences in the routine of daily life may result in different patterns of social interactions. Some of these differences are predictable functions of population density, and, thus, it is natural to expect different levels of social contact in urban than in rural areas. But daily routines may also be quite different across major cities: more than half of daily commuters take public transit in New York City, whereas only ten percent do so in Los Angeles.²⁰

Household and building structure may also promote or inhibit the spread of disease. Nursing homes and long-term care facilities may be particularly susceptible because of the heavy degree of interactions between residents and staff that take place on a moment-to-moment basis.

²⁰ https://en.wikipedia.org/wiki/List_of_U.S._cities_with_high_transit_ridership

Some early statistics are consistent with the hypothesis that residents in these facilities are at unusual risk of infection: as of April 14, residents and workers at long-term care facilities accounted for 13.8% of all positive tests (3,907 of 28,163) in Massachusetts.²¹ Similarly, 36% of residents at Pine Street Inn, a homeless shelter in Boston, tested positive for SARS-CoV-2.²²

It also appears that some types of social contacts may be unusually likely to transmit SARS-CoV-2. A single “Beer Pong” party where participants shared drink glasses at an Austrian ski resort is credited with producing hundreds of infections in Denmark, Germany, and Norway (Hruby, 2020). A soccer game with record attendance may help to explain why Bergamo is an epicenter of the pandemic, perhaps exacerbated by the outcome of the game, as fans hugged and kissed each time Atalanta (the team from Bergamo) scored. Atalanta won the game 4 to 1 (Azzoni and Dampf (2020)). A March 10th choir practice in Washington State with 60 attendees resulted in 45 infections and two deaths.²³ The phenomenon of “superspreaders” is well known in the epidemiology literature: as Stein (2011) summarizes, “The minority of individuals who infect disproportionately more susceptible contacts, as compared to most individuals who infect few or no others, became known as super-spreaders, and their existence is deeply rooted in history.” (See Hethcote and Van Ark (1987) for a model that incorporates some of these heterogeneities.)

Early studies of the spread of SARS-CoV-2 highlight the importance of networks and the advantages of flexible jobs that enable some people to work from home (Dingel and Neiman (2020) assess variations across cities and industries in terms of feasibility of working from home.) Geographic concentrations of infections in New York and Italy can be traced to the

²¹ <https://www.mass.gov/doc/covid-19-cases-in-massachusetts-as-of-april-14-2020/>

²² <https://www.wbur.org/commonhealth/2020/04/14/coronavirus-boston-homeless-testing>

²³ <https://www.cnn.com/2020/04/01/us/washington-choir-practice-coronavirus-deaths/index.html>

strength of network ties to early cases (Kuchler et al. (2020), see also Borjas (2020)), while people have taken greatest steps for social distancing in regions with relatively high median incomes and internet access (Chiou and Tucker, 2020). These observations suggest that there may be systematically different trajectories for the spread of disease across incomes and regions, and perhaps even across regions according to political predilections (Allcott et al. (2020)). Case counts to date suggest possible differences for infection rates by race and ethnicity (Kendi (2020)) but testing rates may vary across these groups and data for the race of infected individuals is very spotty.²⁴

Heterogeneous response across people

Observational data suggests strong correlations between age, sex, existing medical conditions, and risk from SARS-CoV-2 infection. On March 26, the Center for Disease Control (CDC) reported that people aged 65 and over accounted for approximately half of hospitalizations and ICU admissions and 80% of deaths and those aged 85 and over were particularly affected.²⁵ The standard SIR model assumes equal mortality rates for all infected people; the SIR models that we discuss typically allow for differential risks by age, but often do not incorporate other factors that appear to be strongly correlated with risk.

Though the United States is not tracking deaths by gender (Gupta (2020)), the observed mortality rate has been at least 50% higher for men than for women in 12 of the 14 countries for which data is available.²⁶ In addition, the CDC warns that “people of any age who have *serious*

²⁴ In Massachusetts, for example, “Hispanic” (22.4%) and “Non-Hispanic Black/African American” (15.7%) groups are disproportionately represented among people with a positive test and a reported category for Race/Ethnicity but Race/Ethnicity is listed as “Unknown” for 45% and as “Missing” for 18% of people with positive tests. (<https://www.mass.gov/doc/covid-19-cases-in-massachusetts-as-of-april-14-2020>)

²⁵ <https://www.cdc.gov/mmwr/volumes/69/wr/mm6912e2.htm>

²⁶ <http://globalhealth5050.org/covid19/>, retrieved April 7, 2020.

underlying medical conditions might be at higher risk for severe illness from COVID-19.”²⁷

These guidelines are consistent with retrospective analysis of multi-center data from Wuhan, China, which found that hypertension, diabetes, and heart disease were the most common comorbidities and that Sequential Organ Failure Assessment (SOFA) and d-dimer scores were significantly associated with death for hospitalized patients (Zhou et al. (2020)). Since the CDC includes “severe obesity” (BMI \geq 40) as an underlying medical condition that increases risk in this case, the high rate of adult obesity may increase the mortality rate from SARS-CoV-2 in the United States (Ludwig and Malley (2020)). Anthony Fauci noted that “comorbidities ... are, unfortunately, disproportionately prevalent in the African American population,” suggesting another reason that there might be variations in outcomes by race and ethnicity.²⁸ It may be possible to account for these variations in risk across cities, countries, and subgroups by post-stratifying the final size of the epidemic by the estimated proportions falling within specific risk groups, but that would require information on the interactions between underlying medical conditions and mortality rates from SARS-CoV-2 that is not yet available.

Medical Capacity by Location

The standard SIR model also assumes consistent medical response independent of one’s geographic location. Yet, there may be important variation in medical capacity across locations. The United States ranks 32nd in the world with a bit less than 3 hospital beds per thousand people. By comparison, South Korea, Japan, and Germany each have at least 8 hospital beds per thousand people, though the United States has greater per-capita ICU capacity than any of those

²⁷ <https://www.cdc.gov/coronavirus/2019-ncov/need-extra-precautions/groups-at-higher-risk.html>

²⁸ <https://www.whitehouse.gov/briefings-statements/remarks-president-trump-vice-president-pence-members-coronavirus-task-force-press-briefing-april-7-2020/>

countries.²⁹ Hospital capacity may be of most importance at moments of peak local infection. To avoid the strain of going overcapacity in some places, France (Chang (2020)) and Germany (Ellyat (2020)) have been moving patients to under-utilized hospitals. By contrast, in the United States, so far there has only been movement of medical personnel to New York City and other hot spots of disease.

Dosage and the Nature of Exposure

The standard SIR model does not adjust the risk of infection or mortality rates conditional on infection to account for the dose of exposure, or more generally, for the nature of social contact. Some researchers have conjectured that exposure to a higher “viral load” can result in more severe illness, perhaps motivated in part by the death of Li Wenliang, the 33 year old ophthalmologist who worked to publicize the disease in December, 2019. As American doctors Rabinowitz and Bartman comment, “Dose sensitivity has been observed for every common acute viral infection that has been studied in lab animals, including coronaviruses”. (Rabinowitz and Bartman (2020)). Early evidence from China and Italy only finds a strong association between viral load and the likelihood of infecting others, not an association between viral load and the outcome for the person with that viral load. (See Cereda et al. (2020), He et al. (2020), Geddes (2020), and Heneghan, Brassey, and Jefferson (2020) for discussion of this point.)

Multiple strains of virus

At present, eight strains of SARS-CoV-2 have been identified by researchers, though it is not clear how many other strains might exist. These original eight strains are sufficiently similar that it seems that the virus is not mutating quickly enough to produce a version that is more

²⁹ https://en.wikipedia.org/wiki/List_of_countries_by_hospital_beds

deadly to humans (Weise (2020)). While anti-viral immunity can be long lasting, it is not certain that a person who has been exposed to one strain of SARS-CoV-2 will be immune to exposure to other strains. Contrary to some statements in the popular media, it remains unclear whether less virulent (i.e. less deadly) SARS-CoV-2 strains will emerge on time scales that are relevant for this epidemic. First, although these sorts of changes are well-supported by theory³⁰, in practice they seem to occur very slowly, especially for diseases with relatively low mortality rates. Moreover, the same geographic heterogeneity that may work to slow the spread of the virus will almost certainly slow the spread of less virulent strains, even if they have a considerable fitness advantage over other strains.

IV. Data

The World Health Organization (WHO) defines a pandemic as “the worldwide spread of a new disease”.³¹ Data collection in the early stages of a pandemic is both especially important and especially challenging because a pandemic typically involves a relatively novel disease, and by definition involves a globally occurring stress on public institutions which can make information gathering and dissemination difficult.

Early aggregate statistics for infections and hospitalizations in an epidemic are notoriously unreliable. Many regions will be slow to identify and test the first people who contract a new disease, and those who are infected might not be inclined to go to a hospital for treatment and might not be admitted given the seemingly common nature of their symptoms. By contrast, death rates have been seen as more reliable measures for tracking the initial spread of diseases such as Ebola and SARS.

³⁰ See Read (1994) and Lively (1999) for discussions of evolutionary models of virulence.

³¹ https://www.who.int/csr/disease/swineflu/frequently_asked_questions/pandemic/en/

Several factors limit the value of death rates for short-run analysis in modeling the spread of SARS-CoV-2. Deaths seem to lag exposures by at least two weeks, so any analysis based on death rates requires a substantial time lag and reduces one's ability to assess the effect of a policy change. There may also be an additional time lag before deaths are incorporated in administrative data; official revisions have resulted in upward adjustments of previously reported deaths by as much as 78% in the United Kingdom (Giles (2020)). Further, there is often ambiguity in the cause of death from SARS-CoV-2 because it interacts so frequently with an underlying condition and because pneumonia is frequently the proximate cause of death. There may even be systemic differences across countries - for instance, doctors in the United States have much more discretion in the choice of cause of death than doctors in the UK (Henriques, 2020). The difficulty of assessing death rates from historical data is underscored by the ongoing debate about the effects of the 1918 Flu,³² and similar uncertainty exists even for more recent widespread influenza outbreaks.

An alternative method for assessing death rates is to use year-over-year difference-in-difference comparisons of all deaths at a regional level, comparing areas with known infections to other areas, though that approach is often not possible in real time. Using this method, there is growing evidence that the number of deaths attributed to SARS-CoV-2 is a substantial underestimate of the actual effect of the disease.³³ The difference-in-difference method also likely includes the effect of negative externalities of the epidemic, such as the loss of resources for doctors to treat other medical conditions, in the death rates ascribed to SARS-CoV-2. While

³² For example, Johnson and Mueller (2002) and Spreuwenberg, Kroneman, and Paget (2018) provide dramatically different estimates of these death rates.

³³ <https://www.economist.com/graphic-detail/2020/04/03/covid-19s-death-toll-appears-higher-than-official-figures-suggest>

it is essential to account for negative externalities in cost-benefit analyses, they likely distort estimates of the spread of disease that are based on death rates in overwhelmed areas.

There is now an abundance of data at a local and daily level on test results and hospitalizations. Unfortunately, these data are incomplete and inconsistent across cities and countries because testing is not universal and because there is no universal standard for hospitalization. Asymptomatic individuals may make up a substantial share of the infected, and are now believed to be contagious, but are infrequently tested. (See Stock, Aspelund, Droste, and Walker (2020) for estimates from Iceland, which has tested 6% of its citizens). Since we don't know how the tested population is related to the general population in almost any location, it is not clear what we learn from observing aggregate test results by region or country.

Sero-surveys

Properties of the standard “polymerase chain reaction” (PCR) test for SARS-CoV-2 are not known with certainty. The proportion of false negative tests may vary by site (Cummins, 2020) and may be as large as 30%. Given the absence of reliable test information for representative samples of the population, Dushoff (2020) and others suggest that we should try to conduct *retrospective* blood-based tests known as sero-surveys to determine the population infection rate. Sero-surveys test for the existence of antibodies in the blood as evidence of past infection. The Johns Hopkins Center for Health Security estimates that the sensitivity—the fraction of infected people who test positive—is 93.8% and the specificity—the fraction of non-infected who test negative—is 95.6% for the RDT serology test, the first test approved for use in the U.S.³⁴ The

³⁴ <http://www.centerforhealthsecurity.org/resources/COVID-19/Serology-based-tests-for-COVID-19.html>

CDC plans to conduct national sero-surveys, but will not start them until summer 2020 (Branswell (2020)). One initial study provides serology results for a sample of 3,330 people in Santa Clara County, California (Bendavid et al. (2020)), but the interpretation of those results are complicated by uncertainty about the specificity of the test that was used (Gelman (2020)).

We provide a more comprehensive list of data sources that might be of interest to researchers in the Appendix.

V. Estimating Parameter Values for the SIR Model

The basic reproduction number R_0 describes the expected number of new infections that will be produced by a single infected individual in a “naive” population (i.e. where all other individuals are susceptible). It is a very helpful statistic for comparing different models and tends to be one of the rate parameters that is easiest to accurately measure. In general, there are two methods that can be used to estimate R_0 .

The first, more mechanistic, approach is to use contact-tracing data from early in an epidemic. Health organizations often deploy contact tracing to attempt to slow or eradicate the spread of disease – identifying and quarantining all people who come in contact with the earliest-known patients. Because very few individuals are infected, and because institutions are comparatively unstressed, it is ideally possible to identify and test all individuals who came into contact with an infected individual, and from that to accurately compute R_0 (Ferretti et al., (2020)). However, this approach is less effective as the number of infected individuals grows large, as it can be difficult to definitively identify the source of new infections, and because comprehensive testing of contacts may no longer be feasible. Additionally, this method of estimation may fail to accurately capture heterogeneity in infectiousness among individuals or regions, especially if the distribution has a long tail (e.g. due to “super-spreaders”). The second,

more phenomenological, approach is often straightforward to implement for the early stages of an epidemic, e.g. fitting the rate of growth to a known functional form. However, as case numbers grow, this approach becomes less effective, and because it is parametrized based on the number of confirmed cases, biased or temporally inconsistent testing will reduce the reliability of this estimate. Nevertheless, there has been high similarity in phenomenologically estimated R_0 values across different countries, suggesting that there is some underlying biological reality to these estimates.³⁵

Even under ideal circumstances, it is important to note that parameterizing an SIR model with reported case data is necessarily a tricky endeavor. A classic method used in many textbook exercises would be to fit the basic model we introduce at the beginning of this text by setting the recovery rate $\gamma = 1/(\text{mean infectious period})$ and $R_0 = \beta/\gamma$. This approach is especially well-suited for some types of models, e.g. individual-based stochastic simulations, as it treats the rate terms as waiting times drawn from an exponential distribution. Thus, if one were to simulate an initial cohort of infected individuals within an arbitrarily large susceptible population, and track them until all of the original infected individuals recovered, one would find that the average duration of infection across those original individuals was $1/\gamma$, and that the total number of new infections directly caused by the original individuals, divided by the total number of original individuals, would be R_0 . However, this definition is somewhat problematic in practice, as it requires the ability to fully track outcomes of all individuals in the initial cohort, until the entire cohort has recovered. If any individuals are omitted – e.g. because they had especially mild cases and recovered before they could be identified or infect others, or because they had especially severe

³⁵ Ridenhour, Kowalik, and Shay (2014) estimate somewhat different values of R_0 across countries for the 2009 flu epidemic.

cases and did not recover before the end of the study, the resulting rate coefficients will be biased. One partial solution is to fit rates based on “typical” rather than “average” outcomes. For example, if we know that 50% of patients recover within $t_{1/2}$ days, then we can estimate γ from the half-life of the infection – i.e. $\gamma \approx \log(1/2)/t_{1/2}$. Similarly, we can roughly approximate R_0 in terms of the average number of new infections per half-life of the infection – i.e. $R_0 + 1 \approx \exp(\beta t_{1/2})$, and thus $\beta \approx \log(R_0 + 1)/t_{1/2}$, where “ $R_0 + 1$ ” signifies R_0 new infections over $t_{1/2}$ days, plus the original infection. The advantage of this approach is that it can be applied even when only partial data from a cohort is available, and does not require information about individual-level outcomes, which can be difficult to acquire for legal and logistical reasons. We stress, however, that in modern applications of disease models, parameterization approaches are generally much more sophisticated than those that we present here – see Massaud *et al.* (2010) for examples of some suitable methods.

Economists might consider other approaches for estimating R_0 or might develop models for which this parameter is not fundamental, as the similarity in existing estimates of R_0 across countries can be viewed as suspicious. A structural economic model would provide microfoundations so that R_0 emerges endogenously in equilibrium from a model of social interactions. Further, economists might focus on a more specific parameter, such as the probability of transmission per interaction (or possibly per unit of time of interactions) to facilitate analysis of policies that limit or penalize interactions. An alternative reduced form model might estimate R_0 as a function of city characteristics, policies adopted, and the degree to which people self-modify behavior as a function of prevalence.

While R_0 for SARS-CoV-2 can be estimated with comparatively high accuracy relative to other rate coefficients, there remains considerable uncertainty even about the (average) value of

this parameter. For example, a collation of studies by the Germany's Robert Koch Institute³⁶ suggests that in the absence of control measures, R_0 falls between about 2.4 and 3.3 (although they acknowledge excluding several studies with especially high estimates). If we assume that half of infected individuals recover within 5-6 days, these rates would imply a doubling-time of about 2.5-3.5 days, which closely matches initial growth rates of confirmed case numbers in most countries (though again, all of these estimates are highly uncertain). Although this is a deceptively small range, it is important to remember that small changes in rates have major consequences. After a month of uncontrolled growth starting from a single infection, a doubling time of 2.5 days would yield about 4000 infections, whereas a doubling time of 3.5 would lead to about 400. Note, however, that much of this uncertainty is ultimately related to the generation time of the infection, rather than the rate at which it spreads in a population. Thus, phenomenological estimates of overall growth rates in confirmed case numbers have generally been less variable than estimates of R_0 , and, correspondingly, tend to provide more accurate short-term forecasts of case numbers.

There is minimal data for estimating parameters other than R_0 , and in many cases uncertainty spans a wide range of possible values. Using the summary statistics of the Robert Koch Institute, as reported on 7 April 2020, as an example: incubation time has been reported to be about 5-6 days, but with observed values ranging from 1 to 14 days; reported case mortality rates (i.e. deaths divided by confirmed cases) range from 0.1% to 22%; the fraction of confirmed cases relative to actual cases (i.e. including undetected cases) has been roughly estimated at 5-9.2% based on early reports from China; the duration of contagiousness is so uncertain that no

³⁶ See https://www.rki.de/DE/Content/InfAZ/N/Neuartiges_Coronavirus/Steckbrief.html for summary statistics and related citations (in German):

range is provided. Because each of these ranges of values contributes additional uncertainty to model forecasts, it seems likely that predictions which fully incorporate all of these sources of uncertainty will likewise span at least an order of magnitude, and likely much more.

Estimates for the fraction of infections that have been detected include at least two sources of uncertainty: first, that in most locations, not all symptomatic people can be tested, and second, not all people infected with SARS-CoV-2 are symptomatic (or at least, their infections are “sub-clinical”, in that they are not sick enough to seek medical care). Stock (2020) defines the “asymptomatic rate” as the fraction of the infected who are not recorded as infected. Early estimates of asymptomatic rates ranged from 18% to 31% (Stock, 2020), but more recent data suggests values in a much higher range, with Stock et al. (2020) estimating a value of 90% based on more representative sampling in Iceland. Stock (2020) also observes that different reports of the asymptomatic rate may not be comparable because they do not also use the same definition of “asymptomatic”. Uncertainty about the asymptomatic rate is closely linked with uncertainty about the mortality rate for SARS-CoV-2; Atkeson (2020b) shows that existing data are consistent with a wide range of mortality rates.

VI. Existing Models and Predictions

A handful of models that have received attention in the media and yield a fairly disparate set of results; we review five of them here.³⁷

Epstein Estimates

Richard Epstein predicted that the “adaptive response” and other factors would dramatically limit the spread of disease: “Even if there is some undercounting, it is highly unlikely, given the relatively short (two-week) incubation period, that the number of current cases will more than double or triple.” (Epstein (2020)) This model is largely mechanistic, in that it hypothesizes specific processes that drive dynamics, although the specific rates chosen for the model have been heavily criticized. Epstein initially predicted 500 deaths or fewer in the United States, then later revised that estimate to 5,000 deaths. Needless to say, these predictions have already been falsified.

Oxford Study

Lourenco et al (2020) consider three scenarios with a range of point estimates from 36% to 68% of UK adults infected as of March 19. The twin implications of these estimates was that the country was already likely close to the herd immunity point and – dividing observed deaths by the implied number of infected people – that the mortality rate for the disease must be quite low and thus that the Imperial College projections were dramatic overestimates. But the paper essentially assumes these conclusions: “Our overall approach rests on the assumption that only

³⁷ We chose to review the Epstein, Ferguson, and Murray models because they have been reported to be used by governments, the Lourenco model because it gained attention for providing a prediction contrary to that of Ferguson, and the Kissler model because one of the co-authors of that paper, Marc Lipsitch, has been publicly identified as one of the modelers consulted by the U.S. government. Rogers and Molteni (2020) provide a very detailed popular account of the relationship between policy choices and the projections of various disease models.

a very small proportion of the population is at risk of hospitalisable illness.” (Their MCMC estimation starts from a prior that the case fatality rate is about 1/700 in some versions and about 1/7000 in others.) This model is largely mechanistic, and follows a relatively standard SIR form.

IHME

The IHME paper (Murray et al. (2020)) is designed to estimate the demand for hospital beds, and thus to guide local-level policy decisions. At the same time, it yields estimates for caseload and deaths, and those estimates have been widely publicized and used by state governments for planning purposes.

This model takes as data the observed death rates on a day-by-day basis for each U.S. state and fits a cumulative Gaussian cdf to produce the predicted future course of the log of death rates in each state after accounting for state-level variations in age distribution. It then works backwards from future predicted death rates to project the demand for hospital beds and ICU admissions per capita for each state through the course of the epidemic.³⁸

These results rely on two critical modeling assumptions. First, the authors assume a single change in behavior (and thus the nature of the fitted curve) at the time of the formal announcement of social distancing measures in each state. Second, the authors assume that social distancing induces changes in death rates that are analogous to the rises and falls indicated by the official statistics for Wuhan. As the paper comments, “*Modeling for US states based on one completed epidemic, at least for the first wave, and many incomplete epidemics is intrinsically challenging. The consequent main limitation of our study is that observed epidemic*

³⁸ The general method of fitting a deterministic curve to a cumulative set of observed cases tends to produce biased results because it does not adequately account for uncertainty. King et al. (2015) use simulation results to demonstrate this bias for predictions of the spread of Ebola.

curves for COVID-19 deaths define the likely trajectory for US states.” Over time, IHME has been using Bayesian methods to incorporate data from the U.S. as it is released, while still putting heavy weight on the entire cycle of results observed in Wuhan.

The lockdown in Wuhan was more strict and more strictly enforced than any in the US—even after the recent reopening there are limits on how long residents can spend out of their residential compounds and checkpoints at which authorities take residents’ temperatures and check a phone app that indicates if they are considered at risk of being infected. Accordingly, it seems likely that US infection rates will not decline as quickly as in Wuhan given our less stringent practices of social distancing. For this reason, the IMHE model may underestimate the fraction of deaths which will occur after the peak.

The IHME model is probably the most phenomenological of all of the models that we consider here, in that its predictions are primarily based on fitting curves to historical data, rather than on *a priori* biological assumptions about disease dynamics. This property has allowed the model to provide real-time location-specific forecasts. However, it has also led to intense criticism based on the suggestion that its predictions are more strongly driven by statistical and functional form assumptions than they are indicative of actual likely trends.

Kissler Seasonality Model

Kissler et al., (2020) use historical data from five years of US hospital admissions for two strains of coronavirus (HCoV-HKU1 and HCoV-OC43, which typically induce common colds) to fit parameters for a mechanistic SIR model that vary on a seasonal basis and allows for periods of immunity and cross-immunity after recovery from a given infection. They then expand the model to include SARS-CoV-2 as a third coronavirus under the assumption that it has

similar seasonality to the two milder strains, which have been shown to spread in any season, though with different peaks by season. Focusing on scenarios where policies do not prevent the herd immunity threshold from being reached, they note that substantial uncertainty about the strength and duration of immunity following an infection and levels of cross-immunity makes many future patterns possible. SARS-CoV-2 outbreaks might recur in regular annual or biennial intervals, or the virus could seemingly be mostly eliminated and then resurge five years later.

Ferguson / Imperial College Model

Ferguson et al. (2020) is likely the closest of the five we examine to a structural economics model. It uses Census data to model the spread of disease in a mechanistic SIR model at a fairly local level, considering four sources of social interactions that could produce “transmission events”: within household, at work, in school, or in the community at large. The model is parameterized so that initial transmission events occur approximately equally at home, at school/work, or in the community, matching a stylized fact from previous studies suggesting that approximately one-third of transmissions occur in each of these places.

The model accounts for social distancing by applying particular rules for the effect of five different interventions: “Case Isolation”, “Voluntary Home Quarantine”, “Social Distancing of those over age 70”, “Social Distancing of entire population”, “Closure of schools and universities”. For example, it assumes that closing schools increases contact rates within affected families by 50% and also increases contact rates by 25% in the community in general. Any combination of these five interventions is assumed in the model to induce specific changes in social contacts, and in turn determines the transmission rate R_t in one’s local area at time t .

The changes in contact rates assumed in this model are never justified and, in fact, appear to be entirely arbitrary and in some cases clearly inaccurate.³⁹ Economists would likely prefer the transparency of Atkeson (2020a), which characterizes the behavioral and institutional response to the spread of infection in reduced form with a change in the value of R_0 .

Comparison of the Five Models

One critical distinction between the predictions of IMHE and Ferguson is that IMHE predicts a much earlier peak and steeper decline in the rate of infection, presumably because IMHE puts so much weight on the experience of Wuhan, where an extreme form of social distancing was enforced.⁴⁰ Much as this modeling choice for IMHE seems to have a huge influence on the qualitative nature of the results, the contrarian predictions of Epstein and Lourenco appear to follow directly from their assumptions. The conclusion of Kissler et al. (2020) that SARS-CoV-2 could have peaks in any season appears to follow similarly from their modeling choice to use patterns from other coronaviruses to model the spread of SARS-CoV-2.

Confidence Intervals and Scenarios

Stock (2020) argues that there is too little information to calibrate SEIR models of SARS-CoV-2 at this time, and in particular, that due to selection into testing, we know neither the true number of people who are infected in any country, nor the mortality rate conditional on infection.⁴¹ Stock then demonstrates that two distinct plausible values of the asymptomatic rate (the proportion of infected people who have no symptoms if illness) suggested by data from different subsamples in Wuhan yield dramatically different predictions for the effects of three

³⁹ The scenario for the closing of schools and universities assumes that 25% of universities remain open, though in the United States, it seems that all but Liberty University will be closed from mid-March through at least May 1.

⁴⁰ Bergstrom suggests that this feature of the model has led to systemic errors in its predictions for Italy and Spain. https://twitter.com/CT_Bergstrom/status/1250304069119275009.

⁴¹ See also Bendavid and Bhattacharya (2020).

social distancing polices. These simple simulations cast profound doubt on the precision of any existing model that provides predictions of the future course of the spread of SARS-CoV-2.

Unfortunately, there seems to be a general tendency for researchers to report a greater degree of confidence than is warranted for an existing model, in part because it is not straightforward to quantify parameter uncertainty or to trace the effect of those uncertainties in a non-linear model. Realistic confidence intervals in this context would also be so wide as to seem vacuous. As Jonathan Dushoff commented in an e-mail communication, *“Retrospectively, the most successful looking model is also likely to be a model with narrow confidence intervals, where there was some luck involved in making the model forecasts look good. In cases where there's a lot of models, and not so many realizations, this is very likely to happen.”*

Among the five models we discuss in the previous section, only Murray's (2020) IMHE model provides explicit confidence intervals for any predicted results. For example, that paper reports a 95% confidence interval of (38,242, 162,106) for total deaths in the United States. While this is a reasonably wide range, it is not clear from the paper how this interval was computed, and it still seems likely to be too narrow.⁴² In particular, the paper seems to make no attempt to allow for uncertainty over the degree of reduction in contacts in the United States from mid-March on: presumably this confidence interval is the range of possible values

⁴² The paper explains that uncertainty in the forecast is driven primarily by variance in the "fixed" and "random" effects in the model. It does not provide further details of the computation of any confidence intervals. In a standard hierarchical modelling framework, this version of forecasting errors would imply that their estimates include both variability attributable to differences in mortality rates (of reported deaths) among localities, modelled as "random" parameters drawn from a hyperprior (typically a Gaussian distribution), and uncertainty in the mean mortality rate computed across all localities, modelled as a "fixed" effect. One caveat for such an approach is that this estimate of uncertainty is predicted on two assumption: (1) the data used to fit the model represents an unbiased sample of the true underlying process; and (2) that the true process generating dynamics matches the analytical function hypothesized in their model. See also Bergstrom's critique of this approach.

https://twitter.com/CT_Bergstrom/status/1243838213086539776

conditional on the assumed functional form for deaths, which maintains that deaths rise and fall symmetrically about some peak. The functional form is said to have been chosen because others “did not fit the data as well.” This is implicitly relying heavily on reported data from the two epidemics that have risen and fallen, Wuhan and South Korea, and this dependence is presumably not taken into account.

The four other papers report mean estimates for a range of scenarios, suggesting a confidence interval interpretation without using that language. For example, Lourenco (2020) states that the results from analysis of various scenarios of the Oxford Model indicate “*significant population level immunity accruing by mid March in the UK as ρ is decreased to plausible values*”, where ρ is the proportion of the population susceptible to hospitalisable illness. Kissler et al (2020) report five qualitative implications of their analysis, justifying them with statements such as “*In all modeled scenarios*”, “*many scenarios lead to*”. Similarly, Ferguson et al (2020) states, “*Such policies are robust to uncertainty in both the reproduction number, R_0 (Table 4) and in the severity of the virus (not shown)*,” but allowing for a small degree of variation in a single parameter is hardly a reasonable test of robustness.⁴³

In sum, the language of these papers suggests a degree of certainty that is simply not justified. Even if the parameter values are representative of a wide range of cases within the context of the given model, none of these authors attempts to quantify uncertainty about the validity of their broader modeling choices.

⁴³ The paper only provides the results for variations in the value of R_0 and only appears to consider variation in this parameter in a range from 2.0 to 2.6, which appears to be a much narrower range than a 95% confidence interval for R_0 .

VII. “Flattening the Curve” and Future Policy Considerations

The policy response to SARS-CoV-2 in the US (and several other nations) has been a tragic failure. By the end of this month the US will have lost tens of thousands of lives. The two trillion dollar relief package will ameliorate only a portion of economic suffering that has resulted from the outbreak. And the most rosy realistic forecast for the future is that after another month of social distancing we might be roughly back to a state of the world that is about as promising as where we could have been in mid-March given better policy. Even in that scenario, the number of infections will likely be much higher in mid-May than they were in mid-March. The infectious will be more dispersed. Production of personal protective equipment will have increased, but this will likely be offset by depletion of the existing stockpile. At present, testing capacity in the United States is dramatically higher than it was, but still lags what South Korea achieved.

With this sobering realization in mind, we still need to address the question of policy going forward. A starting point for thinking about optimal policies is to work backwards from the end. In the long run we will presumably reach one of three endgames.

First, we may work to hold infections to a moderate level until an effective vaccine becomes available, perhaps in 12 to 24 months. An optimal policy reaching this endpoint would make tradeoffs between holding down infection levels and sustaining economic activity until the vaccine is available.

Second, we may get some good news that leads us to choose to allow the infection to spread at a more rapid, but controlled, rate. This news could involve learning that some antiviral treatment or some practice like universal wearing of N95 masks reduces the health consequences

of infection or slows its spread (Abaluck et al. (2020) summarize existing evidence on the virtues of universal adoption of masks). Or we might learn that the virus is less harmful than is currently believed, e.g because asymptomatic cases are more common than we know.

Third, we may reach the point where we cannot or choose not to adopt policies that keep transmission rates in check, and infection proceeds to the point where herd immunity stops the spread. Current simulations suggest that the health consequences of reaching this endpoint rapidly would be very, very bad (Greenstone and Nigam (2020), see also Rowthorn (2020)). And there would likely be severe economic disruption even if we abandoned all government-mandated distancing efforts—few people would want to go to work, send their children to school, or eat in a restaurant if we reached a peak where 10% of the US population was currently infectious.

Flattening the rate at which the disease expands has the obvious advantage of reducing the possibility of overloading hospital systems and thereby reducing negative externalities. It also buys time, potentially allowing a larger fraction of those who will eventually become infected to benefit from more effective treatments that may be developed. “Flattening the curve” will also reduce the extent to which a disease is able to over-shoot the herd immunity threshold (since the incidence of the disease at the point that the threshold is reached determines the number of additional cases that will still occur before the disease dies out). For example, if we take the simple three-compartment model described at the beginning of the paper with R_0 of 2.6 and a 5.5 day half-life for the infectious period (i.e. $\beta = \log(R_0+1)/5.5 = 0.23$, $-\gamma = \log(1/2)/5.5 = -0.13$) and 0.0013% of the population (i.e. 100,000 out of 7.5 billion people) infected at the start of the simulation, then we would expect an “uncontrolled” spread of the disease to ultimately infect about 75% of the population, a peak incidence of about 13%, and for 99% of all infections

to occur within about 6 months. In contrast, if control methods are put into place to prevent the incidence from ever exceeding 1% of the population, only about 56% of the population will ultimately be infected (roughly equivalent to the herd immunity threshold in this case), but it would take more than one and a half years for 99% of the infections to occur.

That is, “flattening the curve” is no panacea, as it lengthens the time required to reach the herd immunity threshold and still is projected to produce a large number of deaths. It is possible that we will learn that asymptomatic infections are currently much larger than is known and this endpoint is not so bad in less densely populated areas. Some countries, especially in the developing world may find that they cannot avoid it even if it is as bad as feared. The same could be true in the US given the incomplete social distancing in place and rising discord regarding even these incomplete measures. But we focus here on thinking about optimal policies that reach one of the first two endpoints.

We think it is useful to think separately about two types of policies that may be implemented to alleviate health-related and economic suffering while waiting for an effective vaccine/treatment.

First, we may pursue an aggressive policy of testing and contact tracing. Such policies are thought to be a key part of how South Korea, Hong Kong, and Singapore have mostly kept infection rates in check. Absent failures to build testing capacity in the US, we could perhaps have achieved comparable success. Effective contract tracing, however, requires a capacity for testing and public health outreach that is far beyond what could be implemented today in the US. And once contract tracing alone cannot effectively mitigate the spread, and strong social distancing policies are adopted, the incremental value from contact tracing is reduced.

If, in the future, active cases have been reduced to a point where contact tracing is again feasible, it could become a central part of our approach to SARS-CoV-2. Controlling the disease while distancing only those who are known to be infected is extremely attractive from an economic viewpoint. Whether the US could hope to keep R_0 below one mostly by these means, however, is unclear. Many countries practicing contact tracing have done so in part by making extensive mandated use of phone-based and other geolocation data that the US has not yet contemplated. And even with such aggressive use of technology, reports from Hong Kong suggest that it has been challenging to keep R_0 below one without also implementing some social distancing measures. Much of the US is much less densely populated than is Hong Kong, so less aggressive contact tracing might suffice. It is even conceivable that testing-based policies may allow us to simultaneously reduce infections and sustain economic activity even if distancing cannot first reduce infections to a low level. For example, the cost and availability of home-pregnancy-style tests could fall to such a degree that a substantial fraction of the population could take one every day.⁴⁴ But we are now very far from that level, and have not been making the massive investments in testing research and developing the public health infrastructure that such an aspiration would seem to require.

Second, we will pursue some set of social distancing measures. Many variants of social distancing have been implemented. We find it useful to think of individual restrictions as lying on a continuum in terms of the ratio of economic costs incurred per unit of virus-slowness. At one extreme, there are some policies with a negative ratio – society would be better in both dimensions if long lines and crowded waiting rooms at the department of motor vehicles were

⁴⁴ Silver (2020) provides a detailed analysis of the implications of future testing scenarios given the possibilities of both false positives and negatives. See also https://twitter.com/zbinney_NFLinj/status/1245789672833417217.

replaced with expanded online services and scheduled appointments with minimal waits. Consensus seems to be shifting that universal wearing of facemasks may also provide substantial benefits – the slower spread of the virus and low number of deaths in Japan is intriguing and suggests that universal adoption of high-grade facemasks could be incredibly valuable. Costs are certainly minimal compared to most other interventions, so any substantial benefit would make the ratio attractive. Work-from-home and distance education have more substantial costs, but also presumably provide substantial benefits. Shutting construction sites, at least relative to continuing with modified work practices that increase social distancing, mask usage, and handwashing, seems like a yet larger ratio of economic costs to disease slowing benefits. Many other activities have not been suspended presumably because the economic costs are seen as prohibitive. Recent reductions in new hospitalizations are encouraging, but to some extent we have been taking it on faith that R_0 can be reduced below one without taking the more extreme measures used in China.⁴⁵

Absent successful contact tracing, reductions in social distancing would likely induce a transmission rate $R_0 > 1$ and could lead to recurrence of disease. Cyclical peaks are well known from prior diseases: measles outbreaks took place every few years for decades in England and the U.S. (Dalziel et al., 2016). There is reason to believe that subsequent outbreaks of the disease will be less intense, since people who previously contracted it may remain immune for some time, and will grow more slowly than the initial outbreak (and presumably, that they will therefore be easier to control). Nevertheless, in biological models there are rarely any guarantees, and a well-known counterexample to this phenomenon is the 1918 influenza

⁴⁵ Flaxman et al. (2020) examining European data up through March 28, 2020 estimate confidence intervals for R_0 which include a fairly wide range of values above and below one.

pandemic, for which the second wave of infections was more severe than the first in many locations. Theoretically, we would also not expect the effect to be very large unless the susceptible fraction has moved substantially away from one. If our current distancing efforts limit infections to just a few percent of the population, as we hope, then this effect would not be large.

Serology testing will be a third component of policy response. Having at least some random serology testing seems very important – the information it provides on how widespread infection has been is tremendously important for thinking about how the virus will spread and how severe the health consequences will be. It can also provide economic benefits in that those who are immune can safely return to work. But this economic benefit may not be large in practice, both because we hope that we can avoid a situation where a large share of the population is infected, and because there will be uncertainty about how strongly and for how long prior infection is protective.

An optimal policy for the period until vaccines/treatments become salient will presumably involve choosing some point on the social distancing spectrum and whether to attempt aggressive contact tracing at each point in time.⁴⁶ The costs and benefits of policies will naturally vary with the current prevalence of active infections, both because the number of future infections prevented by a unit decrease in R_0 is proportional to the number currently affected and because economic impacts will vary. For example, the loss of learning when schools move online will be smaller if the alternative was to have teachers calling in sick and leaving classes in

⁴⁶ It is unclear a priori whether the optimal control policy should be stationary. (We are grateful to Michael Kremer for suggesting this point.) Morris et al. (2020) and Alvarez, Argente, and Lippi (2020) provide separate analyses for SIR models where the optimal control policy is not stationary.

the hands of substitutes. Hence, the optimal policy will be an adaptive one with measures that adjust to prevailing conditions.

If the optimal policy keeps society in the region where the susceptible fraction S is close to one, then the optimal control problem of choosing an optimal policy for each (I, S) might be thought of as mostly just a one-dimensional problem where one must choose an optimal policy for each I . Depending on what we learn about fatality rates, rates of spread, and the effects of various policies on these characteristics, it is plausible that the optimal policy might take on one of several forms.

In an optimistic scenario, we may find that social distancing policies similar to those we have adopted, plus more face masks, can drive R_0 far below one. In that case, continuing these policies for a moderate period of time may get us to the point where we can switch to primary reliance on contact tracing combined with just the most efficient distancing policies on the spectrum. To be prepared to take advantage of such an opportunity, should it present itself, it might be sensible to start practicing for it right away by attempting to immediately institute aggressive contact tracing in areas of the country that are relatively isolated and where prevalence is low.

It is also plausible that the combination of contact tracing and highly efficient social distancing measures will not be enough to keep cases from expanding. In that case, it may be that we need to also adopt social distancing practices intermediate on the spectrum alongside contact tracing for quite some time. For example, it may be that schools will not open for 1-2 years, and we should start preparing for such eventualities.

There may also be inherent nonconvexities in the set of social distancing policies available to us. For example, it may be that the combination of shutting schools, shutting bars, restaurants, and retail stores, encouraging work from home, and limiting gathering is more efficient than an intermediate policy involving a subset of these elements. An optimal control framework might say that the best thing to do in such a situation is to adopt a policy where we mix in continuous time between the aggressive distancing and testing plus moderate distancing alternatives.

But in practice, there are presumably transition costs incurred whenever stores open and close, firms shift to work from home, and so on. In this case, the optimum could potentially be an alternation between more and less strict shutdown policies, with new shutdowns triggered every time the disease prevalence reaches some threshold. An appealing side-benefit of extensive testing and contact tracing is that it can make it easier to recognize when prevalence has moved into the region where testing hits its capacity constraint and more aggressive distancing is needed. Alternation between periods of low and high prevalence of disease could also be optimal in a world where occasional returns to more normal activity can help to maintain worker-firm connections and allow firms to occasionally clear bottlenecks in production and shipping that arise from social distancing.

A limitation of any such strategy is that periods with $R_0 > 1$ must be balanced with periods with $R_0 < 1$ to keep prevalence in check. Given that R_0 is well above one under minimal social distancing, we will be unable to spend much time in such a regime if the R_0 in the social distancing regime is only a little below one. There will also be a temptation to loosen restrictions in the summer if seasonal differences slow the spread at that time. Whether that is advisable depends on whether the effect of a restriction is smaller in proportional terms.

New policy discussions will need to occur whenever the best available treatment changes or when knowledge about the epidemiology changes. Most consequentially, as the set of available treatments improves we will need to think about whether social distancing restrictions then in force are still warranted.

An important consideration for all policy discussions is that we must recognize that we are making decisions under substantial uncertainty. For all the reasons we noted earlier, we still know very little about the epidemiology of SARS-CoV-2. And we are also following an economic path that is unprecedented in modern times and should recognize that there is great uncertainty about how quickly and well modern economies will recover from shorter vs. longer and more vs. less extensive shutdowns. Some heuristics that we may want to keep in mind for dealing with such uncertainty are that it may be useful to try to keep the economy and disease situation as close as we can to situations with which we have some experience, and that we should think about worst-case consequences of policies. Adaptive policies that implement stricter policies whenever prevalence is increasing are appealing in part because they are well adapted to uncertainty in these ways.

VIII. Conclusion and Opportunities for Further Research

There are some ways in which we are now better positioned than when we first faced the SARS-CoV-2 threat. Most importantly, we have more information about what we are facing. Current and future research is extremely valuable because we still face great uncertainty and there are momentous decisions to come.

In the past month, economists and political scientists have produced new and valuable research with papers that may guide the way for future efforts. There is an urgent need for better data (Stock (2020), Stock, Aspenlund, Droste, and Walker (2020)), and for creative and entrepreneurial methods of interpreting the limited data that is available (Fang et al (2020), Harris (2020)). It is not clear that models of the spread of disease have played a wholly positive role in shaping policy in the first three months of 2020. Similarly, some reports suggest that a misguided understanding of behavioral science (the fear of “behavioral fatigue”) delayed the response of the UK government to the spread of the pandemic.⁴⁷ But there have already been clear contributions in recent weeks for both behavioral approaches (Barari et al. (2020), Bricese et al. (2020)) and theoretical modeling (Atkeson (2020), Berger, Herkenhoff, and Mongey (2020)).

After an initial policy debate about whether and when to close down businesses and schools in February and March 2020, we can anticipate the next debate about how and when to reopen them in May and June 2020 and beyond. Using the information we have and will gain to make better policy choices in our second opportunity is critically important.

⁴⁷ <https://behavioralscientist.org/why-a-group-of-behavioural-scientists-penned-an-open-letter-to-the-uk-government-questioning-its-coronavirus-response-covid-19-social-distancing/>

References:

Abaluck, J. et al. (2020). “Universal Cloth Mask Adoption,” *Covid Economics*. 5:147-163, April 16, 2020.

Acemoglu, D. & Ozdaglar, A. (Fall, 2009). 6.207/14.15: Networks Lecture 8: Diffusion through Networks. MIT Open. <https://economics.mit.edu/files/4626>

Allcott, H., et al. (2020). “Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic,” NBER Working Paper #26946.

Alvarez, F. E., Argente, D., & Lippi, F. (2020). “A Simple Planning Problem for COVID-19 Lockdown,” NBER Working Paper #26981.

Atkeson, A., (2020b). “How Deadly Is COVID-19? Understanding the Difficulties with Estimation of Its Fatality Rate,” NBER Working Paper #26965.

Atkeson, A., (2020a). “What Will Be the Economic Impact of Covid-19 in the US? Rough Estimates of Disease Scenarios,” NBER Working Paper 26867.

Azzoni, T, and A. Dampf, (2020, March). “Game Zero?’ Soccer Game Attended by 40,000 Fans Likely Made This Italian City a Coronavirus Epicenter,” *Time*. <https://time.com/5809848/game-zero-soccer-game-italy/>

Barari, S. et al., (2020). “Evaluating COVID-19 Public Health Messaging in Italy: Self-Reported Compliance and Growing Mental Health Concerns,” Working Paper. <https://j.mp/39btyT2>

Bendavid E., & Bhattacharya, J. “Is the Coronavirus Epidemic as Deadly as They Say?” *Wall Street Journal*, March 24, 2020.

Bendavid E., et al. “COVID-19 Antibody Seroprevalence in Santa Clara County, California,” *Medrxiv*, <https://www.medrxiv.org/content/10.1101/2020.04.14.20062463v1.full.pdf>

Berger, D., Herkenhoff, K., & Mongey, S. (2020). “An SEIR Infectious Disease Model with Testing and Conditional Quarantine,” NBER Working Paper 26901.

Blackwood, J., & Childs, L. (2018). “An introduction to compartmental modeling for the budding infectious disease modeler,” *Letters in Biomathematics*, 5(1), pp.195-221.

Bolker, B. (1999). “Analytic models for the patchy spread of plant disease,” *Bulletin of Mathematical Biology*, 61(5), 849-874.

Booth, W. (2020). “A chilling scientific paper helped upend U.S. and U.K. coronavirus strategies,” *Washington Post*. March 17, 2020.

Borjas, G. (2020). “Demographic Determinants of Testing Incidence and COVID-19 Infections in New York City Neighborhoods,” *Covid Economics*: 3:12-39, April 10, 2020.

Branswell, H., (2020). “CDC launches studies to get more precise count of undetected Covid-19 cases,” *STAT*. April 4, 2020.

Brauer, F., Van den Driessche, P., & Wu, J. (2008). *Mathematical epidemiology* (Lecture notes in mathematics (Springer-Verlag). 1945. Berlin: Springer.

Bricese, G., Lacetra, N., Macis, M., & Tonin, M. (2020). “Compliance with COVID-19 Social-Distancing Measures in Italy: The Role of Expectations and Duration,” NBER Working Paper 26916.

Cereda, D., et al., (2020). “The early phase of the COVID-19 outbreak in Lombardy, Italy,” *ArXiv.org*. Mar 20, 2020. <https://arxiv.org/ftp/arxiv/papers/2003/2003.09320.pdf>

Chang, B. (2020). “France has converted its high-speed TGV trains into ambulances to transport critically ill coronavirus patients across the country,” *Business Insider*. April 1, 2020.

Chiou, L. & Tucker, C. (2020) “Social Distancing, Internet Access, and Inequality”. NBER Working Paper #26982.

Cummins, E. (2020). “Why the Coronavirus Test Gives So Many False Negatives,” *Slate.com*, April 6, 2020.

Dalziel, B., et al. (2016). “Persistent Chaos of Measles Epidemics in the Prevaccination United States Caused by a Small Change in Seasonal Transmission Patterns,” *Plos Computational Biology*, 12(2), E1004655.

Dawsey, J., Abutaleb, Y., Stein, J., Wagner, J. (2020). “Trump weighs restarting economy despite warnings from U.S. public health officials,” *Washington Post*, March 23, 2020

Dingel J. & Neiman B. “How many jobs can be done at home?” *Covid Economics*, 1:16-25, April 3, 2020.

Dushoff, J., (2020). “Preparing for the Storm,” March 23, 2020, <http://dushoff.github.io/notebook/covid.html>

Easley, D., & Kleinberg, J. (2010). *Networks, crowds, and markets: Reasoning about a highly connected world*. New York: Cambridge University Press.

Ellyat, H. (2020). “Germany has a low coronavirus mortality rate: Here’s why”, *CNBC.com*. April 3, 2020.

Epstein, R., “Coronavirus Perspective,” working paper. *Hoover Institution*. Stanford University, March 16, 2020. <https://www.hoover.org/research/coronavirus-pandemic>

Fang, H., Wang, L. & Yang, Y. (2020), “Human Mobility Restrictions and the Spread of the Novel Coronavirus (2019-NCOV) in China,” NBER Working Paper 26906.

Ferguson, N., et al., (2020). “Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand,” Working Paper, Imperial College London. <https://spiral.imperial.ac.uk:8443/handle/10044/1/77482>

Ferretti, L. et al., (2020). “Quantifying SARS-CoV-2 transmission suggests epidemic control with digital contact tracing,” *Science*. March 31, 2020.

Fink, S., (2020). “White House Takes New Line After Dire Report on Death Toll,” *New York Times*, March 16, 2020.

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

Gallos, L., & Fefferman, N. (2015). “The Effect of Disease-Induced Mortality on Structural Network Properties,” *PLoS One*, 10(8), E0136704.

Geddes, L., (2020). “Does a high viral load or infectious dose make covid-19 worse?” *New Scientist*. March 27, 2020

Gelman, A. (2020). “Concerns with that Stanford study of coronavirus prevalence,” <https://statmodeling.stat.columbia.edu/2020/04/19/fatal-flaws-in-stanford-study-of-coronavirus-prevalence/>

Giles, C. “Deaths from coronavirus far higher in England than first reported,” *Financial Times*. April 7, 2020.

Greenstone, M. & Nigam, V. (2020). “Does Social Distancing Matter?” *Covid Economics* 7:1-22, April 20, 2020.

Gupta, A. (2020). “Does Covid-19 Hit Women and Men Differently? U.S. Isn’t Keeping Track.” *New York Times*. April 3, 2020.

Harko, T., Lobo, F., & Mak, M. (2014). “Exact analytical solutions of the Susceptible-Infected-Recovered (SIR) epidemic model and of the SIR model with equal death and birth rates,” *Applied Mathematics and Computation*, 236(C), 184-194.

Harris, J. (2020). “The Coronavirus Epidemic Curve is Already Flattening in New York City,” NBER Working Paper No. 26917.

He, X., et al., (2020). “Temporal dynamics in viral shedding and transmissibility of COVID-19,” *Medrxiv*. <https://www.medrxiv.org/content/10.1101/2020.03.15.20036707v2.full.pdf>

Hengehan, C., Brassey, J., & Jefferson, T. (2020). “SARS-CoV-2 viral load and the severity of COVID-19,” Working Paper, *Centre for Evidence-Based Medicine*, Oxford University, March 26, 2020.

Henriques, M., (2020). “Coronavirus: Why death and mortality rates differ,” BBC.com. April 1, 2020.

Hethcote, H. (2000). The Mathematics of Infectious Diseases. *SIAM Review*, 42(4), 599-653.

Hethcote, H., & Driessche, W. (1991). “Some epidemiological models with nonlinear incidence,” *Journal of Mathematical Biology*, 29(3), 271-287.

Hethcote, H., & Van Ark, J. (1987). “Epidemiological models for heterogeneous populations: Proportionate mixing, parameter estimation, and immunization programs,” *Mathematical Biosciences*, 84(1), 85-118.

Hruby, D. (2020). “How an Austrian ski resort helped coronavirus spread across Europe,” CNN.com, March 24, 2020

Hurford, A. (2012). “Mechanistic Models – What is the Value of Understanding?” <https://theartofmodelling.wordpress.com/2012/02/19/mechanistic-models-what-is-the-value-of-understanding/>

Johnson, N., & Mueller, J. (2002). “Updating the Accounts: Global Mortality of the 1918-1920 ‘Spanish’ Influenza Pandemic,” *Bulletin of the History of Medicine*, 76(1), 105-115.

Johnston, J., (2020). “Patrick Vallance says keeping coronavirus death toll below 20,000 would be 'good outcome,’” *Politics Home*. March 17, 2020.

Kendi, I. (2020). “What the Racial Data Show. The pandemic seems to be hitting people of color the hardest,” *The Atlantic*, April 6, 2020

Kermack, W., & McKendrick, O. (1991a). “Contributions to the mathematical theory of epidemics—I,” *Bulletin of Mathematical Biology*, 53(1-2), 33-55.

Kermack, W., & McKendrick, O. (1991b). “Contributions to the mathematical theory of epidemics—II The problem of endemicity,” *Bulletin of Mathematical Biology*, 53(1-2), 57-87.

- Kermack, W., & McKendrick, O. (1991c). “Contributions to the mathematical theory of epidemics—III. Further studies of the problem of endemicity,” *Bulletin of Mathematical Biology*, 53(1-2), 89-118.
- King, A., Magpantay, F., & Rohani, P. (2015). “Avoidable errors in the modeling of outbreaks of emerging pathogens, with special reference to Ebola,” *ArXiv.org*, 282(1806), 20150347.
- Kissler, S., et al., (2020). “Projecting the transmission dynamics of SARS-CoV-2 through the post-pandemic period,” *MedRxiv*, <https://doi.org/10.1101/2020.03.04.20031112>.
- Kuchler, T., Russel, R., & Stroebel, J. (2020). “The Geographic Spread of COVID-19 Correlates with Structure of Social Networks as Measured by Facebook,” NBER Working Paper #26990.
- Lazar, K., Ryan, A. (2020) “How Accurate Are Coronavirus Tests?” *Boston Globe*, April 2, 2020.
- Lehman, C., Loberg, Clark, & Open Textbook Library. (2019). *Quantitative Ecology: A New Unified Approach*. University of Minnesota Libraries Publishing.
- Li, M., & Muldowney, J. (1995). “Global stability for the SEIR model in epidemiology,” *Mathematical biosciences*, 125(2), pp.155-164.
- Lively, C. (1999). “Migration, Virulence, and the Geographic Mosaic of Adaptation by Parasites,” *The American Naturalist*, 153(S5), S34-S47.
- Lourenco, J. et al., (2020). “Fundamental principles of epidemic spread highlight the immediate need for large-scale serological surveys to assess the stage of the SARS-CoV-2 epidemic,” *MedRxiv*. <https://doi.org/10.1101/2020.03.24.20042291>.
- Ludwig, D., & Malley, R. (2020). “Americans Are Already Too Diseased to Go Back to Work Right Now,” *New York Times*. March 30, 2020.
- Massad, E., Coutinho, F., Burattini, M., & Amaku, M. (2010). “Estimation of R 0 from the initial phase of an outbreak of a vector-borne infection,” *Tropical Medicine & International Health*, 15(1), 120-126.
- Morris, D., Rossine, F., Plotkin, J., & Levin, S. (2020). “Optimal, near-optimal, and robust epidemic control,” Working paper. 2020. <https://osf.io/rq5ct/>
- Murray, C. et al., (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>
- Rabinowitz, J. & Bartman, C. (2020). “These Coronavirus Exposures Might Be the Most Dangerous,” *New York Times*. April 1, 2020.

- Read, A.F. (1994). "The evolution of virulence," *Trends in microbiology*. 2(3), pp.73-76.
- Ridenhour, B., Kowalik, J.M., & Shay, D. K. (2014). "Unraveling R0: Considerations for public health applications," *The American Journal of Public Health*, 104(2), 32.
- Robinson, M., & Blanchard, S. (2020). "Chief scientific adviser wants 40 million Britons - 60 percent of the entire population - to catch coronavirus for herd immunity to take effect," *Daily Mail*. March 13, 2020
- Rogers, A. & Molteni, M. (2020). "The Mathematics of Predicting the Course of the Coronavirus," *Wired*. March 30, 2020.
- Rowthorn, R. "A Cost-Benefit Analysis of the Covid-19 Disease," *Covid Economics* 11:97-106, April 24, 2020.
- Sattenspiel, L. (1990). "Modeling the spread of infectious disease in human populations," *American Journal of Physical Anthropology*, 33(S11), 245-276.
- Silver, N. (2020). "Coronavirus Case Counts are Meaningless," *538.com*. April 4, 2020. <https://fivethirtyeight.com/features/coronavirus-case-counts-are-meaningless/>
- Spreeuwenberg, P., Kroneman, M., & Paget, J. (2018). "Reassessing the Global Mortality Burden of the 1918 Influenza Pandemic," *American Journal of Epidemiology*, 187(12), 2561-2567.
- Stein, R. (2011). "Super-spreaders in infectious diseases," *International Journal of Infectious Diseases*, 15(8), E510-E513.
- Stock, J. (2020). "Data Gaps and the Policy Response to the Novel Coronavirus," *Covid Economics* 3:1-11, April 10, 2020.
- Stock, J., Aspelund, K., Droste, M., & Walker, C. (2020). "Estimates of the Undetected Rate among the SARS-CoV-2 Infected using Testing Data from Iceland," Working Paper.
- Weise, E. (2020). "8 strains of the coronavirus are circling the globe. Here's what clues they're giving scientists," *USA Today*. March 27, 2020.
- Wood, S. (2001). "Partially Specified Ecological Models," *Ecological Monographs*. 71(1), 1-25.
- Zhou, F., et al., (2020). "Clinical course and risk factors for mortality of adult inpatients with COVID-19 in Wuhan, China: a retrospective cohort study," *The Lancet*, 395:10229.1054-1062. March 28, 2020.

Appendix: COVID-19 Data Resources (as of April 16, 2020)

We have compiled a list of data sources, mostly publicly available, that researchers may find useful. Some are just repackaging of official government statistics on infections, hospitalizations, deaths, and so forth, but, depending on the particular website, they might be from different governments and at different levels of aggregation. Others are entirely novel data sources, often shared by private firms who have collected them in the course of their business practice or as a public service. With the exception of the data on the 1918 Flu Pandemic, all data sets cover a period of time relevant for study of the current pandemic.

We offer the list with the caveat that new data dumps, curated repositories, and official statistics are becoming available on a daily basis. Furthermore, links might become stale as websites are moved or removed. (We try to offer searchable descriptions to help researchers track down the data sources.)

1. Statistics on infections, hospitalizations, recoveries, and deaths at the **country** level
 - a. <https://data.humdata.org/dataset/novel-coronavirus-2019-ncov-cases>
 - b. European CDC data <https://www.ecdc.europa.eu/en/covid-19-pandemic>
 - c. Institute for Health Metrics and Evaluation data <http://www.healthdata.org/covid>
2. Statistics on infections, hospitalizations, recoveries, and deaths at the US **state** level, as well as data on ventilator, hospital, and ICU usage and number of tests
 - a. <https://ourworldindata.org/covid-testing>
 - b. <https://covidtracking.com/data>
3. Statistics on infections, hospitalizations, recoveries, and deaths at the US **county** level
 - a. <https://usafacts.org/visualizations/coronavirus-covid-19-spread-map/> (where it says Download data (Jan. 22- April 6))
4. NY Times data at the US **county** level of confirmed cases and deaths by day
 - a. <https://www.nytimes.com/article/coronavirus-county-data-us.html>, available at <https://github.com/nytimes/covid-19-data>
5. New York City-specific data
 - a. <https://www.nytimes.com/article/coronavirus-county-data-us.html>, available at <https://github.com/nychealth/coronavirus-data>
6. Historic data for 43 US cities on the 1918 Flu Epidemic, including deaths, death rates, and non-pharmaceutical interventions (such as social distancing), from US Census data---used in Markel, Lipman, and Navarro (2007), but not clear if and where the data are available for download
 - a. <https://jamanetwork.com/journals/jama/fullarticle/208354>
7. Data on the timing of non-pharmaceutical interventions (social distancing, closure of schools and universities, closure of nonessential business, gathering size limitations, etc.) in the US
 - a. www.keystonestrategy.com/covid-19/, available at <https://github.com/Keystone-Strategy/covid19-intervention-data/>
8. Data on distances traveled using cell phone movements from Unacast---can view and interact with their mobility map, but not download the data. (They may be offering it to universities and non-profits for free.)
 - a. <https://www.unacast.com/covid19/social-distancing-scoreboard>
9. Mobility data is becoming available from other sources as well, such as Google, CityMapper, and SafeGraph
 - a. <https://www.google.com/covid19/mobility/>
 - b. <https://citymapper.com/cmi>
 - c. <https://safegraph.com>

10. Restaurant reservation data by city and day for Australia, Canada, Germany, Ireland, Mexico, United Kingdom, and the United States from Open Table
 - a. <https://www.opentable.com/state-of-industry>
11. Data on airport usage---can view map but not available for download
 - a. <https://wanderlog.com/coronavirus-airports-effect>
12. Data on fever incidence from Kinsa, a firm that makes smart thermometers---can view infographics, but not available for download
 - a. <https://www.kinsahealth.co/images-and-infographics-from-kinsas-health-weather-map-and-data/>
13. Data on air pollution (at weekly or daily frequency)
 - a. For European countries: <https://www.eea.europa.eu/themes/air/air-quality-and-covid19/monitoring-covid-19-impacts-on>
 - b. Worldwide: <https://www.covidexplore.com/PM25> (<https://github.com/mayukh18/covidexplore>)
 - c. India and China (though they might have data on more countries): <https://energyandcleanair.org/blog/>
 - d. <https://aqicn.org/data-platform/covid19/>
14. Satellite data on air pollution data in Europe
 - a. https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Sentinel-5P/Coronavirus_lockdown_leading_to_drop_in_pollution_across_Europe (<https://scihub.copernicus.eu>)
15. Data from Foursquare on consumer location, derives measures of impacts on flights, stocking up behavior, sit-down restaurants, fast food, etc.---can view article, but data not available for download
 - a. <https://enterprise.foursquare.com/intersections/article/understanding-the-impact-of-covid-19/>
16. Data on planned, ongoing, and completed trials for medical interventions
 - a. <https://covid-evidence.org>
17. Google Trends, a free tool offered to allow researchers to download (normalized) data on search volumes over time and across geographic regions
 - a. <https://trends.google.com/trends/>
 - b. Example article using these data <https://www.nytimes.com/2020/04/05/opinion/coronavirus-google-searches.html>

Which workers bear the burden of social distancing policies?¹

Simon Mongey,² Laura Pilossoph³ and Alex Weinberg⁴

Date submitted: 26 April 2020; Date accepted: 28 April 2020

What are the characteristics of workers in jobs likely to be initially affected by broad social distancing and later by narrower policy tailored to jobs with low risk of disease transmission? We use O/NET to construct a measure of the likelihood that jobs can be conducted from home (a variant of Dingel and Neiman, 2020) and a measure of low physical proximity to others at work. We validate the measures by showing how they relate to similar measures constructed using time use data from ATUS. Our main finding is that workers in low-work-from-home or high-physical-proximity jobs are more economically vulnerable across various measures constructed from the CPS and PSID: they are less educated, of lower income, have fewer liquid assets relative to income, and are more likely renters. We further substantiate the measures with behavior during the epidemic. First, we show that MSAs with less pre-virus employment in work-from-home jobs experienced smaller declines in the incidence of 'staying-at-home', as measured using SafeGraph cell phone data. Second, we show that both occupations and types of workers predicted to be employed in low work-from-home jobs experienced greater declines in employment according to the March 2020 CPS. For example, non-college educated workers experienced a 4ppt larger decline in employment relative to those with a college degree.

1 Replication codes and data for this paper are available on the authors' websites. Thanks to SafeGraph for making their data available to us, as well as other researchers studying the consequences of the Coronavirus epidemic. Thanks to Gianluca Violante and Greg Kaplan for making available their codes. Our measures at the three digit occupation level are available on our websites. The views expressed in this study are those of the author and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

2 Kenneth C. Grin Department of Economics, University of Chicago.

3 Federal Reserve Bank of New York.

4 Kenneth C. Grin Department of Economics, University of Chicago.

1 Introduction

The key public policy implemented during the Coronavirus pandemic is ‘social distancing’. This has the economic consequence that many workers will be forced to work from home if feasible. Moreover, returning to work will likely occur more slowly for jobs that require a large degree of physical proximity to others.¹ To the extent that workers vary systematically across these jobs, social distancing policies will have systematically different effects across individuals. Understanding how individuals vary across these occupations is therefore important for policy makers interested in formulating targeted worker assistance programs.

In this paper, we combine multiple data sources to study how individuals vary across occupations which differ in their exposure to social distancing policies. We merge *individual-level* data from the Bureau of Labor Statistics’ *Current Population Survey* (CPS) and the *Panel Study of Income Dynamics* (PSID) with a version of the [Dingel and Neiman \(2020\)](#) classification of an occupations’ capacity to work from home as well as a measure of physical proximity in the workplace.² We construct these two measures using *occupation-level* data from the Department of Labor’s *Occupational Information Network* (O*NET) data.³ We show that despite being negatively correlated, some outlier occupations such as those related to education are both high work-from-home and high physical-proximity, hence relatively more affected when social distancing policies become targeted.

We validate the measures of work-from-home and physical-proximity using data from the *American Time Use Survey* (ATUS) that are meant to capture similar types of job characteristics. The O*NET work-from-home measure does not explicitly account for working at home, since it is designed to capture whether a job *could feasibly* be done from home, and instead is based on the types of activities conducted at work (e.g. heavy lifting, working outdoors etc). Nonetheless, the measure is highly correlated with the share of time working that *is* spent at home in ATUS. Moreover, the O*NET physical-proximity measure is correlated with the reported fraction of time spent working alone. We hope that this validation is useful for other researchers.

With validated occupation-level measures in hand, we proceed in two steps. First, we study how individual characteristics of workers vary across these types of occupations. Our main result is that workers in occupations that are more likely to be affected by social distancing policies are workers we would consider more economically vulnerable. Workers in these occupations are less likely to have a college degree and are less likely to have health insurance provided by their employer. They are less likely to be white, less likely to work at a large firm, and less likely to be born in the USA. Workers in low work-from-home occupations also have disproportionately low levels of liquid assets,

¹For example, Governor Andrew Cuomo’s policy for NY State consists of a Phase I reopening with *Construction* and *Manufacturing* jobs, which the state views as *low risk* and highly essential.

²See https://bfi.uchicago.edu/wp-content/uploads/BFI.White-Paper_Dingel_Neiman.3.2020.pdf

³In these occupation-level data, occupational classifications are finer than those available in the individual-level data. To make the data conformable we develop a cross-walk that allows us to use the Bureau of Labor Statistics’ *Occupation Employment Statistics* (OES) to employment weight O*NET measures within the coarser occupations defined in the CPS. Code is available on request.

which is especially important for policies that provide liquidity to households. Finally, we show that these effects are monotonic in that occupations that score *relatively lower* (higher) in terms of the work-from-home (personal-proximity) measure, are *even more* economically vulnerable.⁴

These relationships are in general stronger when we split occupations by the work-from-home measure rather than the physical-proximity measure. Relative to occupations that score low in terms of the work-from-home measure, there is greater economic diversity among occupations associated with high levels of physical proximity. For example, salon workers, sales assistants and dentists all work in high personal-proximity environments. This suggests that the economic costs of social distancing policies may be more tightly related to pre-epidemic economic status, while the economic costs of a slow return to work that starts with low physical-proximity occupations may be more broadly distributed.

Second, using the limited data available since the start of the epidemic, we show how different occupations and workers have been effected. To begin, we construct a pre-epidemic MSA-level work-from-home measure and show that it is (i) uncorrelated with pre-epidemic mobility as measured using cellphone data from *SafeGraph*, but (ii) strongly correlated with the change in mobility during the epidemic. We then turn to employment changes across the February and March 2020 CPS surveys. Occupations that rank low in the work-from-home measure and high in the physical-proximity measure experienced larger employment declines relative to pre-epidemic February-March changes. Finally, workers that our main results suggested would be more vulnerable did indeed experience larger declines in employment. For example, non-college educated workers experienced a 4ppt larger decline in employment relative to those with a college degree.

Our results have clear implications for public economic policy. First, they provide guidance as to how income replacement and liquidity injection policies may be targeted. Second, since low work-from-home and high physical-proximity workers tend to have lower incomes and lower liquidity, the marginal social cost of income support is low, while the marginal private benefits are high. Third, the correlation between low work-from-home and high physical-proximity jobs creates a double-edged sword. It induces a correlation between *economic risks* under tight social distancing and *health risks* under relaxed social distancing. Already more economically vulnerable workers are disproportionately exposed to unemployment now, and infection in the future, suggesting the need for on-going policy interventions.

Literature. Our two contributions are to look at the *characteristics of workers* in occupations that differ by the aforementioned measures, and study the experience of these workers in the epidemic, using carefully validated occupation-level measures. [Dingel and Neiman \(2020\)](#) use the OES to ask the important question of what fraction of employment and income is accounted for by jobs

⁴When we compare the top quartile of occupations by the work-from-home measure, to the bottom quartile of occupations by the work-from-home measure, we find that the estimated treatment effects are larger. When we compare the third quartile of occupations by this measure, to the second quartile of occupations by this measure, we find that the estimated treatment effects are smaller but still statistically significant in all cases.

that can be done from home. Leibovici et al. (2020) conduct a similar analysis, instead considering low physical-proximity occupations rather than high work-from-home occupations.⁵ Both use the O*NET to classify occupations, and then employment and income data from the OES to study the geographic distribution of employment and income accounted for by types of jobs. Our focus here is on understanding the characteristics of the underlying workers that comprise employment in these jobs, validating the measures by showing they are consistent with measures from other datasets, and verifying that they are indeed correlated with post-outbreak outcomes. This requires a careful merging of the O*NET data by occupational code with other datasets containing worker characteristics, such as the CPS, the PSID, and the ATUS.

Overview. Section 2 describes how we construct our measures of work-from-home and physical-proximity using the O*NET and OES data. We compare the two measures across occupations, and validate each against comparable measures constructed from the ATUS. Section 3 integrates the CPS and PSID data and gives our main results, which are summarized in Figure 3. Section 4 shows how occupations characterized by their work-from-home and physical-proximity measures behaved over the implementation of social-distancing. Section 5 concludes.

2 Low work-from-home and high physical-proximity jobs

Here we describe how we construct our measures of work-from-home and personal-proximity, how the two measures compare across occupations, and then validate the two measures against ATUS data.

2.1 Characteristics of jobs

To construct occupation-level measures that can be merged into worker-level data from ATUS, CPS and PSID, we combine two data sets:

O*NET - Occupation-level data on work activities by occupation, where occupations are defined at the fine SOC level. SOC codes are finer than the Census OCC codes used to define occupations in ATUS, CPS, and PSID.

OES - Occupation-level data on employment and income at the SOC level. This data is used to employment-weight when aggregating skills across SOC level occupations, within OCC level occupations.

Rather than pooling data over time we take the most recent snapshot of the US economy available to us concurrently from these data, which is 2018. We use O*NET Database 24, OES data from

⁵St. Louis Federal Reserve, *On the Economy* blog:
<https://www.stlouisfed.org/on-the-economy/2020/march/social-distancing-contact-intensive-occupations>

2018, and the 2019 March CPS which asks questions regarding occupation and income in the prior year. In Section 4 we add the 2020 March CPS.

Using the O*NET and OES data we construct two measures for each occupation. We sign these in terms of expected negative economic impacts of the crisis: (i) *low work-from-home (LWFH)*, and (ii) *high physical-proximity (HPP)*. For our main results, we split occupations into groups based on these measures, and then use the CPS and PSID to compare the attributes of workers in occupations in each group.

We first detail how we construct *LWFH* and then describe the construction of *HPP* which follows many of the same steps. Let $j \in \{1, \dots, J\}$ denote a 3-digit OCC-code occupation, which is the measure available in worker-level data. Let $l \in \{1, \dots, L\}$ denote the fine SOC-code categorization of occupations in O*NET and OES. We differ from [Dingel and Neiman \(2020\)](#) in how we aggregate skills, but use their set of O*NET job characteristics.

1. We take the following 17 measures of SOC-level occupation attributes in the O*NET data from [Dingel and Neiman \(2020\)](#). We index them by $k = 1, \dots, K$. In the publicly available data, each takes on a value $m_{lk} \in [1, 5]$, representing the average response of workers to an underlying survey in which the options are $\{1, \dots, 5\}$:

- **Work Activities module:** Performing General Physical Activities; Handling and Moving Objects; Controlling Machines and Processes; Operating Vehicles, Mechanized Devices, or Equipment; Performing for or Working Directly with the Public; Inspecting Equipment, Structures, or Material; Repairing and Maintaining Electronic Equipment; Repairing and Maintaining Mechanical Equipment.
- **Work Contexts module:** Electronic Mail Use;⁶ Outdoors, Exposed to Weather; Outdoors, Under Cover; Deal With Physically Aggressive People; Wear Specialized Protective or Safety Equipment such as Breathing Apparatus, Safety Harness, Full Protection Suits, or Radiation Protection; Wear Common Protective or Safety Equipment such as Breathing Apparatus Safety Harness, Full Protection Suits, or Radiation Protection; Spend Time Walking and Running; Exposed to Minor Burns, Cuts, Bites, or Stings; Exposed to Disease or Infections.

Within each 3-digit OCC occupation j , we take the employment-weighted average of m_{lk} across SOC occupations $l \in j$. This gives a measure for each occupation-attribute pair: $\bar{m}_{jk} = \sum_{l \in j} \omega_l m_{lk}$, where $\omega_l = n_l / \sum_{l' \in j} n_{l'}$. To map SOC code occupations into OCC code occupations we start with a cross-walk obtained from US Census, which we then substantially edit and verify.⁷

2. We convert these into binary variables $m_{jk}^* \in \{0, 1\}$ based on whether $\bar{m}_{jk} \geq 3.5$. Since we employment-weighted when computing \bar{m}_{jk} then $m_{jk}^* = 1$ if “*The average respondent to the*

⁶In the case of Electronic Mail Use, we reverse the values such that a *high value* implies that the occupation is less suited to working from home.

⁷The basic cross-walk from Census is available here: <https://www2.census.gov/programs-surveys/demo/guidance/industry-occupation/2010-occ-codes-with-crosswalk-from-2002-2011.xls>

question in the underlying O*NET survey answered at least 4.”

3. We then construct a single measure for each occupation \overline{LWFH}_j by taking the unweighted mean of m_{jk}^* : $\overline{LWFH}_j = K^{-1} \sum_{k=1}^K m_{jk}^*$. In words, this gives the fraction of the K low work-from-home measures \overline{m}_{jk} for which occupation j has a high score. We rescale this to $[0, 1]$ by subtracting the minimum value and dividing by the maximum minus the minimum values.
4. We then assign the binary variable $LWFH_j^* = 1$ (*low work-from-home*) if occupation j is *above* the employment-weighted median value of \overline{LWFH}_j .

This procedure delivers a continuous variable \overline{LWFH}_j and a binary variable $LWFH_j^*$ that can be mapped into the occupational codes contained in the CPS, ATUS, and PSID.

The procedure to construct \overline{HPP}_j and HPP_j^* is similar to the above. We start with a measure m_l from O*NET at the SOC level that reflects physical-proximity at work and takes on a value $m_l \in [1, 5]$. We use the OES to compute an employment-weighted mean \overline{m}_j for all SOC occupations $l \in j$. We then re-scale to the interval $[0, 1]$ by subtracting \overline{m}_j^{Min} and dividing by $(\overline{m}_j^{Max} - \overline{m}_j^{Min})$, which gives our measure \overline{HPP}_j . We then assign the dummy $HPP_j^* = 1$ (*high physical-proximity*) if the occupation is *above* the employment-weighted median of this variable. For context, before this scaling, the median value of \overline{m}_j is 3.6. Therefore, in the underlying survey, the average worker in jobs we have classified as high physical proximity reports working at or within arm’s length of others.⁸

To recap, by construction HPP_j^* and $LWFH_j^*$ are binary variables that equal 1 for the occupations that are *most* likely to be effected by the epidemic and ensuing policies. Half of employment is in $HPP_j^* = 1$ jobs and half of employment is in $LWFH_j^* = 1$ jobs.

2.2 Which jobs are low work-from-home and high physical-proximity?

Figure 1 shows how occupations vary across these two metrics, and where our cut-offs lie for the binary measures.⁹ Unsurprisingly, there is a strong positive correlation between low work-from-home and high physical-proximity occupations. Typical ‘office jobs’ in financial service provision or the legal profession deliver $\tilde{m}_{jk} = 0$ for almost all of the K features of work used to construct the low work-from-home measure. There is also little work done within arm’s length in these jobs. On

⁸Workers that respond to the survey administered by O*NET choose one of: 1 = ‘I don’t work near other people (beyond 100ft)’, 2 = ‘I work with others but not closely (e.g. private office)’, 3 = ‘Slightly close (e.g. shared office)’, 4 = ‘Moderately close (at arm’s length)’, 5 = ‘Very close (near touching)’. Publicly available O*NET data consists of an average of these responses. Since the cut-off value of \overline{m}_j is 3.6, then $HPP_j^* = 1$ for occupations in which the average response to the survey question is at least 4. Our high physical-proximity occupations therefore represent occupations for which the average respondent said they worked at arm’s length or less away from others. For additional information regarding this question, see <https://www.onetonline.org/find/descriptor/result/4.C.2.a.3>.

⁹For readability of this figure, we employment-weight using the OES to aggregate \overline{LWFH}_j and \overline{HPP}_j to the 2 digit level. We then linearly transform each measure \overline{X}_j using its minimum and maximum: $(\overline{X}_j - \overline{X}_j^{Min}) / (\overline{X}_j^{Max} - \overline{X}_j^{Min})$.

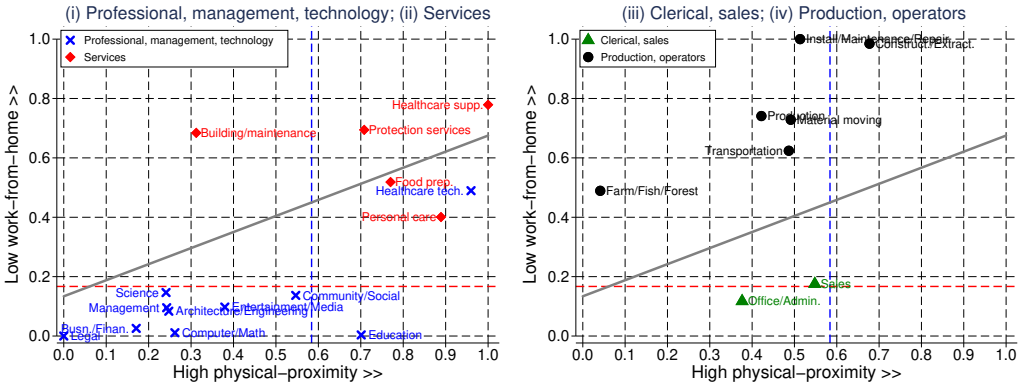


Figure 1: Occupations by Work-from-home and Physical-proximity (2 digit, Census OCC)

Notes This figure compares groups of 2 digit OCC code occupations. To construct this figure, we employment-weight using the OES to aggregate $LWFH_j$ and HPP_j to the 2 digit level. The gray line gives fitted values are from an employment-weighted linear regression across 2 digit occupations. Occupations *above* the red-dashed line have $LWFH_j^* = 1$, and account for half of employment. Occupations *to the right* of the blue-dashed line have $HPP_j^* = 1$, and account for half of employment.

the other hand, construction, material moving, and healthcare jobs are low work-from-home and high physical proximity.

A number of occupations stand out as deviations from this pattern. Education jobs require close physical-proximity, but little of the features that would prevent the job being conducted at home. Under broad public policies of social-distancing, workers in these jobs can successfully stay employed while operating from home, which has occurred through virtual teaching. More targeted social-distancing policies in the recovery phase could be expected to feature education jobs late to reintegrate. Agricultural jobs (Farm/Fish/Forest), meanwhile, may pose lower contagion risk due to low physical-proximity, but are difficult to be done from home. Such jobs may be punished somewhat unduly by indiscriminate social-distancing policies.

2.3 Comparison to measures constructed from ATUS

We construct ATUS-based measures of a job’s ability to be done from home as well as the degree of physical-proximity to others involved. From the 2018 microdata files we use information on the share of work time spent in certain places and with certain people. As a work-from-home measure, we compute the aggregate share of work hours that are spent at home. As a physical-proximity measure, we compute the aggregate share of work hours spent alone.¹⁰

¹⁰We use the question in the ATUS “who” file which asks - for each activity the respondent recorded - “Who was in the room with you/Who accompanied you?” for the measure of hours working spent alone. We use the question from the interview file which asks “where were you during this activity?” for the measure of hours spend working at home.

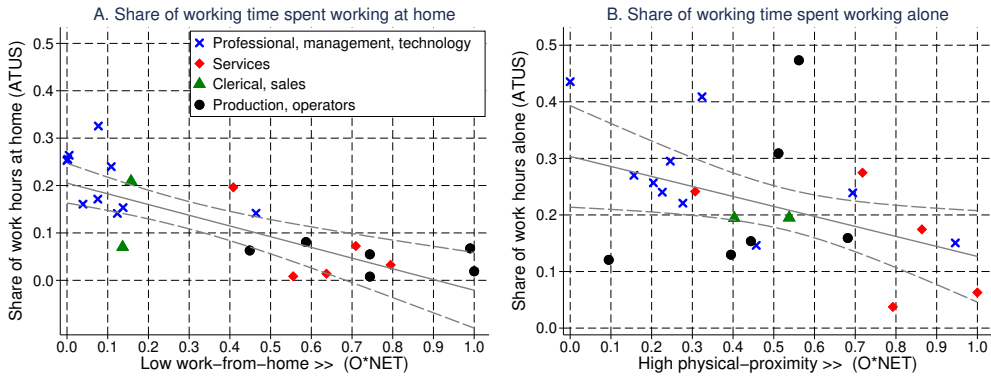


Figure 2: Comparing work-from-home and physical-proximity measures to ATUS

Notes This figure compares the fraction of individuals reporting that they can work from home in the ATUS against the O*NET WFH measure (Panel A). Panel B compares the physical proximity measures constructed from the two datasets. The share of adjusted work hours accounts for the fact that respondents can answer that they are with multiple individuals while performing a particular activity. Fitted values are from employment-weighted linear regressions, and display 95 percent confidence intervals for the conditional expectation of the dependent variable.

Figure 2A plots the ATUS measure of share of hours worked from home against the O*NET $LWFH_j^*$ measure for 2 digit occupations (Panel A). Figure 2B compares physical proximity measures. Both sets of measures are negatively correlated, which we view as validation of the O*NET measures. In particular, the work-from-home measures have a correlation of -0.83 . The physical proximity measure is less tightly linked, with a correlation of -0.44 , but this is to be expected given that the ATUS measure uses information on whether people are around you when you work, while the O*NET measure uses information on how close by those individuals are.

Figure 2 provides preliminary evidence on the distributional effects of social distancing public policies. Workers in professional services jobs (blue markers) already spend a significant fraction of time working from home and more time working alone. These types of workers—usually higher income and college educated—will be less likely impacted by social distancing policies. We now study this in detail using individual-level data.

3 Characteristics of workers by jobs

With $LWFH_j$ and HPP_j measured for occupations that are consistent with the CPS and PSID, we can compare the characteristics of workers in occupations for which these measures are high and for which these measures are low. Throughout we use the sample selection criteria of Heathcote et al. (2010) as well as their construction of wages.¹¹

¹¹ We follow their sample selection criteria for their Sample C which is as follows. We construct wages by taking total wage and salary income plus two-third of self-employment income, and dividing this by total hours worked which we compute as the product of weeks worked last year and usual weekly hours. We keep individuals that have: age

Our approach is simple. Let y_{ij} be a characteristic of a worker i that reports that they mostly worked in occupation j last year.¹² For simplicity, we only consider binary variables in the CPS; for example we construct a variable $y_{ij} = 1$ if the continuous variable ‘wage’ is above the median. We then estimate the following regression for each of our observables, using $LWFH_j^*$ as an example:

$$y_{ij} = \alpha_y + \beta_y LWFH_j^* + \varepsilon_{ij}. \quad (1)$$

We then plot the values for $\hat{\beta}_y$. This sample moment gives

$$\hat{\beta}_y = \mathbb{E} [y_{ij} | LWFH_j^* = 1] - \mathbb{E} [y_{ij} | LWFH_j^* = 0] \quad ,$$

where \mathbb{E} is the sample mean. Given that y_{ij} is binary, $\hat{\beta}_y$ simply gives the fraction of workers for which $y_{ij} = 1$ in *low* work-from-home occupations, relative to the fraction of workers for which $y_{ij} = 1$ in *high* work-from-home occupations. Clearly $\hat{\beta}_y \in [-1, 1]$ and takes the maximum value of 1 when $y_{ij} = 1$ for all individuals for which $LWFH_j^* = 1$, and $y_{ij} = 0$ for all individuals for which $LWFH_j^* = 0$. Comparing estimates across measures y and y' , a higher value of $\hat{\beta}_y > \hat{\beta}_{y'}$ can be interpreted as

“Workers in occupations for which $LWFH_j^ = 1$ are relatively more different from workers in occupations for which $LWFH_j^* = 0$ along dimension y than along dimension y' ”.*

We estimate (1) for a number of individual characteristics. In each case we assign $y_{ij} = 1$ to the individuals with the characteristic most related to being in a low work-from-home occupation. This gives $\hat{\beta}_y \in [0, 1]$. With this approach, we have the following characteristics of workers, all of which take on a value $y_{ij} = 1$:

- **Demographics.** (i) Non-white, (ii) No college degree, (iii) Age below 50, (iv) Male, (v) Single, (vi) Born outside USA, (vii) Non-US citizen, (viii) Rent their home
- **Work.** (i) No healthcare provided by employer,¹³ (ii) Employed at a small firm (< 500 employees), (iii) Part-time employed
- **Income.** (i) Below median wage,¹⁴ (ii) Experienced a spell of unemployment in the last year.

Work and income characteristics are associated with the job at which the worker was employed for the longest period of time in 2018.

In Figure 3A, in blue, we plot the estimates for each of these characteristics for the low work-from-home regression, ordering these attributes from the highest to the lowest point estimate. Figure 3B repeats the exercise for the high personal-proximity regression. We discuss low work-from-home first, and defer our discussion to the PSID measures in red.

25-65, wages that are greater than half of the Federal minimum wage, and total hours worked more than one working month of 8 hour days (176hrs).

¹²We use the IPUMS coding of the March CPS in which this is *OCCLY*.

¹³We set the indicator for employer provided healthcare to 1 if the employer pays for any part of the individual’s health insurance premiums.

¹⁴See footnote 11 for computation of wage.




Figure 3: Characteristics of workers in *Low Work-from-home* and *High physical-proximity* jobs

Notes This figure plots estimates of $\hat{\beta}_y$ for 10 characteristics y from regressions in which $LWFH_j^* \in \{0,1\}$ is the independent variable (Panel A), and in which $HPP_j \in \{0,1\}$ is the independent variable (Panel B). If $x\%$ of workers in *high work-from-home* occupations have, for example, *no college degree*, then Panel A shows that $(x+38)\%$ of workers in *low work-from-home* occupations have *no college degree*. A high value of $\hat{\beta}_y$ therefore means that workers in *low work-from-home occupations* are more likely than workers in *high work-from-home occupations* to be in the category listed on the vertical axis. Point estimates are given by the circle markers, and 95 percent confidence intervals are given by the lines through each marker. All blue results are derived from the CPS, red results are derived from the PSID.

Economic and demographic. Our main result is that occupations that score low in terms of the work-from-home measure feature workers that by all measures are economically more vulnerable. Workers in these occupations are less likely to be white or to have a college degree, which relate to the fact that they are 25 ppt more likely to be below median income. They are more likely to work in smaller firms, which are on average less financially robust and so less likely to suffer from the financial effects of the crisis (Chodorow-Reich, 2014). They are more likely to rent rather than own their homes and so will not be in positions to take advantage of interest rate cuts, and have fewer collateralizable assets to borrow against to compensate for earnings losses.

Workers in these jobs are also less likely to have access to *informal insurance channels* that may help them weather the crisis. They are less likely to be married, which diversifies household income against individual income risk. They are less likely to be US citizens or born in the US, which may lead to less family support, as well as restricted access to emergency government programs. Finally, workers in low work-from-home occupations are more likely to have unstable employment. They are less likely to be employed full-time and more likely to have recently experienced unemployment.

Healthcare. Availability of healthcare is obviously a key insurance mechanism in a health crisis. Workers in low work from home occupations are less likely to have any employer-provided healthcare.

Meanwhile those in jobs that are more readily able to be performed from home are more likely to have employer provided healthcare.

Age. The mortality rate for those with COVID-19 is significantly higher for older individuals.¹⁵ However, we find that the age of workers across these high- and low- work-from-home occupations does not systematically differ. Workers in low work-from-home jobs have the same fundamental health risks, but to the extent that these are also high physical-proximity jobs, higher occupation related health risks.

High physical-proximity. For most of the individual characteristics the results for high work-from-home occupations and low physical-proximity occupations are the same in terms of their sign. For example, workers in both high physical-proximity occupations and low work-from-home occupations are less likely to have a college degree than workers in low physical-proximity and high work-from-home occupations, respectively. The results, however, are less stark, as evidenced by the magnitudes of the coefficients. Differences in workers across high and low personal-proximity occupations are less pronounced than the differences between workers in low and high work-from-home occupations. If we consider high personal proximity occupations to be slower to be brought back as social distancing policies unwind, this suggests that the slow recovery may be more broadly experienced than the concentrated effects of unconditional social distancing.

Nonetheless, the correlation between low work-from-home and high physical-proximity jobs creates a double-edged sword. It induces a correlation in *economic risks* due to policy and *health risks* due to transmission of Coronavirus. More vulnerable workers are therefore relatively more exposed to both.

Sex. The results differ across these two measures most sharply for sex. Individuals in occupations that score highly in terms of work-from-home are more likely to be women, but this is also true for occupations that have high physical-proximity. This relates to the earlier example of Education jobs from Figure 1, which are female-dominated. Taking these results at face value, female workers may be relatively less affected by the universal social distancing measures currently in place, but could be relatively more affected in the future as these restrictions are targeted toward occupations with higher personal-proximity.

Liquidity. We expect that low access to liquid savings will compound the economic consequences of job loss or reduction in hours. To understand whether workers in low work-from-home jobs have disproportionately lower levels of liquid savings we add data from the PSID and construct measures of *hand-to-mouth-ness* of households following Kaplan and Violante (2014).¹⁶ Hand-to-

¹⁵See <https://www.cdc.gov/nchs/nvss/vsrr/covid19/index.htm>.

¹⁶We use the code made publicly available from Kaplan et al. (2014).

mouth households are households with liquid assets that are less than half of one month's income.

Results are shown in red in Figure 3. We find that households in which the highest earner is employed in a low work-from-home or high physical-proximity job are disproportionately hand-to-mouth. Conditional on being hand-to-mouth, households may be poor-hand-to-mouth or wealthy-hand-to-mouth depending on whether they have positive or negative net-assets, respectively. Conditional on being hand-to-mouth, workers in jobs most likely affected by social distancing policy are disproportionately poor-hand-to-mouth.

The magnitudes of the point estimates are economically significant. Hand-to-mouth low work-from-home households are 10 percent more likely to be *poor hand-to-mouth* than hand-to-mouth high work-from-home households. To put this in perspective we could compare this to how, as households age, the composition of hand-to-mouth households shifts from poor- to wealthy-. Starting at age 30, one would need to move all the way up to age 50—a period of high income growth—in order to obtain a 10 percent decline in the fraction of hand-to-mouth households that are poor hand-to-mouth (Kaplan et al., 2014, Figure 6). Despite not being significantly younger, low work-from-home households have finances *as if* they are twenty years further back in their lifecycle.

Comparing extremes. A policy maker might not be interested in policies targeted below and above the median of the indexes we create since they rule in too many individuals. We therefore verify that if we make more extreme comparisons using the tails of our measures, then we get more extreme results in terms of the economic situation of households. Figure A1 in Appendix A compares the lower quartile to the upper quartile (dropping the middle quartile, in red), and the second quartile to the third quartile (dropping the upper and lower quartile, in green). When we compare workers in *very low* work-from-home occupations to workers in *very high* work-from-home occupations (in red), the coefficients are uniformly larger in magnitude. For example, workers in the lowest quartile of work-from-home occupations are nearly 50 ppt more likely to not have a college degree. Targeting policies into the lower tail of the distribution is both cheaper (lower incomes to replace) and more effective (lower resources initially available).

4 Employment during the epidemic

We now use the limited data available since the start of the epidemic for three purposes. First, we validate our measures by checking that MSAs with more employment in low work-from-home jobs—which we classified in Section 2 using pre-virus data—saw smaller increases in the rates at which individuals stayed at home. Second, we show that occupations with low work-from-home scores experienced larger employment losses. Third, we return to the characteristics of workers associated with low work-from-home jobs in Figure 3 and show that these workers experienced larger declines in employment.

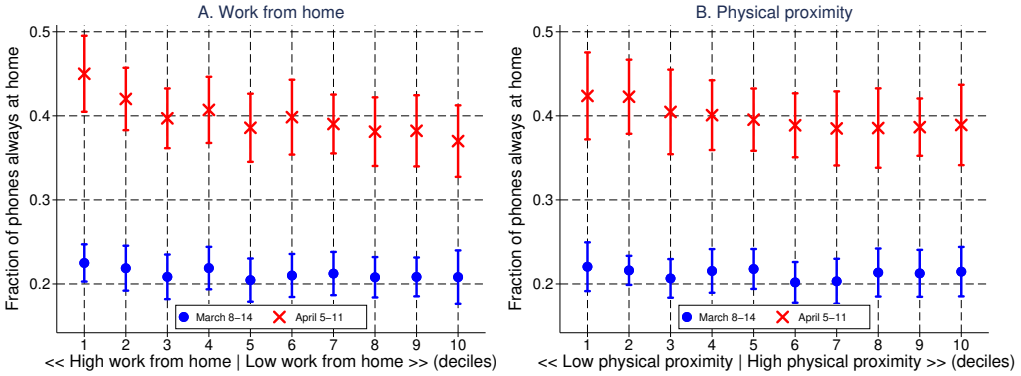


Figure 4: High work-from-home and Low physical-proximity MSAs had higher stay-at-home rates

Notes This figure compares the fraction of phones staying-at-home across MSAs. MSAs are grouped into population-weighted deciles of the work-from-home measure (Panel A), and physical-proximity measure (Panel B). Two weeks of data are plotted. The MSA level work-from-home measure is the employment-weighted average of occupation-level work-from-home measures. The MSA level phones-always-at-home measure time-aggregates over the week and geographically-aggregates over Census tracts, all phone-days in which a phone remains in a 153m×153m space, and divides by total phone-days. For example, in MSAs that are in the lowest (highest) decile of employment in work-from-home jobs, the fraction of phones at home increased by 24 ppt = 46%-22% (0.16 ppt = 37% - 21%).

4.1 Cell-phone behavior by work-from-home at the MSA level

Does the work-from-home measure formulated using pre-virus features of occupations actually relate to outcomes measured after the outbreak began? One way to assess this is to compute a work-from-home measure at the level of a geographic unit, and then study changes in stay-at-home behavior across these units. To do this we study MSAs and use SafeGraph cellphone data. SafeGraph provides daily data on the total number of cellphones that end the day in a Census block group. They also provide the number of these cellphones that stay-at-home, as measured by cellphones that do not leave a small area.¹⁷ We do not have occupation employment data at the Census block group level, but the OES provides employment at the MSA level. We therefore aggregate phone-days and phone-days-at-home within each MSA to construct a weekly measure of the fraction of phones always at home. We construct a work-from-home measure at the MSA level by using OES data on MSA-occupation employment to take the employment-weighted average of \overline{LWFH}_j at the MSA level.

Figure 4 is supportive of the measures being relevant for understanding changes in employment in the epidemic. We bin MSAs into deciles based on this measure, and then compute the average fraction of phones always-at-home. In the second week of March there is no correlation across MSAs between phone behavior and the work-from-home measure. Between the second week in March and

¹⁷The variable we use is *completely_home_device_count*, which gives the number of devices that do not leave an approximately 153m×153m block.

the second week in April, the average fraction of phones that are always at home jumps by around 15 ppt. However, low work-from-home and high physical-proximity MSAs see smaller increases in phones that stay-at-home, as workers are less easily able to relocate from their usual place of employment to their usual place of residence.

4.2 Employment losses by occupation

Excess employment losses from February to March of 2020 show a clear pattern: occupations with low WFH scores had relatively larger declines in employment than occupations with high WFH scores. Jobs that are more easily done at home are more likely to remain intact through the economic shutdown. We show that this is the case using CPS (2-digit) occupational employment data covering January 2010 to March 2020, the latest available CPS data. To account for seasonal factors in occupation employment changes, we construct excess employment losses by taking the log change in employment from February to March 2020 and subtracting the average February to March change in employment. Figure 5A compares the relationship between this excess decline in employment against \overline{LWFH}_j , showing that low work-from-home jobs experienced larger excess employment losses.

An important exception to this relationship, as expected, are those jobs deemed essential by public policy.¹⁸ These are unlikely to have employment losses that correlate with whether the job can be done from home or not. For example, front line medical workers have low WFH measures (healthcare supplemental workers have a WFH index of around 0.2), but because they have been declared essential they can continue to work. We do not have information on which occupations are deemed essential, so instead we use industry data created by Tomer and Kane (2020), who categorize certain 4-digit NAICS industries as “essential”.¹⁹ For each 2 digit occupation we use the 2018 OES to calculate the share of employment in essential industries, and categorize an occupation as essential if more than 75 percent of employment is in an essential industry. Occupations that meet this criterion are depicted in red in Figure 5. Among these occupations there is no significant relationship between the WFH measure and employment growth.²⁰

4.3 Employment losses by worker characteristics

As a final exercise, we study how excess employment losses vary across the worker characteristics considered in Section 3. For each group of workers we compute the total employment change over

¹⁸See <https://www.cisa.gov/publication/guidance-essential-critical-infrastructure-workforce>

¹⁹Tomer and Kane (2020) use job descriptions from the government statement which announced guidelines for categorizing essential jobs. The text for this document can be found at <https://www.cisa.gov/publication/guidance-essential-critical-infrastructure-workforce>.

²⁰The metric we use to categorize occupations as essential is able to pick up certain obvious 2-digit occupations such as healthcare technicians and healthcare support. However, some occupations seem to be left out despite having numerous mentions in the aforementioned government text. For example, the word *construction* is mentioned thirty three times; the word *legal* is mentioned only once. In other work we intend to incorporate a more direct mapping between the government text and the SOC occupation codes.




Figure 5: Employment declines by occupations and worker characteristics (February-March 2020)

Notes Both figures use employment data from the CPS. **Panel A.** This figure plots employment changes by 2-digit OCC occupation against \overline{LWFH}_j . Employment change is the Feb-Mar log change in employment in 2020 of each occupation, relative to the average Feb-Mar log change in employment over 2010-2019 for the occupation. Occupations marked with red diamonds are defined as “essential” using the grouping from [Tomer and Kane \(2020\)](#). Fitted values are from an employment-weighted linear regression estimated on non-essential occupations, and gives 95 percent confidence intervals for the conditional expectation of the dependent variable. **Panel B.** This figure plots employment changes by type of worker. The variables on the *y*-axes are used to split workers into two groups: those in the group given by the label (‘Yes’, marked with a green circle), and those outside that group (‘No’, marked with a red cross). We plot the log change in employment across February and March in 2020, adjusted by subtracting the average February-March change in employment for that group over 2010 to 2019. For the whole sample, we obtain a total decline in employment of -2.3 log points (black dashed line). (*) For the last two cases the sample is restricted to those reporting hours—in the case of *Part time employed*—and hours+earnings—in the case of *Below median wage*. In these subsamples, the average employment change is not -2.3 log points.

February-March 2020, and subtract off the mean total employment change for February-March for 2010-2019. We focus on employment rather than unemployment due to issues associated with the labeling of workers as unemployed versus out of the labor force. Figure 5 shows the results.²¹

Once again, a clear pattern emerges: those groups of individuals who have a higher employment share in low WFH occupations (as identified using the methodology of Section 2) experienced, on average, more severe employment outcomes in March 2020 relative to those in occupations with high work-from-home capability. The characteristics with the largest differential in employment outcomes between groups are citizenship, nativity, marital status, education, and age. Interestingly, while gender was an important margin which predicted the likelihood of being in low work-from-home jobs, the employment changes for February-March 2020 are not as extreme as other observable groups, which may be related to the issue of “essential” jobs discussed earlier.

²¹We check that the total employment losses that we construct using survey weights lines up with total employment losses reported by the BLS. We obtain a value of -1.82 percent. The official value from the BLS is -1.90 percent, which we compute as $\log(158,759/155,772)$ from Table 6 of the following: [BLS ‘Employment Situation’ report - March, 2020](#).

5 Conclusion

We show that workers systematically differ across the types of occupations that were most likely to be hit by the public policies around social distancing required to stop the spread of the Coronavirus. Workers in occupations that are *most likely* to be affected—those with a low score in the work-from-home measure, or a high score in the O*NET measure of personal-proximity—are predominantly characterized by traits associated with the more economically vulnerable in the US economy. These workers are disproportionately less educated, have limited healthcare, are toward the bottom of the income distribution, and have low levels of liquid assets. We showed that this was a useful way of understanding job losses following the start of the outbreak in 2020.

Given the occupation-level indicators we have constructed and made available with this paper, our measures can be used to capture geographic or group level exposure to social distancing policies. Moreover, our simple approach can be extended to individual economic indicators in any microdata that records occupation. An obvious example would be individual level data on wealth and asset portfolios beyond what is available in the PSID, such as the microdata underlying the *Survey of Consumer Finance* (SCF) which is available to researchers at the Federal Reserve Board.

References

- Chodorow-Reich, Gabriel**, “The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008-9 Financial Crisis,” *The Quarterly Journal of Economics*, 2014, 129 (1), 1–59.
- Dingel, Jonathan and Brent Neiman**, “How Many Jobs Can Be Done At Home?,” BFI White Paper, Becker Friedman Institute 2020.
- Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante**, “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967-2006,” *Review of Economic Dynamics*, January 2010, 13 (1), 15–51.
- Kaplan, Greg and Giovanni L. Violante**, “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, July 2014, 82 (4), 1199–1239.
- , – , and **Justin Weidner**, “The Wealthy Hand-to-Mouth,” *Brookings Papers on Economic Activity*, 2014, 45 (1 (Spring), 77–153.
- Leibovici, Fernando, Ana Maria Santacruce, and Matthew Famiglietti**, “Social Distancing and Contact-Intensive Occupations,” Technical Report, St. Louis Federal Reserve Bank - On the Economy Blog March 2020.
- Tomer, Adie and Joseph W. Kane**, “How to protect essential workers during COVID-19,” Brookings Report, Brookings Institute March 2020.

APPENDIX

A Additional figures and tables




Figure A1: Comparing different groups of occupations on the Work-from-home measure

Notes This figure extends Figure 3. The blue markers replicate Figure 3. In constructing the estimates plotted in green, we set $LWFH_j = 0$ for the second quartile of our continuous measure \bar{z}_j , and $LWFH_j = 1$ for the third quartile of \bar{z}_j . In constructing the estimates plotted in red, we set $LWFH_j = 0$ for the first quartile of \bar{z}_j , and $LWFH_j = 1$ for the fourth quartile of \bar{z}_j .

On Covid-19: New implications of job task requirements and spouse's occupational sorting¹

Warn N. Lekfuangfu,² Suphanit Piyapromdee,³
Ponpoje Porapakkarm⁴ and Nada Wasi⁵

Date submitted: 24 April 2020; Date accepted: 25 April 2020

*The Covid-19 pandemic has disrupted working life in many ways, the negative consequences of which may be distributed unevenly under lockdown regulations. In this paper, we construct a new set of pandemic-related indices from the Occupational Information Network (O*NET) using factor analysis. The indices capture two key dimensions of job task requirements: (i) the extent to which jobs can be adaptable to work from home; and (ii) the degree of infection risk at workplace. The interaction of these two dimensions help identify which groups of workers are more vulnerable to income losses, and which groups of occupations pose more risk to public health. This information is crucial for both designing appropriate supporting programs and finding a strategy to reopen the economy while controlling the spread of the virus. In our application, we map the indices to the labor force survey of a developing country, Thailand, to analyze these new labor market risks. We document differences in job characteristics across income groups, at both individual and household levels. First, low income individuals tend to work in occupations that require less physical interaction (lower risk of infection) but are less adaptable to work from home (higher risk of income/job loss) than high income people. Second, the positive*

1 We thank Peter Spittal for helpful comments.

2 Assistant Professor, Department of Economics, Universidad Carlos III de Madrid.

3 Assistant Professor, Department of Economics, University College London.

4 Associate Professor, National Graduate Institute for Policy Studies (GRIPS).

5 Associate Research Director, Puey Ungphakorn Institute for Economic Research, Bank of Thailand

occupational sorting among low-income couples amplifies these differences at the household level. Consequently, low-income families tend to face a disproportionately larger risk of income/job loss from lockdown measures. In addition, the different exposure to infection and income risks between income groups can play an important role in shaping up the timing and optimal strategies to unlock the economy.

1. Introduction

Unlike previous economic shocks, Covid-19 has disrupted labor markets around the world along two new dimensions. First, workers in certain jobs are at higher risk of infection and transmission, particularly those working in close physical proximity to other people. Second, workers whose jobs are not adaptable to work from home may have a higher risk of income loss due to drastic measures (e.g. sectoral lockdowns and social distancing) to curb the spread of virus. Identifying the tradeoffs between the risks of economic loss and public health across occupations is essential to understanding the potentially heterogeneous impact across workers of Covid-19 and policies designed to contain it. At the household level, such risks may be intensified if both spouses sort into similar occupations, and so face common shocks. Socially desired exit strategies require a substantial balance between pandemic containment and economic burdens – both of which may involve rather different sets of stakeholders.

In this paper, we exploit the information on job task requirements of each occupation from the Occupational Information Network (O*NET) to construct a set of new pandemic-related indices using factor analysis. Specifically, these indices measure (i) the degree of job flexibility in terms of work location (due to job reliance on machinery or specific location; and adoption of ICT into task performance), and (ii) the extent to which jobs require the worker to perform tasks in close physical proximity with others. We show that these statistically-constructed indices can represent two important risks posed by the Covid-19 pandemic on workers: the risk of earnings losses when a worker is away from their regular workplace, and the risk of contracting or spreading the virus at the workplace.

Further, the interactions of these indices along the earnings distribution can be informative for designing programs to support different groups of affected workers as well as strategies to reopen the economy. For example, workers who cannot adapt to work from home may require more support than those who can, especially if work location flexibility is negatively correlated with earnings. On reopening the economy, the debate is around how to minimize the economic losses while controlling the spread of the virus. Our analysis suggests that workers in jobs which are not adaptable to work from home, but do not require frequent physical contact with others, should be allowed to return to their workplaces first. On the other hand, those who usually work in close physical proximity to others, but whose jobs are well-suited to work from home, may be the last to return to normalcy.

As an application, we focus our analysis of the impact of the pandemic on a developing economy. With relatively low social safety nets and large shares of workers in the informal sector with weak labor protections, workers in developing countries stand a higher risk of earnings loss in the presence of a global economic and public health crisis such as the Covid-19 pandemic. To investigate such potential impact, our analysis focuses on Thailand.¹ We map the latest release of Thailand's Labor Force Survey (2019) with the O*NET-derived indices, and evaluate the labor market risks arising from the Covid-19 crisis at both individual and household levels.² In developing countries, risk-

¹ Despite the relatively few Covid-19 cases at present, Thailand was one of the countries with the highest number of Covid-19 cases outside China at the onset of the crisis (January 2020), owing to the largest number of daily direct passenger flights from Wuhan. By mid-March of 2020, the Thai government declared a state of emergency – with the implementation of strict sectoral lockdown regulations, and social distancing practice.

² The structure variable definitions and survey conduct in the Thai Labor Force Survey are analogous to the European Union's Labor Force Survey (EU LFS) and the US's Current Population Survey (CPS). We use the third quarter data because it includes seasonal workers. Workers included in the LFS work in both formal and

sharing within households plays a central role in absorbing shocks (e.g., Chiappori et al. 2014, Samphantharak and Townsend 2018). Therefore, if Covid-19 exposes both primary earners in a household to common shocks, the impact on their livelihoods can be severe. Insights from our analysis on Thailand are highly relevant for other countries with similar labor market structures – specifically, a relatively large share of self-employment and low social safety net.

First, we document that there are noticeable differences in occupational indices among individuals from different income groups. Specifically, people with lower earnings tend to face a lower infection risk at the workplace, but a higher risk of income or job losses due to the difficulty in adjusting their working arrangement following a sectoral lockdown. Second, the occupational sorting within married couples reinforces these differences at the household level. Married couples from the lower end of earnings distribution are much more likely to sort into occupations with similar indices, and are highly concentrated in jobs not adaptable to work from home. In effect, earnings within low-income households are highly correlated, which can lead to large losses in household income during the lockdown regulations. This suggests that means-tested emergency relief programs would be more suitable than universal support programs in terms of targeting those working in most adversely affected occupations.

This paper is closely related to works studying the labor market consequences of lockdown measures using occupational characteristics. Hicks (2020) recently uses the O*NET data on the degree of physical proximity to assess which occupations are more likely to be affected. Focusing on work flexibility characteristics, Dingle and Neiman (2020) and del Rio-Chanona et al. (2020) manually classify occupations into a binary variable whether they can be performed at home.³ Our main contribution is to show that (i) physical proximity, (ii) work-location flexibility and (iii) their interactions are crucial for impact evaluation and policy design in response to the pandemic. Our work is directly complementary to Adams-Prassl et al. (2020) which provides evidence from real-time surveys that workers with limited work arrangement are highly exposed to less favorable job outlooks. Additionally, we complement the assortative mating literature by showing that the labor market risk induced by Covid-19 at the household level can be mitigated or amplified depending on the occupational sorting pattern between husbands and wives.

The paper proceeds as follows: Section 2 describes how the indices are constructed using the O*NET. Section 3 applies the indices to evaluate labor market risks at individual and household levels in the Thai context. Section 4 discusses policy recommendations and Section 5 concludes.

2. Methodology

We select 24 task-based occupational variables from the O*NET data on ‘Work Context’ and ‘Work Activities’ to capture (i) the extent to which a job can be done at home, and (ii) whether a job requires working in close proximity with other people. The latter group of characteristics is particularly important for policy decision-making during the pandemic as the virus can easily be transmitted from person to person. (See the Appendix for the list of the selected O*NET variables.)

informal employment (defined by social security and health insurance status), as well as those in agricultural sector.

³ Other papers using task characteristics to classify occupations to evaluate structural changes of labor markets include seminal work by Autor and Dorn (2013) and Blinder (2009).

To reduce the dimensionality of the O*NET variables, we perform an exploratory factor analysis with rotation method to establish a factor retention criterion. We impose oblique rotation of factor loadings to allow for correlation between the factors (Heckman et al., 2013). We retain three factors with eigenvalues greater than 2, following the criteria outlined by Gorsuch (1988).⁴ Table A1 in the Appendix presents the factor loadings on the predicted factors. A larger factor loading (in absolute terms) reflects higher correlation between the selected O*NET variable and the factors. The factors are standardized to have mean zero and standard deviation one.

The first factor encapsulates tasks related to repairing, maintaining, or inspecting equipment, structure or materials and operating vehicles or mechanized devices. Thus, we interpret this factor as a measure of both machine and location dependence of jobs. The second factor captures tasks that frequently utilise ICT - for example interacting with computer, analyzing data or processing information. The last factor captures whether the job often requires workers to perform tasks in close physical proximity to other people or to assist or care for others. For conciseness, in the rest of the paper, we will refer to these factors as indices for ‘machine-dependent’, ‘ICT-enabled’ and ‘physical proximity’, respectively.

We compute the three factor indices for over 900 detailed six-digit occupations (based on the US SOC 2010). We present a selected list of occupations with the highest and lowest scores in each factor in the Appendix. Note that the partial correlations of machine-dependent and ICT-enabled; machine-dependent and physical proximity; and ICT-enabled and physical proximity among the occupation list in the O*NET database are -0.40, 0.05, and 0.16 respectively. Small and statistically insignificant correlations of machine-dependent and physical proximity of occupations suggest that a lockdown restriction in response to the pandemic crisis may involve a trade-off along multiple dimensions, e.g. saving jobs versus preventing infection. The effects of the Covid-19 shock on jobs are therefore likely to be quite different from other economic shocks in past recessions.

Table 1 summarizes the average indices of the three factors derived from the O*NET by the broad occupation groups in columns 2, 3 and 5. While machine-dependent and ICT-enabled are separate factors, the ease of shifting work location from ‘office’ to home are highly depended on both factors in opposing directions.⁵ To ease our analysis, we also report an equally-weighted average of the scores of machine-dependent (reversed) and ICT-enabled factors in column 4, and refer to the additional index as the score of overall work-location flexibility. Broadly speaking, managers and professionals have relatively high degrees of work-location flexibility. Service and sale workers have the highest average indices of physical proximity.

The last three columns compare occupational compositions of workers in Thailand, EU-27 and the U.S. While the occupational shares of EU-27 and the U.S. are similar, the shares of Thailand reflect a common pattern of a middle-income economy – relatively large agricultural and manufacturing sectors with a lower share of workers in the high skill service sector (e.g. managers, professionals, technicians, and associated professionals).

Our analysis draws attention to the interaction between the degree of work-location flexibility and close physical proximity. While a lack of work-location flexibility indicates the risk of income losses

⁴ Statistical criteria for factor retention include the Scree Test, Onatski’s Test and Horn’s Test.

⁵ For instance, a market research survey interviewer has a low index of machine-dependent, but because interviews were typically done face-to-face before the pandemic, this occupation is associated with a low score of ICT-enabled. Without ICT infrastructure, it is unlikely that these interviewers could easily perform their work from home.


due to the inability to work during a lockdown, the physical proximity factor emphasizes the risk of virus infection and transmission in performing such tasks. In the event of a pandemic, performing such tasks is seriously discouraged; therefore, jobs with a high value of physical proximity index may also be exposed to income losses.

Table 1. Average Score of Factors and Occupational Distribution

Occupational Groups (1 Digit)	Work Location-Flexibility Indices			Physical Proximity	Share of workers, 2019 (%)		
	i. Machine-Dependent (-)	ii. ICT-Enabled (+)	Average of [-2] & [3]		US	EU-27	Thailand
	[2]	[3]	[4]		[6]	[7]	[8]
1. Managers	-0.22	0.87	0.67	0.56	11.07	4.23	3.66
2. Professionals	-0.73	0.66	0.85	-0.01	22.65	19.38	5.41
3. Technicians and associate professionals	0.09	0.35	0.16	0.33	14.28	17.81	4.32
4. Clerical support workers	-0.74	-0.14	0.36	0.09	9.89	10.89	4.42
5. Service and sales workers	-0.16	-0.61	-0.28	0.79	17.89	16.69	20.06
6. Skilled agricultural, forestry, fishery workers	0.99	-0.92	-1.17	-0.86	0.17	0.95	31.50
7. Craft and related trades workers	0.73	-0.98	-1.05	-0.41	8.38	11.53	10.59
8. Plant and machine operators, assemblers	1	-1.23	-1.36	-0.59	5.76	8.54	9.40
9. Elementary occupations	0.54	-1.25	-1.09	-0.35	9.90	9.97	10.64

Note: The indices are standardized to have mean zero and standard deviation one. Source: EU-27 occupational shares come from EU Statistics, US shares come from Bureau of Labor Statistics, and Thailand shares come from Thai Labor Force Survey.

Figure 1: Occupation Classification



Note: Factors are standardized to have mean zero and standard deviation one.

Figure 1 depicts the interactions between these two indices. The vertical axis represents the degree of work-location flexibility, and the horizontal axis shows the degree of physical proximity. The left panel illustrates that the 968 occupations from O*NET are distributed across all the quadrants. The right panel shows a selected list of occupations in each quadrant. Workers with occupations in quadrant IV (bottom-right) are arguably the most vulnerable group with respect to both income losses and getting infected at workplaces because they have relatively low degree of work-location flexibility and high degree of physical proximity. Workers with occupations in quadrant III (bottom-left) are also limited in their work arrangements but have jobs with less physical contact and correspondingly lower infection risk at their workplaces. Those in quadrants I and II (top-right and top-left) have jobs which are more flexible. We discuss the policy implications in more detail in Section 4.

3. Evaluating the New Labor Market Risks

Our analysis focuses on measuring supply-side labor market risks associated with various measures to slow down the infection rates. These include closing businesses in some or most sectors and requiring non-essential workers to work from home. We focus our study on potential labor market risks at occupation level using our occupation classification.⁶ Our case study is based on Thailand. Despite the relatively low official number of infections, the country mobilized to slow down the outbreak of Covid-19 by imposing strict sectoral lockdowns and campaigning for social distancing in late March 2020.

In section 3.1, we analyze the potential risks at the individual level. In section 3.2, we extend our analysis to households. Incorporating the role of assortative marriage, this section assesses to what extent occupational sorting of spouses amplifies or attenuates the income risk among different types of households. Insights from our analysis on Thailand are highly relevant for other countries with similar labor market structures – specifically, a relatively large share of self-employment and low social safety net. In term of marriage patterns, Thailand has seen increasing assortative marriage over the past decades, a pattern common to many developed and developing countries (see Chiappori, 2017 for a review).

We map the indices to the Thai 2019 Labor Force Survey (LFS), a quarterly nationally representative sample. For each sampling household, detailed information from all members is collected. This includes demographic characteristics, marital status, employment status, work hours, occupations and sectors. While the complete information on occupation is available for all types of workers (wage or salary workers, self-employed and unpaid workers), earnings data were collected only for wage or salary workers. For individual analysis, we restrict the sample to workers aged between 15 and 65 years old. For household analysis, we further use the subsample of married couples.

3.1 Individual Heterogeneity

Table 2 reports the average indices across genders, age groups and education levels. On average, occupations held by older and lower educated workers tend to be more machine-dependent,

⁶ Because the actual lockdown sectors differ across countries, we do not explicitly incorporate the sectoral lockdown to analyze the labor market risks. See a companion analysis in Lekfuangfu et al. (2020), where we documented the differences in the lockdown sectors in Thailand and European countries.

less ICT-enabled, and have a lower degree of close physical proximity. On average, jobs held by Thai men are less flexible but require less physical contact than jobs held by Thai women. This is because a higher proportion of men work as assemblers or machine operators in factories and agricultural activities, while a higher proportion of women are in sales and services.

Table 2. Average Factors by Worker Characteristics


	Work-Location Flexibility Indices				Total number of workers (millions, %)
	i. Machine-Dependent (-)	ii. ICT-Enabled (+)	Average Index	Physical Proximity	
National	0.41	-0.61	-0.62	-0.25	37.3
Gender					
Male	0.60	-0.59	-0.73	-0.30	20.4 (55%)
Female	0.17	-0.63	-0.49	-0.19	17.1 (45%)
Age groups					
15-25	0.41	-0.75	-0.71	-0.22	4.3 (12%)
26-35	0.21	-0.51	-0.43	-0.09	8.3 (23%)
36-45	0.31	-0.53	-0.51	-0.18	8.7 (24%)
46-55	0.40	-0.61	-0.67	-0.39	8.8 (25%)
56-65	0.62	-0.72	-0.82	-0.46	5.5 (16%)
Education Levels					
Secondary or lower	0.53	-0.73	-0.77	-0.28	31.6 (85%)
College	-0.47	0.25	0.44	0.29	6.1 (15%)

Notes: The indices are standardized to have mean zero and standard deviation one.

To understand how the job task requirements are mapped into earnings, **Figure 2** plots the indices across earnings deciles. Workers with lower earnings work in occupations that are more machine-dependent and less ICT-enabled, making them less flexible to work remotely. Thus, lower earning workers are more exposed to the risk of losing income during the pandemic than higher earning workers. The degree of physical proximity, however, is reversed. Lower paid workers tend to be the laborers (e.g., fixing streets, construction site) and those who work in factories whose work naturally involves less close physical interaction.

Since the Covid-19 shock adversely affected people's health and income, it creates political tensions between people from different groups which can play an important role in shaping policy in response to the crisis. Glover, et al. (2020) emphasize the tension between people outside the labor market ("the old") and those participating in the labor market ("the young"). The old face a higher mortality risk of being infected but little (or no) earnings risk; the opposite is true for the young. Consequently, the old may prefer more drastic measures or delays to opening the economy. Our findings reveal an additional tension among workers in the labor market which has not previously been discussed. While those from a lower earnings bracket face a lower infection risk, they may endure a larger economic loss from having a lockdown imposed on them due to the difficulty in adjusting their work arrangements. Therefore, these workers would prefer an earlier removal of the lockdown. The opposite may be true for the high-income group.

Figure 2. Work Characteristics by Earnings Decile

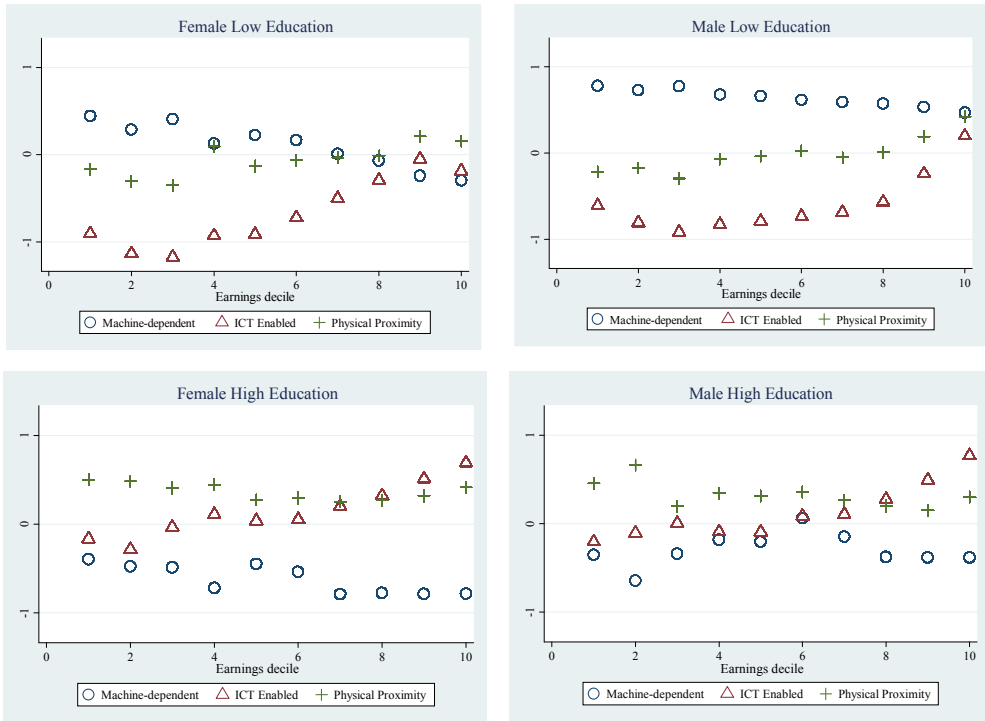


Notes: The figure shows average score in each factor index along the earnings distribution. The horizontal axis is the ranking position of individual wage earners on the earnings distribution of all wage earners observed in the LFS 2019, quarter 3.

Figure 3 plots the indices across earnings decile conditional on gender interacted with education level. There are stark differences between occupational characteristics of workers with and without a college degree along the earnings distribution. The average indices shown in Table 2 capture the characteristics of workers with at most secondary education who have a much larger share in the workforce. However, conditional on having a college degree or higher, the differences between occupations of males and females are modest. Nevertheless, the pandemic-induced risk of earnings losses at the household level depends on *the composition of the household*. That is, households with more *dissimilarity of occupations* with respect to the flexibility to work from home and physical proximity are in a better position to smooth the negative income shock. On the other hand, households with both primary earners working in low physical proximity and flexible jobs would be best off, while households with both partners having limited work-location flexibility would face much harsher economic implications.⁷

⁷ Moreover, other structures of households, such as whether there are young children or not, could be vital. Given school and day care closures, mothers are more likely to be affected. Being able to work from home might alleviate the impact (see Alon et al. (2020) for a discussion).

Figure 3. Work Characteristics by Gender and Education



Notes: We use common earnings decile as in Figure 2.

3.2 Household’s Correlated Risks

The impact of the pandemic on individuals’ earnings discussed in the previous section can be either mitigated or magnified at the household level through occupational sorting within households. We examine this point by focusing our analysis on households in which both spouses work. To shed light on the pattern of occupational sorting, we report the correlation of each index between the husband’s and wife’s job separately for different types of households.

To account for a large share of Thailand’s informal sector (33% self-employed, 17% unpaid workers and 47% paid employees), we classify working married couples in our sample into four types as the following:

- Type A: both work as employees
- Type B: one as employee, another as self-employed
- Type C: both spouses are self-employed
- Type D: both work, and at least one works as unpaid family worker ⁸

⁸ Unpaid family workers are people working without actual pay in an enterprise or farm owned by a family member.

Table 3 displays the spousal correlations of each pandemic-related index (machine-dependent, ICT-enabled, physical proximity). For households of types A and D (76.4% of total), the correlations between indices of jobs held by married couples of all three factors are highly positive. For households of type D, the spousal correlations are close to one – suggesting that most unpaid workers tend to work in the same or similar jobs as their spouses. Thus, for type D, the occupational impact at the household level would be similar to that at the individual level. For type B (one spouse is self-employed), the negative correlations of machine-dependent and physical proximity factors indicate that these households may have a higher degree of risk diversification through less assortative occupational choices.

Table 3. Married couple correlations by types of employment

Types of households	Total married workers in millions (%)	Machine-dependent	ICT Enabled	Physical Proximity
A. Both as employees	4.1 (30)	0.39	0.56	0.44
B. One employee and one self-employed	1.9 (15.3)	-0.05	0.11	-0.03
C. Both as self-employed	1.1 (8.3)	0.19	0.08	0.15
D. One or both as unpaid workers	6.2 (46.4)	0.9	0.92	0.97
All households	13.3 (100)	0.51	0.61	0.66

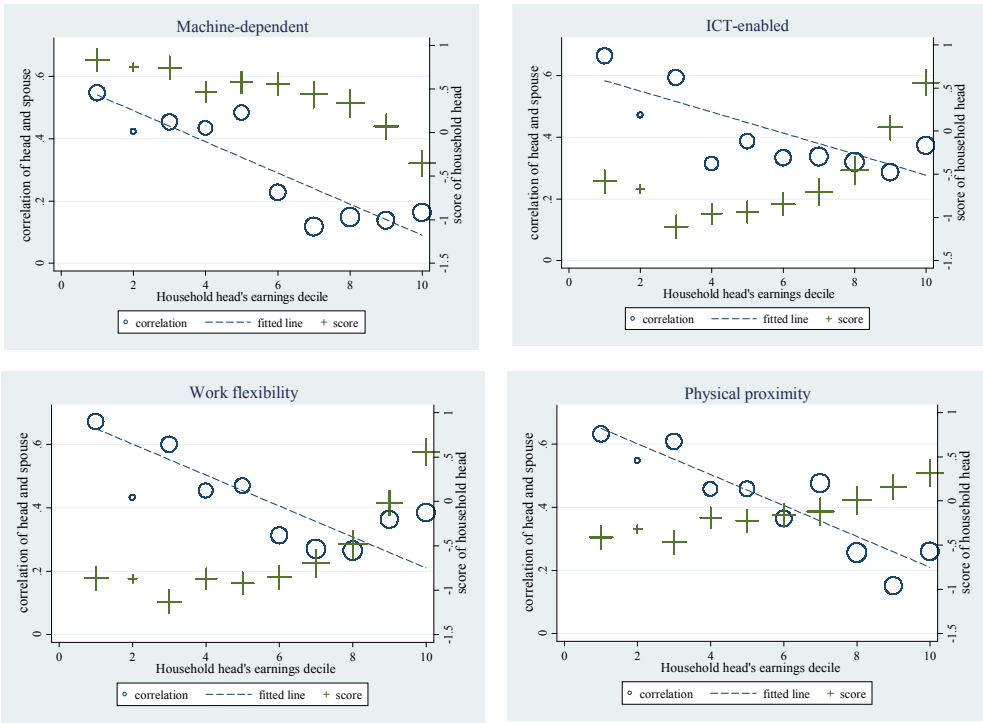
Notes: The correlations are weighted by the sum of individual survey weight of the head and of the spouse.

Whether the positive spousal correlations for the job flexibility or physical proximity factors reflect the scale of labor market risk depends predominantly on the magnitude of these indices. For instance, a household is better insulated from a negative shock from lockdown measures when at least one spouse is in an occupation with a low degree of machine dependent and/or high degree of ICT-enabled - which implies a higher probability of being able to work from home. In contrast, the Covid-19 crisis could cause larger losses in household's income if both spouses lose their jobs because they cannot work from home. Moreover, the impact can substantially worsen income and consumption inequalities if such positive occupational sorting (into jobs with limited locational flexibility) is more prevalent among poor households.

In what follows, we investigate the pattern of spousal correlations along the earnings distribution. We focus on households of type A for which earnings of both spouses are observed in the data. **Figure 4** depicts the spousal correlations of their occupational factor (on the left vertical axis), and the average score of household head for a given factor (the right vertical axis).⁹ It shows that the spousal correlations are strongly positive particularly at the lower-end of the earnings distribution. This suggests that married couples from low-income households work in occupations with common levels of machine-dependent, ICT-enabled and physical proximity. Further, the size of the positive correlation decreases along the earnings decile, in particular for machine-dependent factor. These plots present compelling evidence that labor market risks due to the Covid-19 are heterogeneous across households – and that those at the bottom-end of the earnings distribution are more at risk to the pandemic shock than others.

⁹ We define household head as the highest earner of the couple. The horizontal axis in Figure 4 represents earnings decile, calculated from all wage workers aged 15-65 years old.

Figure 4. Type A household: both are employees



Notes: The figures show spousal correlations in each factor index. The horizontal axis is the ranking position of household heads on the earnings distribution of all wage earning individuals observed in the LFS 2019 (quarter 3). On the horizontal axis, we use the common earnings decile as in Figures 2 and 3.

4. Policy Implications

In response to the Covid-19 pandemic, affected countries around the world have introduced various forms of supports to aid their citizens, including emergency cash transfer programs, social assistance, in-kind food and utility and financial obligation waivers. The cash transfer programs appear to be most common with some countries launching a one-off universal transfer whereas others used a means-tested cash transfer, i.e., the cash amount is conditional on household's income or people working in certain occupations (Gentilini et al., 2020).

In Section 3, we show that the degree of potential impacts of Covid-19 on workers depend on the two new dimensions of their job characteristics (the degrees of work-location flexibility and close physical proximity), and these impacts can be intensified by occupational sorting in marriage. Our findings can be used to guide efficient supporting schemes for different targeted groups of workers, and strategies for reopening the economy. At the time of writing, some countries announced that they have been able to slow down the virus outbreak; thus, the recent debate has shifted towards how to open up the economy without jeopardizing the public health systems.

Table 4 demonstrates examples of occupations in the four quadrants deriving from the cross-dimensionality between work-location flexibility and physical proximity (as seen in **Figure 1**). Workers with occupations in quadrant IV (bottom-right) are the most vulnerable group. Due to the high degree of close physical proximity, these jobs have been the first to be restricted, and potentially will be the last to return to normalcy. Unlike workers in quadrants I and II (top-right and top-left), workers in quadrant III (bottom-left) could ‘produce’ only if they are allowed to return to their workplaces. In fact, those in quadrant II (top-left) may experience relatively mild impacts from the lockdown measures since their jobs are more flexible and do not require frequent physical contact with others.

Table 4: Selected occupations in four impact groups

Work Location Flexibility	Physical Proximity	
	Low	High
High	Quadrant II: Mild impact	Quadrant I: Medium impact
	Sociologist	Human resources manager
	Programmer	Fitness manager
	Website developer	Business strategy manager
	Economist	Information coordinator
	Financial advisor	Head hunter
Low	Quadrant III: Medium impact	Quadrant IV: Severe impact
	Legal councillor	Secondary school teacher
	Garment factory worker	Cleaner
	Metal worker	Restaurant server
	Planter, Grower	Travel organiser
	Construction worker	preschool teacher
	Machine controller	Tour guide
Production worker	Dentist	
Painter and polisher	Veterinarian	

Notes: The impact level is derived based on the interactions of work location flexibility and physical proximity.

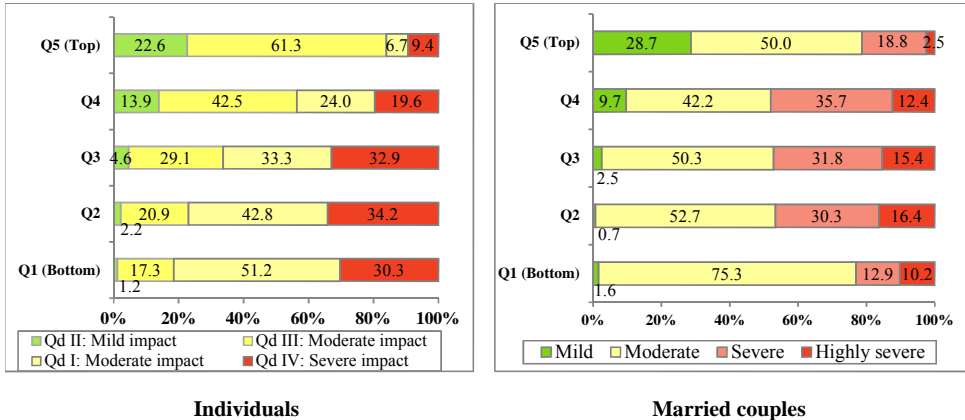
The composition of workers across the four quantiles is not equally distributed on the earnings distribution. As can be seen in the left panel of **Figure 5**, there is a substantially larger fraction of workers with occupations in quadrant IV (red) — the severely impacted group— at the lower-end of the earnings distribution. In contrast, workers at the top quintile have the largest share of jobs considered to be mildly affected by the pandemic crisis (based on our occupational classifications) in the short run. This means that without adequate government intervention to support income or employment for the poor, the adverse impact of Covid-19 could worsen income inequality.

The findings suggest that mitigation interventions should be targeted based on job characteristics when possible.¹⁰ For example, ICT-related support would assist those in quadrant I (top-right) to

¹⁰ In some sectors such as finance, businesses can continue without all workers being able to work from home since core and non-core activities are divisible e.g. face-to-face customer relation activities can be put on hold while main banking activities can run remotely. However, in some other sectors such as hotels and restaurants,

maintain their work activities. In contrast, the measure would be less effective for those in quadrant III (bottom-left) since their jobs are required to be performed at specific locations. In effect, potential schemes providing substitutions for income losses would be more suited for workers holding jobs in quadrants III and IV. Given that the Covid-19 pandemic is likely to disproportionately affect the low earners, the introduction of means-tested relief programs targeting those working in most adversely affected occupations, rather than a universal program, would be socially desirable.

Figure 5. Fraction of the derived impact levels by earnings distribution



Notes: The left panel shows the proportion of wage earning individuals by the derived impact level (based on the interactions of work location flexibility and physical proximity) for each earnings quintile. The right panel shows the proportion of married couples in wage earning employment by the derived impact level. ‘Mild’ is for households with both spouses’ jobs in quadrant II; ‘moderate’ when no spouse’ jobs in quadrant IV; ‘severe’ when one spouse’s job in quadrant IV; ‘highly severe’ when both spouses’ jobs in quadrant IV. For both panels, the ranking position is based on the earnings distribution of all wage earning individuals observed in the LFS 2019 (quarter 3).

The above argument is reinforced when taking into account the high degree of occupational sorting among married couples at the bottom-end of the earnings distribution (as discussed in Section 3.2). The right panel in **Figure 5** shows the fractions of at-work married couples (household-level) according to the derived severity of the pandemic impact on their jobs. In this case, we define impact as ‘mild’ (green) for households with both spouses in jobs of quadrant II, ‘moderate’ (yellow) for households without any spouse’s job in quadrant IV; ‘severe’ (orange) for households with one spouse’s job in quadrant IV, and ‘highly severe’ (red) when both spouses’ jobs are in quadrant IV.

The top earnings quintile has noticeably the largest fraction of households classified as mildly impacted (green) and the smallest fraction of households classified as ‘highly severe’ (red). In contrast, married couples who both work in occupations in quadrant IV are of the highest fraction in the middle quintile groups. A large fraction of the bottom quintile couples are classified as moderate impact because many low wage occupations are based in factories which require less physical interaction. Overall, our findings suggest that suitable relief schemes, for instance income transfers,

this kind of division is unlikely to be possible. Although their managers and sales may be able to work from home, most service staff work in close physical proximity to others.

should be means-tested with criteria based on specific occupational characteristics as well as joint household earnings.

As for reopening the economy, other things being equal (for instance, health, age of household members, the infection rate and healthcare capacity in the area), our results and the application of the occupation indices, discussed earlier, indicate that the highest priority to relax lockdown regulations should be given to workers in occupations in quadrant III (bottom-left). Without returning to their workplace, these workers face a high risk of income losses. Additionally, allowing them to return to work may involve minimal infection transmission risk since their works require limited physical contact with others.

5. Conclusion

The Covid-19 pandemic has posed new types of risks on workers around the world. Given the rapid transmission from person to person of the virus, drastic measures such as lockdowns and social distancing have been imposed to control the spread of infection. Despite differences in the scope of sectoral lockdowns across countries, these measures undoubtedly come with sizable costs to the economy.

The direct effect of a lockdown can have different impacts on workers with different job characteristics. To understand such heterogeneous impacts, we use a factor analysis to construct a set of occupational indices that are general but relevant to study the impacts of the Covid-19 pandemic. These indices feature two key dimensions of job task requirements: the degrees of work-location flexibility and working in close physical proximity to others. The former captures the risk of the worker's income loss, and the latter captures the infection risk posed to the worker and the public. We show that occupations in the O*NET are broadly distributed over these two dimensions.

Using the data from Thailand, we document that low earners tend to work in occupations that are less adaptable to work from home, but their jobs usually do not require frequent physical interaction with others. Furthermore, we show evidence that spouses in low-income households sort into jobs with similar characteristics. This occupational sorting amplifies income risk at the household level during the lockdown period. Our findings offer evidence supporting the use of *means-testing* in assistance programs to ease the burden of those immediately affected by the drastic measures. Our indices can also be useful when designing a policy to reopen the economy with the goal of minimizing the income and job losses while controlling the spread of the virus.

Finally, our study takes the first step to analyzing the impact of the pandemic from the labor supply side. Fruitful avenues for future research include (i) incorporating the labor demand side (incorporating, for example, the decline in consumption and supply-chain effects); (ii) allowing for substitutions – cases in which workers switch to jobs requiring similar skills or, over the longer term, adjust their skills; and (iii) using our constructed indices as supplementary classifications of jobs in order to further track the labor market adjustments as a result of the pandemic in the long run.

6. References

- Adams-Prassl, A., Boneva, T., Golin, M., and Rauh, C. (2020) Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys, *mimeo*
- Alon, T. M., Doepke, M., Olmstead-Rumsey, J., and Tertilt, M. (2020). The Impact of COVID-19 on Gender Equality. *CEPR Covid Economics: Vetted and Real-Time Papers*, 4.
- Autor, D.H. and Dorn, D. (2013). ‘The growth of low skill service jobs and the polarization of the US labor market’, *American Economic Review*, 5(100), 1553–97.
- Blinder, A.S. (2009). How many US jobs might be offshorable? *World Economics*, 2(10), 41–78.
- Chiappori, P. A. (2017). *Matching with transfers: The economics of love and marriage*. Princeton University Press.
- Chiappori, P. A., Samphantharak, K., Schulhofer-Wohl, S., and Townsend, R. M. (2014). Heterogeneity and risk sharing in village economies. *Quantitative economics*, 5(1), 1-27.
- del Rio-Chanona, R. M., Mealy, P., Pichler, A., Lafond, F., & Farmer, D. (2020). Supply and demand shocks in the COVID-19 pandemic: An industry and occupation perspective. *CEPR Covid Economics: Vetted and Real-Time Papers*, 6.
- Dingel, J. and Neiman, B. (2020), How many jobs can be done at home? *CEPR Covid Economics: Vetted and Real-Time Papers*, 1.
- Gentilini, U., Almenfi, M., Orton, I, and Dale, P. (2020). Social Protection and Jobs Responses to COVID-19: A Real-Time Review of Country Measures. *Live Document. World Bank, Washington, DC*
- Glover, A., Heathcote, J., Krueger, and D., Rios, V. (2020), Health vs Wealth: On the distributional effects of controlling a pandemic, *CEPR Covid Economics: Vetted and Real-Time Papers*, 6.
- Gorsuch, R. L. (1988). Exploratory factor analysis. In *Handbook of multivariate experimental psychology* (pp. 231-258). Springer, Boston, MA.
- Heckman, J., Pinto, R., & Savelyev, P. (2013). Understanding the mechanisms through which an influential early childhood program boosted adult outcomes. *American Economic Review*, 103(6), 2052-86.
- Hicks, M. J., Faulk, D., & Devaraj, S. (2020). Occupational Exposure to Social Distancing: A Preliminary Analysis using O* NET Data. *Becker Friedman Institute white paper*
- Lekfuangfu, N.W., Piyapromdee, S., Porapakkarm, P., Wasi, N. (2020). ‘Sectoral Lockdowns: Labor Supply Analysis in Thailand’, *aBRIDGED article, Puey Ungphakorn Institute for Economic Research*.
- Samphantharak, K., & Townsend, R. M. (2018). Risk and return in village economies. *American Economic Journal: Microeconomics*, 10(1), 1-40.

Appendix: Construction of the indices from O*NET variables

We select a list of *Work Activities* (14 items) and *Work Context* (14) from the O*NET work characteristics as described in Table A1.

Table A1. Selected O*NET variables

Work Activities variables: we take the Importance Score of each activity (measured on the 0-100 scale).	Work Context variables: we use the original scale (0,25,75,100) ¹¹
<ul style="list-style-type: none"> Assisting and Caring for Others Performing for or Working Directly with the Public Repairing and Maintaining Electronic Equipment Repairing and Maintaining Mechanical Equipment Operating Vehicles, Mechanized Devices, or Equipment Performing General Physical Activities Interacting With Computers Handling and Moving Objects Documenting or Recording Information Controlling Machines and Processes Thinking Creatively Processing Information Analyzing Data or Information Inspecting Equipment, Structures, or Material 	<ul style="list-style-type: none"> Structured versus Unstructured Work Pace Determined by Speed of Equipment Freedom to Make Decisions Spend Time Walking and Running Physical Proximity Outdoors, Under Cover Outdoors, Exposed to Weather Telephone Work With Work Group or Team Public Speaking Responsible for Others' Health and Safety Electronic Mail Face-to-Face Discussions Contact With Others

Table A2: Selected list of occupations with highest and lowest scores (3 factors)

Machine-Dependent	ICT-Enabled	Physical Proximity
Panel A: Top scores		
Metal workers	Chemical engineers	Nurses*
Fire-fighters*	Chief executives	Personal care workers
Refrigeration mechanics	Community leaders	Child care services managers
Well drillers	Mining engineers	Midwives*
Freight handlers	Supply distribution managers	Traditional medicine professionals
Miners and quarries	Police officers*	Ambulance workers*
Ships' engineers	Inspectors and detectives	Customs and border inspectors*
Boiler operators	Mechanical engineers	Paramedical practitioners*
Electronics mechanics	Biologists	Veterinarians*
Forestry plant operators	Services managers	Police officers*
Panel B: Bottom scores		
Legal professionals	Weaving machine operators	Visual artists
Economists	Laundry machine operators	Livestock farm laborers*
Mathematicians	Shoemaking operators	Subsistence crop farmers*
Credit and loans officers	Subsistence crop farmers	Weaving machine operators
Higher education instructors	Livestock farm labourers	Shoemaking machine operators
Health professionals*	Tobacco products makers	Economists
Arts teachers	Pelt dressers	Garment makers
Language teachers	Sewing machine operators	Sewing machine operators
Human resource managers	Horticultural labourers	Subsistence fishers*
Survey interviewers	Fibre machine operators	Hunters and trappers

Notes: *denote occupations regarded as 'essential' in the Covid-19 pandemic.

¹¹ The scale indicates either the frequency of task, or the importance of the task required in each occupation.

Words can hurt: How political communication can change the pace of an epidemic¹

Lucas Argentieri Mariani,² Jessica Gagete-Miranda³ and Paula Rettl⁴

Date submitted: 22 April 2020; Date accepted: 25 April 2020

Citizens' compliance with measures enacted by health authorities can have an important effect on the state of public health, particularly during epidemics. How much can political leaders influence compliance with such measures? In this paper, we analyze this question in the context of Brazil, where the president Jair Bolsonaro disrespected the recommendations and measures implemented by health authorities during a country-wide pro-government demonstration that took place amid the COVID-19 outbreak. We conclude that Bolsonaro's behavior increased the pace of COVID-19 diffusion. In particular, after the day of the manifestations, the daily number of new COVID-19 is 19% higher in cities that concentrate Bolsonaro's voters as compared to cities that concentrate opposition voters. The impact is verified even in cities where no demonstration took place, which indicates that the quicker spread of COVID-19 was not only due to people agglomerating during the manifestation, but also due to the changed behavior of Bolsonaro's supporters regarding social distancing measures. We directly test this later mechanism exploring an index of social isolation and find that citizens' compliance with social distancing decreased among pro-Bolsonaro cities after the demonstrations.

1 We thank In Loco S/A for kindly making available their social isolation index to the scientific community.

2 PhD Student; University of North Carolina at Chapel Hill.

3 PhD Student; Bocconi University.

4 PhD Student; Bocconi University.

1 Introduction

Citizens' compliance with health measures and policies can have an important effect on the state of public health. This is the case of mass immunization through vaccines, for example [Larson, 2016]. Because social distancing is key in preventing a rapid diffusion of COVID-19, the pandemic exposes the importance of co-production¹ for public health in democracies: the efficacy and efficiency of the measures are dependent upon citizens' compliance with them. What determines public compliance with such measures?

In this article, we address this question by focusing on the specific role of political leaders in persuading the public to comply. Recent evidence for the U.S. suggests that different stances taken by leaders of the Democratic and Republican parties on issues such as how big is the danger represented by COVID-19 and what would be an adequate response to it, made blue and red voters comply at different levels with social distancing measures and hold different attitudes towards the disease [Kushner Gadarian et al., 2020, Allcott et al., 2020, Grossman et al., 2020, Barrios and Hochberg, 2020].

These recent findings confirm that source cues and persuasion can have powerful effects on citizens' attitudes, support for public policies [Tesler, 2012, Nicholson, 2012, Brader and Tucker, 2012, Samuels and Zucco Jr, 2014] and behavior [Ajzenman, 2018] both in the U.S. and in newer democracies with less established party systems. Importantly, this literature concludes that citizens tend to follow the cues of their preferred leader or party when the issue that they have at hand is new or complex and hence heuristics can be particularly helpful in decision-making.

The question of how to reach a balance between controlling the spread of COVID-19 while mitigating the negative impact on the economy is one example of how a new complex policy issue suddenly becomes salient and voters find themselves having to make up their

¹Co-production is the way through which residents of a community contribute to the production of public goods. This contribution does not need to be voluntary, nor the citizen needs to be a direct beneficiary of the service towards which she contributes. In short, co-production is part of "the duties and rights that residents within a community have towards the public administration" [Bertelli and Cannas, 2019]

minds about it with no prior. In this article, we analyze how politicians persuade voters in making up their minds on this issue by exploiting a pro-government manifestation that happened amid the COVID-19 outbreak in Brazil.

On March 15th, 2020, the president Jair Bolsonaro ignored recommendations from the health ministry, his doctors, and the OMS by joining the manifestation, greeting and taking pictures with his supporters. Importantly, before this day Bolsonaro's position on social distancing was unclear as he had urged people to avoid agglomerations on one day and called the new coronavirus "just a flu" on the other. We compare pre and post-trends in daily numbers of new COVID-19 cases at the municipality level and conclude that Bolsonaro's clear taking of stand on that day caused the disease to spread more quickly in cities that concentrate his voters, but not in cities where he did poorly in the previous presidential election.

Jair Bolsonaro is not the only leader to have downplayed the risk of the pandemic for his country. Donald Trump in the U.S. [Lopez, 2020], Andres Manuel Lopez Obrador, in Mexico [Ward, 2020], and Daniel Ortega, in Nicaragua [Rivers and Gallon, 2020] are other examples. However, we argue that Brazil is a particularly interesting case for three main reasons. First, the clear turning point of the president's stance on March 15th, 2020 allows for a clear identification of the effect.

Secondly, the Brazilian case is an example of how cues from a single political leader can have important effects on public opinion and behavior. Indeed, Bolsonaro is increasingly the only relevant political leader in the country who voices skeptical stands on social distancing measures. His own then health ministry - Luiz Henrique Mandetta - repeatedly confronted the president on the issue and 25 out of the 27 governors declared that they would maintain the social distancing measures in spite of the president's open discontentment with them². Moreover, when the outbreak emerged and evolved into an epidemic, Bolsonaro had no party, nor a majority in parliament. Therefore, we can be quite confident that the effect we observe

²<https://g1.globo.com/politica/noticia/2020/03/25/governadoras-reagem-ao-pronunciamento-de-bolsonaro-sobre-coronavirus.ghtml>.

is due to a single man displaying a clear stand on the issue and not due to a coordinated response by an administration or party.

Thirdly, the specific context of Brazil allows us to shed light on how quickly political identities can form and serve as basis for heuristic reasoning. This is a topic that is still understudied. In fact, just recently scholars have started to analyze party cues in fluid party systems [Brader and Tucker, 2012], where political identities are less crystallized. In Brazil, Samuels and Zucco Jr [2014] show that identifying with the two main Brazilian parties (PSDB and PT) play a role in defining political attitudes through source cues, but they did not find the same effect for less established parties. Since Bolsonaro was barely known before 2018 presidential election campaign, our results can contribute to this literature by providing evidence on how quickly political identities can form and become the basis for heuristic reasoning.

In fact, our results show that in municipalities that concentrate Bolsonaro's supporters, the daily number of COVID-19 new cases increases at a higher pace post-manifestations as compared to the pre-manifestations period. Furthermore, this change in trend is not observed in municipalities that concentrate opposition voters. We also find that citizens' compliance with social distancing decreased in pro-Bolsonaro cities after the demonstrations. This suggests that at least in some contexts political identities can form and serve as basis for heuristic reasoning quite quickly, including in cases where the decisions at stake may incur huge costs.

This study contributes to three strains in the literature. First, it adds to the extended literature on source cues and political persuasion [DellaVigna and Gentzkow, 2010, Nicholson, 2012]. In particular, we contribute to the less developed literature on the influence of party and political leader cues on political behavior in new democracies where party systems are weaker [Brader and Tucker, 2012, Samuels and Zucco Jr, 2014]. We also add to the literature on how uninformative persuasion influence attitudes [Mullainathan et al., 2008, Bassi and Rasul, 2017]. Secondly, we contribute to quickly rising literature on political responses to

COVID-19. For example, Allcott et al. [2020], Kushner Gadarian et al. [2020] find that republicans worry less about COVID-19 and comply less with social distancing measures. In addition, Bisbee and Honig [2020] show that not only politics can affect epidemics, but also that epidemics can affect politics. More specifically, COVID-19 decreased support for Sanders in the primaries, because he is seen as a less secure option by anxious voters. We provide new evidence on the relationship between politics and the COVID-19 crisis in another context: Brazil.³

2 Context

The initial cases of COVID-19 happened in the province of Wuhan, in China, on December 2019 [Li et al., 2020]. Due to the highly contagious nature of the disease, in less than three months there were more than 81 thousand cases spread over almost forty countries [WHO, 2020c]. On January 30th 2020, the World Health Organization declared COVID-19 outbreak a global health emergency [WHO, 2020a], further updating the status of the disease to a pandemic on March 11th [WHO, 2020b].

Brazil was the first Latin American country to confirm a COVID-19 case. The outbreak of the virus in the country started on February 26th 2020, in the city of São Paulo [De Sousa and Savarese, 2020]. By the second week of March, more than 400 people had tested positive in the country, mainly in the Southeast region, and state governors started to take action. Eight states started to implement measures of social distancing, such as closing schools, museums, and libraries, banning gatherings with more than 500 people, and limiting business hours [Cerioni, 2020].

Even though Bolsonaro has not taken the virus very seriously - on March 10th, for instance, he said that the COVID-19 crisis was minor and that it was mainly "a fantasy from the media" [Globo, 2020] -, he gave some signals at first suggesting that people should follow

³Ajzenman et al. [2020], a paper developed independently and at the same time uses a similar approach to ours to show that president Jair Bolsonaro influenced his supporters to comply less with social distancing measures.

social distancing measures. While his supporters were planning demonstrations on March 15th against the Congress over an ongoing budget dispute, Bolsonaro went on TV on March 12th to urge the organizers to postpone the demonstrations. He highlighted that at that moment the priority should be people's lives and that going to the demonstration would risk the health of many Brazilians [DW, 2020].

The protest organizers, however, decided to ignore the president's requests and to keep the demonstrations on the original date. They widely broadcast it on social media as a civil disobedience movement, where they apologized to the president, but said that they were going to protest anyway [Uribe and Linhares, 2020].

On March 15th protests happened in around 250 cities in Brazil. In a move that surprised the media, the population, and his own health ministry, Bolsonaro contradicted his own advice and joined the protests in Brasilia to meet and greet demonstrators [Marshall, 2020]. This was particularly striking since he was supposed to be self-isolating after members of a Brazilian delegation to the US led by him were tested positive with COVID-19 [Phillips and Agren, 2020]. After this day, Bolsonaro shifted his attitudes towards COVID-19, and his discourse became increasingly critical to social distancing measures. On March 20th, he criticized governors for closing business, on the ground that this would be detrimental to the economy [Militão, 2020]. On March 24th Bolsonaro went a step further by urging governors in an official to re-open business and arguing that Brazil should implement a vertical quarantine. Moreover, he also mentioned that for the majority of the population, COVID-19 would not be more than a just sniffle [Economist, 2020].


As one can see, Bolsonaro's speech of skepticism regarding coronavirus has escalated quite fast. However, before March 15th, his position was still unclear. Only when he joined the demonstrations disregarding all coronavirus warnings, his attitude towards the crises became clear enough. In section 5, we show how we explore this shift to estimate the impact of this change in stance on the growth of COVID-19 cases.

3 Voters differences in Brazilians' response to COVID-19

In this section, we show some motivating evidence that Bolsonaro's supporters responded to the outbreak of COVID-19 differently from his opponents. A nationally representative poll taken between March 18th and March 20th asked Brazilians about their concerns and behaviors regarding COVID-19 [DataFolha, 2020]. Figure 1 shows responses by those who voted for Bolsonaro in the runoff of 2018 elections (hereinafter referred as "supporters") and those who did not vote for him (i.e., those who either voted for his opponent or who canceled their vote - hereinafter referred as "non-supporters"). Panel A shows that non-supporters were slightly more concerned about being infected by COVID-19: 39% of them were very afraid of being infected and 25% were not afraid. For supporters, these figures were 36% and 28%, respectively.

Panel B brings respondents' speculations of how many deaths will happen in Brazil due to COVID-19. The difference between supporters and non-supporters is more sizeable here: while 50% of non-supporters predicted that many people will die due to COVID-19, only 43% of supporters made the same prediction. Panel C presents support for different measures to stop the spread of the virus. The most striking differences between supporters and non-supporters regard closing non-essential business - interestingly, a higher share of supporters favored this measure (49%, against 42% of non-supporters) - and teleworking - where we see the opposite pattern: 75% of supporters favored such measure, while the share is quite larger among non-supporters (85%). Finally, panel D brings differences in self-reported behavior changes due to COVID-19. The differences between supporters and non-supporters are quite large here: while 77% of non-supporters decreased school/university attendance, only 70% of supporters did so. While 47% of non-supporters decreased the number of days going to work, only 39% of supporters did so. Their behavior regarding outside leisure activities is also different: while 80% of non-supporters decreased such activities, only 76% of supporters did so.

Figure 1: Electorate differences in beliefs and behavior regarding COVID-19



Notes: (i) This figure shows responses to a nationally representative poll [DataFolha, 2020] by Bolsonaro’s voters and non-voters; (ii) The poll took place between March 18th and March 20th (iii) Panel A shows responses about the fear of being infected by COVID-19. Panel B shows priors about the number of deaths caused by COVID-19 in Brazil. Panel C shows support for measures aiming to stop the spread of COVID-19. Panel D shows self-reported changes in behavior due to COVID-19.

In general, Bolsonaro’s supporters seem to take COVID-19 less seriously than non-supporters, and behave accordingly, taking less social distancing measures. Such differences between Bolsonaro’s supporters and non-supporters could stem from different reasons, such as different risk attitudes - Bolsonaro’s supporters might be less risk-averse, for instance -

Covid Economics 12, 1 May 2020: 104-137

or different costs and benefits from social distancing (as derived by Allcott et al. [2020] for Democrats and Republicans in the United States). However, part of these differences might emerge due to the influence that Bolsonaro has over his electorate. His behavior could change their risk perception, or reinforce an apparent trade-off between health and economy, which could induce his supporters to not follow social distancing as closely, eventually leading to a higher spread of the virus in more pro-Bolsonaro areas. In the next sections, we show evidence of this second source of difference. That is, we show that Bolsonaro's behavior did have an impact on his electorate, which led to an increase in COVID-19 cases in more pro-Bolsonaro cities. We also show that the mechanism in place is likely a decrease in social distancing.

4 Data

Our analysis uses data from all Brazilian municipalities that had at least one confirmed COVID-19 case in the period between March 8th and March 26th. The data on COVID-19 cases come from daily updated reports of the State Health Secretariats with information about new confirmed cases, total cases, and deaths related to COVID-19 for each municipality.⁴

We use the results of the 2018 presidential elections to measure cities' support for Bolsonaro.⁵ Brazilian presidential elections are run under a dual-ballot system, where unless a candidate gets more than 50% of votes in the first round, the two most voted candidates dispute a second round or runoff. For each city, we use the runoff results to define support for Bolsonaro in two ways.⁶ The first measure is a binary variable that takes value equal to one if Bolsonaro had the majority of votes in that city, and the second measure is a continuous variable equal to the margin of votes above the 50% cutoff.

⁴This information was compiled by <https://covid19br.wcota.me/> and <https://brasil.io/dataset/covid19/caso>.

⁵Electoral results at the municipal level come from Tribunal Superior Eleitoral (TSE). For more details, see de Leon et al. [2014].

⁶The results for the first round are used in a robustness check specification.

We also explore data on the location of the March 15th demonstrations to check for heterogeneous impacts of Bolsonaro's behavior. These data come from a document sent to the media by the demonstration's organizers ("Movimento Avança Brasil"), listing all the municipalities where protests were confirmed to happen. Out of the 257 municipalities listed in this document, 46 had at least one confirmed COVID-19 case in the time-frame of our analysis and are included in our specification.

Our estimations use the most recent data on cities' GDP per capita and population⁷ as control. These data come from the Brazilian Bureau of Statistics.

Finally, to test whether social distancing is indeed the mechanism in place, we use an index of social isolation that explores data from over 60 million cellphone devices in Brazil and, for each municipality, measures the daily percentage of such devices that remained within a radius of 450 meters from the location identified as their home. This index was developed by In Loco, a Brazilian technology company.⁸

Table 1 presents municipalities' characteristics before the demonstrations. Column (1) brings the average characteristic for all municipalities while columns (2) and (3) report the figures for pro-Bolsonaro and against-Bolsonaro cities, respectively. Pro(Against)-Bolsonaro cities are cities where Bolsonaro won (lost) the runoff in 2018. Column (4) shows the p-value of the difference between columns (2) and (3), after controlling for States' fixed effects. It is reassuring that pro- and against-Bolsonaro cities do not display any striking difference regarding the number of COVID-19 cases before the demonstrations, GDP per capita, and population since the pace of COVID-19 spread and the number of tests performed in suspect cases heavily depend on these characteristics. Still, as shown in the next section, we perform several different checks to ensure that these two groups of municipalities are indeed comparable. The Table also shows that before March 15th, pro- and -against Bolsonaro cities did not have any significant difference regarding their level of isolation.

⁷2015 for the GDP per capita, and 2017 for the population.

⁸For more details about this index, visit <https://www.inloco.com.br>.

	Pro Bolsonaro		Against Bolsonaro	
	All	(runoff)	(runoff)	P-Value
	(1)	(2)	(3)	(4)
Pro Bolsonaro (runoff)	0.82	1.00	0.00	
Pro Bolsonaro (1st round)	0.86	1.00	0.21	0.00
Margin over 50% (runoff)	0.13	0.19	-0.16	0.00
Margin over 50% (1st round)	0.03	0.09	-0.21	0.00
Pre-demonstration COVID-19 cases	0.13	0.14	0.10	0.12
Demonstrations	0.40	0.40	0.38	0.35
GDP per capita (in thousands of 2015 BRL)	33.40	36.30	20.30	0.99
2018 population (in thousands)	413.9	418.4	393.2	0.08
Pre-demonstration isolation index	0.262	0.262	0.264	0.25
N	215	176	39	


Notes: (i) "Pro-Bolsonaro" cities are cities where Bolsonaro got the majority of votes in the 2018 elections; "Margin over 50%" is a continuous variable indicating by how much Bolsonaro's votes exceed 50% in the elections; (ii) Column (4) shows the p-value of the difference between columns (2) and (3) after controlling for States' fixed effects.

Table 1: Descriptive Statistics

Figure 2 presents a motivating illustration for our analysis. It shows the growth in COVID-19 cases between March 15th and March 26th in pro- and against-Bolsonaro cities. Initially, pro-Bolsonaro cities had a slower growth rate if compared to against-Bolsonaro cities. However, while against-Bolsonaro cities keep their growth pace throughout this period, pro-Bolsonaro cities present a sharp increase in their growth rate after March 19th. This gives some evidence of a shift in the proliferation of COVID-19 in this later group of cities. The fact that the shift happened only four days after the demonstrations is consistent with the timing for the symptoms to show up after someone is exposed to the virus.⁹ In the next section, we present the identification strategy to give a causal interpretation to this shift.

⁹The time varies between 2 and 14 days, see more at <https://www.cdc.gov/coronavirus/2019-ncov/about/symptoms.html>.

Figure 2: Average growth in the number of COVID-19 cases for cities with different support for the President during the 2018 Elections



Covid Economics 12, 1 May 2020: 104-137

5 Identification Strategy

We explore the facts described in section 2 to estimate the impact of Bolsonaro’s behavior on the spread of COVID-19 among his supporters. More specifically, we identify the demonstrations on March 15th - when the president disregarded all coronavirus warnings and joined the protests - as the event where Bolsonaro’s attitude regarding the virus became clear. Therefore, we analyze trends in COVID-19 cases around this date, comparing cities with higher and lower support for Bolsonaro. The idea behind this comparison is that cities that concentrate Bolsonaro’s voters are more responsive to his influence than cities that

concentrate opposition voters.

Our identification strategy is a difference-in-differences approach, described by the following estimation:

$$y_{i,s,t} = \alpha_i + \alpha_{st} + X_{i,s,t} * \beta + \delta * (Pro-Bolsonaro_i) \times (Post-March 15th) + \epsilon_{i,s,t} \quad (1)$$

where $y_{i,s,t}$ is the number of COVID-19 cases in municipality i , in State s , at time t . To make municipalities as comparable as possible, it is important to control for time-varying unobserved heterogeneity, such as trends among municipalities with similar characteristics. Hence, besides controlling this estimation for municipality fixed effects, α_i , we also control it for common State by time trends, α_{st} , which absorbs local spillovers among states. Moreover, vector $X_{i,s,t}$ controls for two other important variables. The first variable is municipalities' population interacted with time and with the number of cases one day before the demonstrations - that is, on March 14th. The inclusion of this variable absorbs common trends among cities with a similar population and the same initial number of infected people. This is relevant since population density affects the spread of COVID-19 [Rocklöv and Sjödin, 2020, Stie et al., 2020]. The second variable is municipalities' GDP per capita interacted with time and with the number of cases right before the demonstrations. This absorbs common trends among cities with similar income levels and the same initial number of infected people. This is important since richer municipalities might test suspect cases at a higher rate than poorer municipalities.

Our parameter of interest is δ . As explained in section 4, the variable "*Pro-Bolsonaro*" is defined in two different ways: it is either a binary variable indicating that Bolsonaro had the majority of votes in a municipality, or a continuous variable equal to the margin over 50% of votes. The variable "*Post-March 15th*" is a binary variable that takes value one after

March 15th and zero before this date. Since we estimate equation 1 using the log number of cases, δ can be interpreted as the additional percentage growth in the number of cases in municipalities with higher support for Bolsonaro after the demonstrations on March 15th.

The validity of this identification relies on the assumption that Bolsonaro's supporters were not able to learn about his attitude towards COVID-19 before the protests on March 15th. The fact that the president urged his supporters three days before the protests to do not demonstrate due to health concerns helps to build this argument. Still, we test this assumption looking at dynamic effects of Bolsonaro's behavior before and after March 15th.

We implement the following specification:

$$y_{i,s,t} = \alpha_i + \alpha_{st} + X_{i,s,t} * \beta + \sum_{k=-7}^{10} \delta_k * (Pro-Bolsonaro_i) \times I(Demonstrations_k = 1) + \epsilon_{i,s,t} \quad (2)$$

where $I(Demonstrations_k = 1)$ is a binary variable indicating each day before and after the demonstrations. We look at a window of one week before and ten days after March 15th.

6 Results

Table 2 brings the results of our main estimation. Columns (1) to (4) show the results using the majority of votes as a measure of support for Bolsonaro, while columns (5) to (8) present the results using the margin over 50% as such a measure of support. Columns (1) and (5) control only for State by time fixed effects and the subsequent columns gradually add more controls.

The results of our baseline estimation are robust across all sets of controls we employ in each column. They point to a disproportional growth in cities in which the president has higher support. Our preferred specification with all controls (Column (4)), shows that after the demonstrations on March 15th, cities where Bolsonaro won the majority of votes in


2018 experienced an increase in COVID-19 cases 18.9% higher than cities where he did not have the majority of votes. Column (8) in turn shows that an increase of 1% of Bolsonaro's margin of votes over 50% increases the growth of COVID-19 in around 0.6%.

Figure 3 shows the dynamic results coming from equation 2. Panel A brings such results for cities where Bolsonaro won versus cities where he lost the runoff in 2018, while Panel B shows the results considering the margin over 50% of votes as a measure of the support for Bolsonaro. The figures show that the trends of COVID-19 cases are similar between pro- and against-Bolsonaro cities for a few days after the demonstrations - at least until March 18th, as shown in Panel B. However, after that, pro-Bolsonaro cities experienced quite persistent higher growth in their COVID-19 cases. As said before, the window between March 15th, when the demonstrations took place, and March 18th, when the trends start to diverge, is consistent with the timing for the symptoms to show up after someone is exposed to the virus.

6.1 Discussion about mechanisms

The effects presented in Table 2 could come from two different sources. First, the demonstration itself, since people agglomerating could increase the spread of the disease in localities where the demonstrations took place. Second, citizens could have changed their behavior after seeing Bolsonaro himself ignoring COVID-19 warnings. This could have happened even in cities where demonstrations did not take place. As the ultimate goal of this analysis is to check if the attitudes of the President shaped the behavior of his supporters, we need to understand whether at least part of these effects are coming from a decrease in citizens compliance with social distancing. We start this investigation exploring the location of the protests to isolate the effects of the demonstrations from the indirect effects of Bolsonaro's behavior, both for the cities where the protests took place and the ones where they did not.

Figure 3: Disproportional growth of COVID-19 cases in pro-Bolsonaro cities



Notes: (i) These graphs show the disproportional growth in the number of COVID cases in cities pro Bolsonaro over time when compared to cities against him. In Panel A, cities pro(against) Bolsonaro are the cities where he won(lost) the runoff in the 2018 elections. In Panel B, the support for Bolsonaro is measured with a continuous variable equal to the margins of votes over 50% in the runoff. On March 15th Bolsonaro's supporters marched against the Congress and Bolsonaro ignored coronavirus warnings to join them. (ii) In this specification we control for locality, State by time, and number of cases before the demonstrations interacted with city characteristics (population and GDP per capita) by time FEs. (iii) Confidence interval of 90%. (iv) Standard errors clustered at state-time level.

Covid Economics 12, 1 May 2020: 104-137

	Dependent variable: $\log(1+\text{COVID-19 cases})$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Majority of votes	0.049*** (0.015)	-0.012 (0.009)						
Post March 15 th × Majority of votes	0.404*** (0.094)	0.295*** (0.069)	0.240*** (0.059)	0.189*** (0.050)				
Margin over 50%					-0.046 (0.056)	0.007 (0.017)		
Post March 15 th × Margin over 50%					0.913*** (0.263)	0.919*** (0.172)	0.883*** (0.127)	0.596*** (0.128)
Observations	5950	5850	5850	5850	5950	5850	5850	5850
R Squared	0.34	0.60	0.80	0.81	0.34	0.61	0.80	0.81
State x Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pre-demo N-cases x 2018 Population x Time FE	No	Yes	Yes	Yes	No	Yes	Yes	Yes
City FE	No	No	Yes	Yes	No	No	Yes	Yes
Pre-demo N-cases x GDP per capita X Time FE	No	No	No	Yes	No	No	No	Yes

Notes: (i) Standard errors clustered at state-time level; (ii) * p<0.10, * p<0.05, ** p<0.01; (iii) "Majority of votes": cities where Bolsonaro won the runoff of the 2018 elections; (iv) "Margins over 50%": continuous variable indicating by how much Bolsonaro's votes exceed 50% in the runoff; (v) On March 15th Bolsonaro's supporters marched against the Congress and Bolsonaro ignored coronavirus warnings to join them.

Table 2: Growth in COVID-19 cases after demonstrations on March 15th

The specification of this test is the following:

$$\begin{aligned}
 y_{i,s,t} = & \alpha_i + \alpha_{st} + X_{i,s,t} * \beta + \delta_1 I(\text{Demonstration}_i = 1) \times (\text{Post March 15th}) \\
 & + \delta_2 (\text{Pro-Bolsonaro}_i) \times I(\text{Demonstrations}_i = 1) \times (\text{Post March 15th}) \quad (3) \\
 & + \delta_3 (\text{Pro-Bolsonaro}_i) \times I(\text{Demonstrations}_i = 0) \times (\text{Post March 15th})
 \end{aligned}$$

where $I(\text{Demonstrations}_i = 1)$ is an indicator variable that takes value one if the city had a demonstration on March 15th and $I(\text{Demonstrations}_i = 0)$ is an indicator variable that takes value one if the city did not have a demonstration on March 15th. The parameter δ_1 represents the disproportional growth in the number of cases due to the direct effects of the demonstrations, which might be caused by the agglomeration. The parameters δ_2 and δ_3 are the indirect effects of the demonstrations for cities with a higher concentration of Bolsonaro's voters. For the cities that had a demonstration, the parameter δ_2 could be an effect of the size of the demonstrations in the locality since it is likely that these cities have more supporters. Parameter δ_3 , however, is likely to represent the indirect effects of the President's behavior on his supporters' behavior after the demonstrations, since there were no agglomeration effects on these localities.

Table 3 presents the results of this estimation. Columns (1) and (2) use the majority of votes as a measure of support for Bolsonaro while columns (3) and (4) use the margin over 50% as such a measure of support. Among cities with a lower concentration of Bolsonaro's voters, the presence of protests on March 15th increases the growth of COVID-19 from 21.8% to 26.2%, depending on the specification. This number is even higher among cities with a larger concentration of Bolsonaro's voters. In column (2), for instance, we see that cities where Bolsonaro won the runoff in 2018 and where a demonstration took place on March

15th, the growth in COVID-19 cases is 43.3% higher than in cities where he lost the runoff and that did not host demonstrations.

Most interestingly, even among cities where demonstrations did not take place, those with a higher concentration of Bolsonaro's voters experienced higher growth in COVID-cases than those with a lower concentration of his voters. In column (2) we see that places in which Bolsonaro had the majority of votes in 2018 experienced a growth in COVID-19 cases 14% higher than places where he did not get such a majority. Column (4), in turn, shows that if the margin over 50% of the votes for Bolsonaro increases by 1%, the growth in COVID-19 cases increases 0.43%. This indicates that the results found in Table 2 are not driven only to the fact that people agglomerated during the protests. This is certainly part of the increase in COVID-19 cases, but it does not explain all of it. The fact that the virus also spread faster in places that concentrate Bolsonaro's voters but did not host protests indicates that residents of these places might have stopped to follow social distancing measures after March 15th, possibly influenced by Bolsonaro's behavior during the demonstrations.

Table 4 brings further evidence that these results are indeed driven by the behavior of Bolsonaro's supporters. We replicate the exercises presented in Table 3 using the results of the first round of the 2018 elections, instead of the runoff. Votes during the first round of a dual-ballot system are more likely to be sincere and less likely to be strategic than in the second round [Duverger, 1954, Fujiwara et al., 2011]. Therefore, we expect Bolsonaro to have more influence on those who voted for him already in the first round since they have a stronger admiration for him. If this is true, the impact of his behavior on March 15th should be even higher among this group of voters. Table 4 shows that this is indeed the case. In particular, if we look at localities that did not host demonstrations in March 15th, cities in which Bolsonaro had the majority of votes in the first round have a growth in COVID-19 cases 25.6% larger than in cities where he did not get this majority.

	Dependent variable: log(1+COVID Cases)			
	(1)	(2)	(3)	(4)
Post March 15 th × Local demonstration	0.262*** (0.080)	0.247*** (0.076)	0.240*** (0.043)	0.218*** (0.040)
Post March 15 th × Supports Bolsonaro × Local demonstration	0.238*** (0.064)	0.186*** (0.054)	0.728*** (0.133)	0.539*** (0.137)
Post March 15 th × Supports Bolsonaro × No demonstration	0.266*** (0.085)	0.140* (0.081)	0.692*** (0.166)	0.434*** (0.167)
Observations	5850	5850	5850	5850
R Squared	0.75	0.82	0.80	0.82
City FE	Yes	Yes	Yes	Yes
Pre-demo N-cases x 2018 Population x Time FE	Yes	Yes	Yes	Yes
State x Time FE	Yes	Yes	Yes	Yes
Pre-demo N-cases x GDP per capita X Time FE	No	Yes	No	Yes

Notes: (i) Standard errors clustered at state-time level; (ii) * p<0.10, * p<0.05, ** p<0.01; (iii) In in Columns (1) and (2), "Supports Bolsonaro" is a binary variable that takes value equal 1 if Bolsonaro won the runoff of the 2018 elections; in columns (3) and (4), it is a continuous variable indicating by how much Bolsonaro's votes exceed 50% in the runoff; (iv) On March 15th Bolsonaro's supporters marched against the Congress and Bolsonaro ignored coronavirus warnings to join them.

Table 3: Heterogeneous impacts by local demonstrations



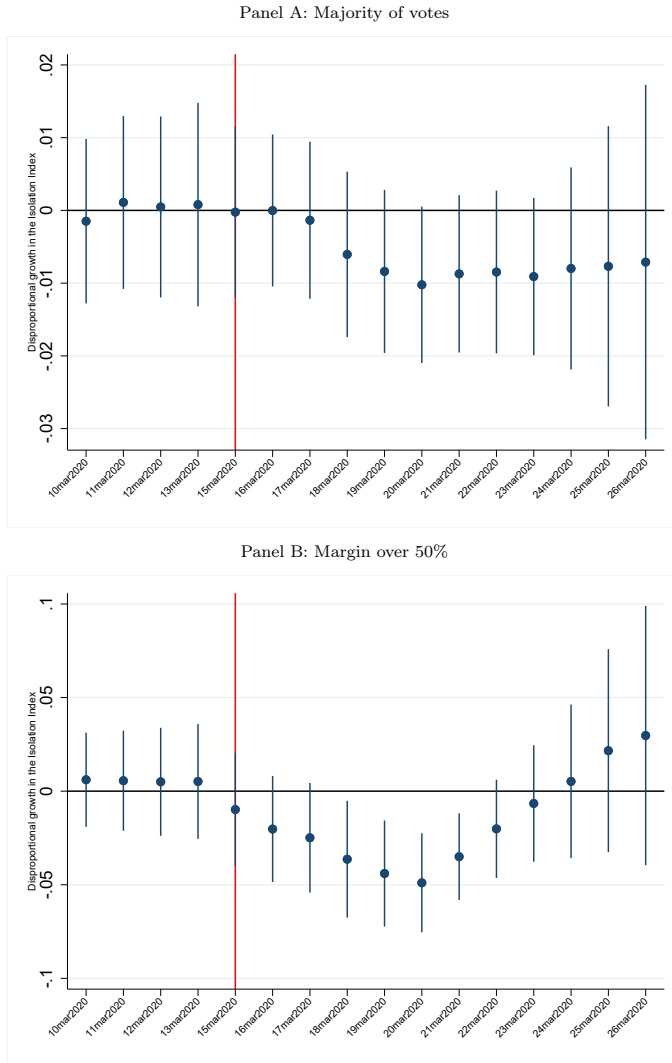
	Dependent variable: log(1+COVID Cases)			
	(1)	(2)	(3)	(4)
Post March 15 th × Local demonstration	0.138 (0.123)	0.206* (0.120)	0.182*** (0.036)	0.161*** (0.035)
Post March 15 th × Supports Bolsonaro × Local demonstration	0.322*** (0.070)	0.302*** (0.065)	1.019*** (0.161)	0.859*** (0.160)
Post March 15 th × Supports Bolsonaro × No demonstration	0.424*** (0.097)	0.256*** (0.093)	1.139*** (0.221)	0.902*** (0.210)
Observations	4896	4896	4896	4896
R Squared	0.76	0.83	0.81	0.83
City FE	Yes	Yes	Yes	Yes
Pre-demo N-cases x 2018 Population x Time FE	Yes	Yes	Yes	Yes
State x Time FE	Yes	Yes	Yes	Yes
Pre-demo N-cases x GDP per capita X Time FE	No	Yes	No	Yes

Notes: (i) Standard errors clustered at state-time level; (ii) * p<0.10, * p<0.05, ** p<0.01; (iii) "Supports Bolsonaro" in Columns (1) and (2) is a binary variable that takes value equal 1 if Bolsonaro was the most voted candidate in the first round of the 2018 elections, and in Columns (3) and (4) is a continuous variable indicating by how much Bolsonaro's votes exceed 50% in the first round of the 2018 elections; (iv) On March 15th Bolsonaro's supporters marched against the Congress and Bolsonaro ignored coronavirus warnings to join them.

Table 4: Heterogeneous impacts by local demonstrations - first round of the 2018 elections.



Figure 4: Disproportional decrease of social isolation in pro-Bolsonaro cities



Notes: (i) These graphs show the disproportional decrease in index of social isolation in cities pro Bolsonaro over time when compared to cities against him. In Panel A, cities pro(against) Bolsonaro are the cities where he won(lost) the runoff in the 2018 elections. In Panel B, the support for Bolsonaro is measured with a continuous variable equal to the margins of votes over 50% in the runoff. On March 15th Bolsonaro’s supporters marched against the Congress and Bolsonaro ignored coronavirus warnings to join them. (ii) The index of social isolation represent the percentage of mobile devices that remained within a radius of 450m from the location identified as their home. (iii) In this specification we control for locality, State by time, and number of cases before the demonstrations interacted with city characteristics (population and GDP per capita) by time FEs. (iv) Confidence interval of 90%. (v) Standard errors clustered at state-time level.

We finally explore the index of social isolation to measure what is driving the increase

in COVID-19 cases in the most direct possible way. As explained in section 4, the index of social isolation gathers data from over 60 million mobile devices and estimates for each municipality/day the percentage of devices that remained within a 450 meters radius from their home. Similar data is used to measure how partisanship [Barrios and Hochberg, 2020, Grossman et al., 2020] and belief in science affects social distancing in the US [Brzezinski et al., 2020].

In addition, in a paper developed independently and at the same time as ours, Ajzenman et al. [2020] use this measure of social isolation to show that Bolsonaro has influenced his supporters to comply less with social distancing measures. Here, we use this index to check whether this also applies to the sub-sample of cities we analyze - that is, cities that had at least one COVID-19 case in the considered window of time and hence were at a similar stage of the epidemic. In our analysis, we also consider different trends depending on the number of cases before the demonstrations and municipality characteristics as controls.

There are some caveats of using this index as a proxy of compliance to social distancing. Firstly, it relies on data collected from mobile devices that have GPS, Bluetooth, and/or Wi-Fi running in their background. Such a sample of devices might not be representative of the population. In particular, it is more likely that people from higher socioeconomic status are over-represented in this sample.

Secondly, the radius established by such an index - 450 meters from people's home - is rather arbitrary and might represent different levels of non-compliance to social distancing depending on the size of the city - in a city like São Paulo or Rio de Janeiro, for instance, someone might have to go further than 450 meters just to do their groceries. Indeed, Barrios and Hochberg [2020] address this issue by using the change in average daily distance traveled from the pre-pandemic period, while Painter and Qiu [2020] do not consider movements from home to work. Unfortunately, the data provided by In Loco does not allow us to compute a similar measure. Finally, staying within such a radius also does not necessarily mean that one is complying with social distancing since people could still be meeting their neighbors.

With these caveats in mind, we perform exercises similar to those described in equation 2 but with the index of social isolation as the dependent variable. In line with Brzezinski et al. [2020], we account for differences in people's propensity to travel during week days and weekends. Figure 4 brings the results of such an estimation. Panel A shows the results for cities where Bolsonaro won versus cities where he lost the runoff in 2018, while Panel B brings the results considering the margin over 50% of votes as a measure of the support for Bolsonaro. In Panel A we observe a decrease in social isolation, even though this is not significantly different from zero.¹⁰ Panel B shows a more clear trend of decrease in cities where Bolsonaro enjoys higher support, starting already on March 16th and persisting until March 22nd. Overall, these results indicate that Bolsonaro's supporters reacted to his behavior during the demonstrations, decreasing their own isolation.

6.2 Robustness

Tables A.1 and A.2 in the Appendix bring some checks of the robustness of our results. Table A.1 presents estimations excluding cities in the State of São Paulo. São Paulo was the State where the outbreak of COVID-19 happened, and consequently had already a quite large number of confirmed cases when the demonstrations took place. Moreover, Bolsonaro had the majority of votes in all São Paulo cities present in our sample. The results observed in Table 3 could, therefore, be driven solely by the fact that the cluster of cities with the highest number of COVID-19 are also cities that support Bolsonaro. However, this does not seem to be the case, as the results are pretty stable when we exclude these cities from our sample.

Table A.2 reproduces our estimations with different measures for the occurrence of COVID-19 cases in the analyzed localities. Columns (1) and (2) present the result using the likelihood of having any confirmed COVID-19 case in the city as the dependent variable. Columns (4) and (5) in turn use the number of cases by thousands of inhabitants. In both

¹⁰Some of the coefficients after March 17th are significant at 15% level.

cases, results are quite similar to our original analysis.

Recent studies indicate that the number of cases of COVID-19 could be much higher due to low number of tests performed in Brazil. The number of cases can be 12 times larger than the official numbers according to studies which use the number of COVID-19 like deaths¹¹. With this caveat in mind, we also analyse the effect of Bolsonaro's behavior on March 15th on growth in the number of deaths caused by COVID-19 and other related symptoms (henceforth COVID-19 related deaths), namely pneumonia, severe acute respiratory syndrome, respiratory failure and septicemia¹². We decided to consider also death caused by these related symptoms because not all severely ill people who present COVID-19 symptoms are being tested in Brazil [Lemos [2020]]. In line with our other results, Figure 5 shows that pro-Bolsonaro cities also displayed a disproportional increase of COVID-19 related deaths. In the future, we plan to compare these numbers with the ones in 2019 and calculate the excess of deaths [Ghislandi et al., 2020].

7 Conclusion

Citizens' compliance with public health measures are extremely important for containing the spread of contagious diseases. Previous work has focused on different causes of compliance with health measures, concluding that trust in health authorities is relevant [Larson, 2016]. However, much less work has been developed on the politicization of health measures. In this paper, we shed light on how a single political leader - the Brazilian president Jair Bolsonaro - can influence compliance with health measures when his stance is in stark contrast with the position of health authorities and experts.

We conclude that Bolsonaro indeed influenced his supporters to comply less with social distancing measures by joining nation-wide demonstrations organized by his supporters on

¹¹For more details, visit <https://www.businessinsider.com/brazils-coronavirus-cases-likely-12-times-higher-than-reported-2020-4>.

¹²Data on number of deaths and their causes come from legal death certificates (<https://cartorionobrasil.com.br/>)

March 15th. His participation to these demonstrations signaled to his supporters that he was not in favor of social distancing measures. As a result, we observe an increase of 18.9% on the daily number of new registered COVID-19 cases in municipalities that concentrate his supporters. This result is robust to a number of measures of municipalities' support for Bolsonaro. We also provide evidence that the increase in the number of cases was not caused only by the agglomeration of people in the day of the protest. Indeed, we show that a similar effect is observed also in pro-Bolsonaro cities where no manifestation was registered. Finally, we provide evidence as to which mechanisms underlie this effect. Specifically, we show that after the demonstrations, cities where Bolsonaro enjoys higher support decreased their levels of social isolation when compared to cities where he has lower support.

References

- Nicolás Ajzenman. The power of example: corruption spurs corruption. 2018.
- Nicolás Ajzenman, Tiago Cavalcanti, and Daniel Da Mata. More than words: Leaders' speech and risky behavior during a pandemic. *Available at SSRN 3582908*, 2020.
- Hunt Allcott, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Y Yang. Polarization and public health: Partisan differences in social distancing during covid-19. *Available at SSRN 3570274*, 2020.
- John M Barrios and Yael Hochberg. Risk perception through the lens of politics in the time of the covid-19 pandemic. Technical report, National Bureau of Economic Research, 2020.
- Vittorio Bassi and Imran Rasul. Persuasion: A case study of papal influences on fertility-related beliefs and behavior. *American Economic Journal: Applied Economics*, 9(4):250–302, 2017.
- Anthony M Bertelli and Silvia Cannas. Law and co-production: The importance of citizenship values. 2019.
- James Bisbee and Dan Honig. Flight to safety: 2020 democratic primary election results and covid-19. *Covid Economics*, 3(10 April 2020):54–84, 2020.
- Ted Brader and Joshua A Tucker. Following the party's lead: Party cues, policy opinion, and the power of partisanship in three multiparty systems. *Comparative Politics*, 44(4):403–403, 2012.
- Adam Brzezinski, Guido Deiana, Valentin Kecht, and David Van Dijke. The covid-19 pandemic: Government vs. community action across the united states. 2020.
- Clara Cerioni. Como os estados brasileiros estão agindo para conter o coronavírus. *Exame*, Mar 2020. URL <https://exame.abril.com.br/brasil/como-os-estados-brasileiros-estao-agindo-para-conter-o-coronavirus/>.

Pesquisa DataFolha. Maioria está bem informada sobre coronavírus e defende isolamento total. <http://datafolha.folha.uol.com.br/opiniaopublica/2020/03/1988617-maioria-esta-bem-informada-sobre-coronavirus-e-defende-isolamento-total.shtml>, 2020.

Lopez de Leon, Fernanda Leite, and Renata Rizzi. A test for the rational ignorance hypothesis: Evidence from a natural experiment in Brazil. *American Economic Journal: Economic Policy*, 6(4):380–98, 2014.

Marcelo De Sousa and Mauricio Savarese. Brazil confirms first coronavirus case in Latin America. *AP NEWS*, Mar 2020. URL <https://apnews.com/fd3d0d0120dd10f3d09bad78a4dd9539>.

Stefano DellaVigna and Matthew Gentzkow. Persuasion: empirical evidence. *Annu. Rev. Econ.*, 2(1):643–669, 2010.

Maurice Duverger. *Political Parties: Their Organisation and Activity in the Modern State*. Methuen & Company, 1954.

DW. Bolsonaro pede que população não vá a protestos do dia 15. *DW Brasil*, Mar 2020. URL <https://www.dw.com/pt-br/bolsonaro-pede-que-populaç~ao-n~ao-vá-a-protestos-do-dia-15>.

The Economist. Brazil's president fiddles as a pandemic looms. *The Economist*, Mar 2020. URL <https://www.economist.com/the-americas/2020/03/26/brazils-president-fiddles-as-a-pandemic-looms>.

Thomas Fujiwara et al. A regression discontinuity test of strategic voting and Duverger's law. *Quarterly Journal of Political Science*, 6(3–4):197–233, 2011.

Simone Ghislandi, Raya Muttarak, Markus Sauerberg, and Benedetta Scotti. News from the front: Excess mortality and life expectancy in two major epicentres of the COVID-19 pandemic in Italy. Technical report, 2020.

Globo. Bolsonaro diz que 'pequena crise' do coronavírus é 'mais fantasia' e não 'isso tudo' que mídia propaga. *Globo*, Mar 2020. URL <https://g1.globo.com/politica/noticia/2020/03/10/bolsonaro-diz-que-questao-do-coronavirus-e-muito-mais-fantasia.ghtml>.

Guy Grossman, Soojong Kim, Jonah Rexer, and Harsha Thirumurthy. Political partisanship influences behavioral responses to governors' recommendations for covid-19 prevention in the united states. *Available at SSRN 3578695*, 2020.

Shana Kushner Gadarian, Sara Wallace Goodman, and Thomas B Pepinsky. Partisanship, health behavior, and policy attitudes in the early stages of the covid-19 pandemic. *Health Behavior, and Policy Attitudes in the Early Stages of the COVID-19 Pandemic (March 27, 2020)*, 2020.

Heidi J Larson. Vaccine trust and the limits of information. *Science*, 353(6305):1207–1208, 2016.

Vinicius Lemos. Coronavírus: mortes suspeitas acendem alerta sobre total de vítimas no brasil. *BBC Brasil*, Mar 2020. URL <https://www.bbc.com/portuguese/brasil-52059165>.

Qun Li, Xuhua Guan, Peng Wu, Xiaoye Wang, Lei Zhou, Yeqing Tong, Ruiqi Ren, Kathy SM Leung, Eric HY Lau, Jessica Y Wong, et al. Early transmission dynamics in wuhan, china, of novel coronavirus-infected pneumonia. *New England Journal of Medicine*, 2020.

German Lopez. The trump administration's botched coronavirus response, explained. *Vox*, Mar 2020. URL <https://www.vox.com/policy-and-politics/2020/3/14/21177509/coronavirus-trump-covid-19-pandemic-response>.

Euan Marshall. Brazil president jair bolsonaro ignores coronavirus warnings to take selfies with demonstrators. *The Telegraph*, Mar 2020. URL <https://www.telegraph.co.uk/news/worldnews/south-america/2020/03/14/bolsonaro-covid-19-warnings-selfies/>.

[//www.telegraph.co.uk/news/2020/03/16/brazil-president-jairbolsonaro-ignores-coronavirus-warnings/](https://www.telegraph.co.uk/news/2020/03/16/brazil-president-jairbolsonaro-ignores-coronavirus-warnings/).

Eduardo Militão. Bolsonaro diz que fechar shoppings e comércio prejudica a economia. *UOL*, Mar 2020. URL <https://economia.uol.com.br/noticias/redacao/2020/03/20/bolsonaro-diz-que-fechar-shoppings-e-comercio-prejudica-a-economia.htm>.

Sendhil Mullainathan, Joshua Schwartzstein, and Andrei Shleifer. Coarse thinking and persuasion. *The Quarterly journal of economics*, 123(2):577–619, 2008.

Stephen P Nicholson. Polarizing cues. *American journal of political science*, 56(1):52–66, 2012.

Marcus Painter and Tian Qiu. Political beliefs affect compliance with covid-19 social distancing orders. *Available at SSRN 3569098*, 2020.

Tom Phillips and David Agren. Bolsonaro and amlo slammed for snubbing coronavirus warnings. *The guardian*, Mar 2020. URL <https://www.theguardian.com/world/2020/mar/15/bolsonaro-amlo-coronavirus-warnings>.

Matt Rivers and Natalie Gallon. Why are these three presidents downplaying coronavirus warnings? *CNN*, Mar 2020. URL <https://www.cnn.com/2020/03/24/americas/coronavirus-latin-america-presidents-intl/index.html>.

J Rocklöv and H Sjödin. High population densities catalyze the spread of COVID-19. *Journal of Travel Medicine*, 03 2020. ISSN 1195-1982. doi: 10.1093/jtm/taaa038. URL <https://doi.org/10.1093/jtm/taaa038>. taaa038.

David Samuels and Cesar Zucco Jr. The power of partisanship in brazil: Evidence from survey experiments. *American Journal of Political Science*, 58(1):212–225, 2014.

Andrew Stie, Marc Berman, and Luis Bettencourt. COVID-19 Attack Rate Increases with

City Size. 04 2020. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3564464.

Michael Tesler. The spillover of racialization into health care: How president obama polarized public opinion by racial attitudes and race. *American Journal of Political Science*, 56(3): 690–704, 2012.

Gustavo Uribe and Carolina Linhares. Mesmo após cancelamento, bolsonaristas falam em manter atos pró-governo. *Folha de São Paulo*, Mar 2020. URL <https://www1.folha.uol.com.br/poder/2020/03/cancelamento-de-atos-nao-acalma-animos-de-bolsonaristas-contra-congresso.shtml>.

Alex Ward. Mexico's coronavirus-skeptical president is setting up his country for a health crisis. *Vox*, Mar 2020. URL <https://www.vox.com/2020/3/26/21193823/coronavirus-mexico-andres-manuel-lopez-obrador-health-care>.

WHO. Who director-general's statement on ihr emergency committee on novel coronavirus (2019-ncov). *World Health Organization*, Mar 2020a. URL [https://www.who.int/dg/speeches/detail/who-director-general-s-statement-on-ihr-emergency-committee-on-novel-coronavirus-\(2019-ncov\)](https://www.who.int/dg/speeches/detail/who-director-general-s-statement-on-ihr-emergency-committee-on-novel-coronavirus-(2019-ncov)).

WHO. Who director-general's opening remarks at the media briefing on covid-19 - 11 march 2020. *World Health Organization*, Mar 2020b. URL <https://www.who.int/dg/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19---11-march-2020>.

WHO. Coronavirus disease 2019 (covid-19): situation report, 19. 2020c.

A Appendix

	Dependent variable: log(1+COVID Cases)			
	(1)	(2)	(3)	(4)
Post March 15 th × Local demonstration	0.267*** (0.080)	0.250*** (0.077)	0.258*** (0.045)	0.247*** (0.043)
Post March 15 th × Supports Bolsonaro × Local demonstration	0.210*** (0.062)	0.179*** (0.054)	1.006*** (0.141)	0.895*** (0.142)
Post March 15 th × Supports Bolsonaro × No demonstration	0.281*** (0.086)	0.178** (0.084)	1.119*** (0.163)	0.911*** (0.166)
Observations	4975	4975	4975	4975
R Squared	0.76	0.83	0.81	0.83
City FE	Yes	Yes	Yes	Yes
Pre-demo number of cases x Time FE	Yes	Yes	Yes	Yes
State x Time FE	Yes	Yes	Yes	Yes
Pre-demo number of cases x 2017 GDP per capita X Time FE	No	Yes	No	Yes

Notes: (i) Standard errors clustered at state-time level; (ii) * p<0.10, * p<0.05, ** p<0.01; (iii) "Supports Bolsonaro" is the margin of votes for Bolsonaro above the 50% cutoff in the runoff of the 2018 elections; (iv) On March 15th Bolsonaro's supporters marched against the Congress and Bolsonaro ignored coronavirus warnings to join them.

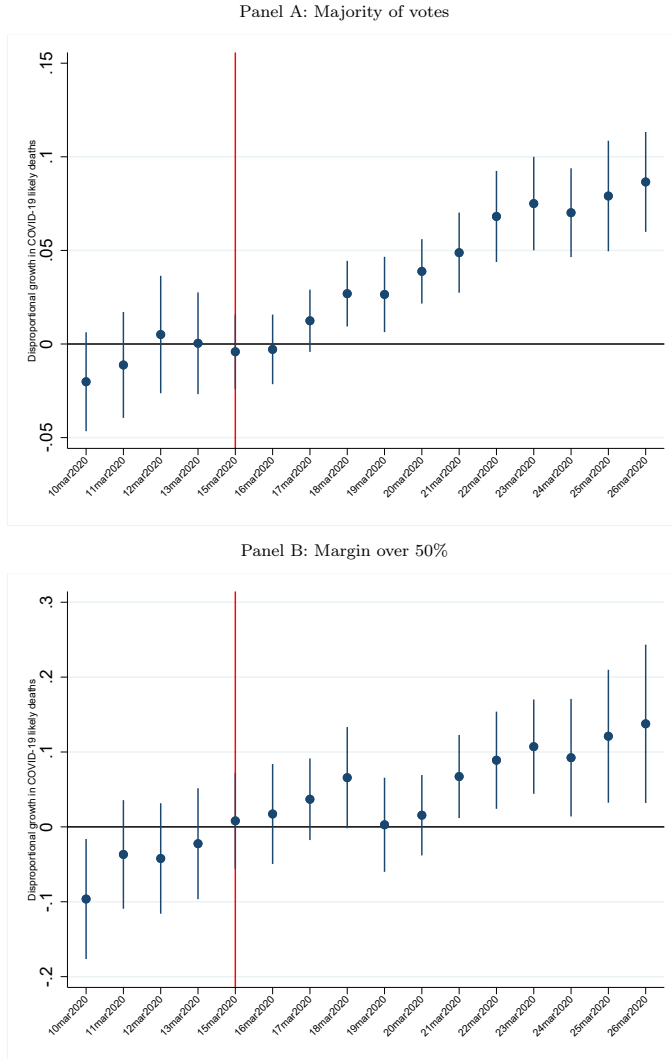
Table A.1: Robustness check - results from Table 3 excluding the state of São Paulo.

	I(Number of cases>0)		Cases by 1k people	
Post March 15 th × Supports Bolsonaro	0.285***		0.030***	
	(0.077)		(0.006)	
Post March 15 th × Local demonstration		0.031		0.001
		(0.024)		(0.001)
Post March 15 th × Supports Bolsonaro × Local demonstration		0.231***		0.033***
		(0.084)		(0.007)
Post March 15 th × Supports Bolsonaro × No demonstration		0.360***		0.017***
		(0.113)		(0.006)
Observations	5850	5850	5850	5850
R Squared	0.82	0.82	0.64	0.64
City FE	Yes	Yes	Yes	Yes
Pre-demo number of cases x Population X Time FE	Yes	Yes	Yes	Yes
State x Time FE	Yes	Yes	Yes	Yes

Notes: (i) Standard errors clustered at state-time level; (ii) * p<0.10, * p<0.05, ** p<0.01; (iii) "Supports Bolsonaro" is the margin of votes for Bolsonaro above the 50% cutoff in the runoff of the 2018 elections; (iv) On March 15th Bolsonaro's supporters marched against the Congress and Bolsonaro ignored coronavirus warnings to join them.

Table A.2: Robustness check - results from Tables 2 and 3 using alternative measures of the number of COVID-19 cases growth.

Figure 5: Disproportional growth in the COVID-19 like deaths in pro-Bolsonaro cities



Notes: (i) These graphs show the disproportional growth in the COVID-19 like deaths in cities pro-Bolsonaro over time when compared to cities against him. In Panel A, cities pro(against) Bolsonaro are the cities where he won(lost) the runoff in the 2018 elections. In Panel B, the support for Bolsonaro is measured with a continuous variable equal to the margins of votes over 50% in the runoff. On March 15th Bolsonaro's supporters marched against the Congress and Bolsonaro ignored coronavirus warnings to join them. (ii) The deaths included here are due to COVID-19, pneumonia, acute breathing insufficiency, and septicemia. (iii) In this specification we control for locality, State by time, and number of cases before the demonstrations interacted with city characteristics (population and GDP per capita) by time FEs. (iv) Confidence interval of 90%. (v) Standard errors clustered at state-time level.

Covid Economics 12, 1 May 2020: 104-137

Will Covid-19 affect inequality? Evidence from past pandemics¹

Davide Furceri,² Prakash Loungani,³ Jonathan D. Ostry⁴
and Pietro Pizzuto⁵

Date submitted: 24 April 2020; Date accepted: 25 April 2020

This paper provides evidence on the impact of major epidemics from the past two decades on income distribution. Our results justify the concern that the current pandemic could end up exerting a significant impact on inequality: past events of this kind, even though much smaller in scale, have led to increases in the Gini coefficient, raised the income shares of higher income deciles, and lowered the employment-to-population ratio for those with basic education compared to those with higher education. We provide some evidence that the distributional consequences from the current pandemic may be larger than those flowing from the historical pandemics in our sample.

- 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.
- 2 Deputy Division Chief, International Monetary Fund and University of Palermo.
- 3 Assistant Director, International Monetary Fund and Johns Hopkins University.
- 4 Deputy Director, International Monetary Fund and CEPR Research Fellow.
- 5 Adjunct Professor, University of Palermo.

I. INTRODUCTION

Deaths from the Covid-19 pandemic have already exceeded 200,000 according to official statistics. This tragic cost has been accompanied by the upending of millions of other lives as governments take necessary steps to limit the spread of the virus. In the United States, for instance, evidence from a large-scale survey of households suggests that 20 million jobs were lost by early April, far more than were lost over the entire Great Recession of 2008-09 (Coibion, Gorodnichenko and Weber 2020). While the global job loss is more difficult to gauge, the decline in working hours thus far, which is easier to track in real-time, is already equivalent to a decline in 195 million full-time jobs (ILO 2020).

While most, if not all, economic classes are adversely affected by the pandemic in one way or another, it is possible that people in low-income deciles and low-skilled workers may end up being disproportionately hurt. Indeed, there is already anecdotal evidence of the substantial effects of the pandemic on these groups, raising concerns that it will end up raising inequality in many countries. There are direct and fairly immediate effects from low-income groups being more prone to the disease; as one example, Schmitt-Grohe, Teoh and Uribe (2020) find that in New York City, poor people are less likely to test negative for Covid-19: moving from the richest to the poorest zip codes is associated with a decline in the fraction of negative test results from 65 to 38 percent. Recent analysis from the Kansas City Fed suggests that workers with non-college education have taken the largest hit in the first wave of job losses due to Covid-19 in the United States.¹

In addition, there are indirect and longer-lasting effects from possible job loss and other shocks to income and diminished employment prospects. The ILO estimates that 1.25 billion workers, representing nearly 40 per cent of the global workforce, are employed in sectors that face high risk of worker displacement. These sectors also have a high proportion of workers in informal employment, with limited access to health services and social protection (ILO 2020). Despite attempts by governments to limit the damage, such workers run a high risk of facing challenges in regaining their livelihoods even after economies start to recover. In many countries, low-income

¹ <https://www.kansascityfed.org/en/publications/research/eb/articles/2020/women-take-bigger-hit-job-losses-covid19>

households can also suffer an impact on non-labor income due to decline in remittances as the pandemic affects the livelihoods of migrants. The World Bank estimates that global remittance flows, which fell 5% during the 2009 financial crisis, will fall 20% this year, which would mark the sharpest decline since 1980.

To shed light on such potential impacts of Covid-19, this paper provides evidence on the impact of pandemics and major epidemics² from the past two decades on income inequality, income shares of the top and bottom deciles, and the employment prospects of people with low education levels (using educational attainment as a proxy for skills). Our results justify the concern that Covid-19 could end up exerting a significant impact on inequality. Past pandemics, even though much smaller in scale, have led to increases in the Gini coefficient, raised the income shares of higher deciles of income, and lowered the employment-to-population ratio for those with basic education compared to those with higher education.

This paper relates to two main strands of literature. The first is the literature on the economic effects of pandemics (for recent contributions, see Atkeson 2020; Barro et al. 2020; Eichenbaum et al. 2020; Jorda et al. 2020; Ma et al. 2020). This literature provides evidence of large and persistent effect on economic activity. In particular, Ma et al. (2020) examined the same set of episodes considered in our paper and found that real GDP is 2.6% lower on average across 210 countries in the year the outbreak is officially declared and remains 3% below pre-shock level five years later. The second strand of the literature is on the role of crises and recessions in exacerbating inequality by depressing employment for those most vulnerable, such as less skilled and youth (see de Haan and Sturm 2017 and references therein).

The remainder of the paper is structured as follows. Section II describes our data and econometric method and Section III presents our results. The last section concludes and outlines avenues for future work on this topic.

² For convenience, we refer to all these events as pandemics.

II. DATA AND ECONOMETRIC METHOD

Income distribution

Our data on various measures of distribution come from three sources. Table A1 in the Appendix provides summary statistics on the variables used in the analysis.

- Gini coefficients are from the Standardized World Income Inequality Database (SWIID), which combines information from the United Nations World Income Database (UNWIDER) and the Luxembourg Income Study (LIS). SWIID provides comparable estimates of market income inequality for 175 countries from 1961 to the present.³
- Income shares by decile are from the World Bank’s World Development Indicators. This source provides internationally comparable statistics for a large number of economies; however, for many countries the time series is rather short, so in the end our results on income deciles are for a limited sample of 64 countries.⁴
- Data on employment by skill levels are difficult to obtain for a large group of countries. The ILO notes that “statistics on levels of educational attainment remain the best available indicators of labor force skill levels.” Hence, we use ILO data on employment-to-population ratios for different education levels—advanced, tertiary and basic.⁵

Pandemic events

Following Ma et al. 2020, we focus on five major events: SARS in 2003; H1N1 in 2009; MERS in 2012; Ebola in 2014; and Zika in 2016. The list of countries in our sample that are affected by each event is given in Table A2 in the Appendix. Among the five events, the most

³ See Solt (2009) for details on the construction of this data set.

⁴ See <https://datacatalog.worldbank.org/dataset/world-development-indicators> for details.

⁵ See <https://ilostat.ilo.org/resources/methods/description-employment-by-education/> for details.

widespread one is H1N1 (Swine Flu Influenza). We construct a dummy variable, the pandemic event, which takes the value 1 when WHO declares a pandemic for the country and 0 otherwise.

Empirical methodology

To estimate the distributional impact of pandemics, we follow the method proposed by Jordà (2005) and estimate impulse response functions directly from local projections:

$$y_{i,t+k} - y_{i,t-1} = \alpha_i^k + \gamma_t^k + \beta^k D_{i,t} + \theta^k X_{i,t} + \varepsilon_{i,t+k} \quad (1)$$

where $y_{i,t}$ is the log of our distribution variables (e.g. the Gini coefficient) for country i in year t ; α_i are country fixed effects, included to take account of differences in countries' average income distribution; γ_t are time fixed effects, included to take account of global shocks such as shifts in oil prices or the global business cycle; $D_{i,t}$ is a dummy variable indicating a pandemic event that affects country i in year t . $X_{i,t}$ is a vector that includes two lags of the dependent variable and the pandemic dummy. In the baseline, we do not control for other factors affecting inequality as the date of occurrence of the pandemic event is likely to be exogenous and uncorrelated to these factors and the error term in equation (1).


Equation (1) is estimated for an unbalanced panel of 175 countries over the period 1961-2017, for each horizon (year) $k=0,\dots,5$. Impulse response functions are computed using the estimated coefficients β^k , and the confidence bands associated with the estimated impulse-response functions are obtained using the estimated standard errors of the coefficients β^k , based on robust standard errors clustered at the country level.

III. DISTRIBUTIONAL IMPACTS OF PANDEMICS

Impacts on Gini coefficients

Figure 1 shows the estimated dynamic response of Ginis to a pandemic event over the five-year period following the event, together with the 90 percent confidence interval around the point estimate. Table 1 reports the associated regressions.

Figure 1. Impact of pandemics on market Gini and net Gini coefficients (%)



Notes: Impulse response functions are estimated using a sample of 175 countries over the period 1961-2017. The graph shows the response and 90 percent confidence bands. The x-axis shows years (k) after pandemic events; $t = 0$ is the year of the pandemic event. Estimates based on $y_{i,t+k} - y_{i,t-1} = \alpha_i^k + \gamma_t^k + \beta^k D_{i,t} + \theta^k X_{i,t} + \varepsilon_{i,t+k}$. $y_{i,t}$ is the log of the Gini coefficient for country i in year t ; α_i are country fixed effects; γ_t are time fixed effects; $D_{i,t}$ is a dummy variable indicating a pandemic event that affects country i in year t . $X_{i,t}$ is a vector that includes two lags of the dependent variable and the pandemic dummy. See Table A2 for the full list of pandemic events. Standard errors in parentheses are clustered at the country level.

Table 1. Impact of pandemics on market Gini and net Gini coefficients

Panel A: Market Gini						
	k=0	k=1	k=2	k=3	k=4	k=5
$D_{i,t}$	0.0683 (0.0781)	0.142 (0.130)	0.132 (0.161)	0.174 (0.219)	0.595** (0.288)	0.658* (0.379)
$D_{i,t-1}$	0.0545 (0.0563)	0.118 (0.0889)	0.157 (0.151)	0.383 (0.245)	0.473 (0.341)	0.869* (0.474)
$D_{i,t-2}$	0.0699 (0.0685)	0.106 (0.137)	0.218 (0.218)	0.300 (0.305)	0.671* (0.399)	0.831 (0.532)
$\Delta y_{i,t-1}$	0.550*** (0.0457)	0.966*** (0.0908)	1.287*** (0.114)	1.456*** (0.147)	1.592*** (0.174)	1.745*** (0.178)
$\Delta y_{i,t-2}$	0.102*** (0.0275)	0.162*** (0.0612)	0.156* (0.0845)	0.194* (0.104)	0.218 (0.135)	0.186 (0.149)
Observations	4,771	4,596	4,421	4,247	4,075	3,906
R ²	0.563	0.576	0.567	0.556	0.559	0.567

Panel B: Net Gini						
	k=0	k=1	k=2	k=3	k=4	k=5
$D_{i,t}$	0.0844 (0.0764)	0.216 (0.148)	0.344* (0.206)	0.576** (0.282)	1.248*** (0.363)	1.283*** (0.443)
$D_{i,t-1}$	0.104* (0.0614)	0.303** (0.123)	0.524** (0.207)	1.096*** (0.320)	1.186*** (0.392)	1.677*** (0.574)
$D_{i,t-2}$	0.145 (0.0926)	0.303* (0.180)	0.659** (0.284)	0.760** (0.362)	1.047** (0.522)	1.125* (0.668)
$\Delta y_{i,t-1}$	0.590*** (0.0428)	1.005*** (0.0896)	1.336*** (0.121)	1.480*** (0.168)	1.588*** (0.192)	1.689*** (0.216)
$\Delta y_{i,t-2}$	0.0520** (0.0249)	0.0723 (0.0586)	-0.00246 (0.0859)	-0.0222 (0.123)	-0.0444 (0.151)	-0.0924 (0.174)
Observations	4,771	4,596	4,421	4,247	4,075	3,906
R ²	0.534	0.521	0.498	0.476	0.473	0.477

Note: Estimates are obtained using a sample of 175 countries over the period 1961-2017, and based on $y_{i,t+k} - y_{i,t-1} = \alpha_i^k + \gamma_t^k + \beta^k D_{i,t} + \theta^k X_{i,t} + \varepsilon_{i,t+k}$. $y_{i,t}$ is the log of the Gini coefficient for country i in year t ; α_i are country fixed effects; γ_t are time fixed effects; $D_{i,t}$ is a dummy variable indicating a pandemic event that affects country i in year t . $X_{i,t}$ is a vector that includes two lags of the dependent variable and the pandemic dummy. See Table A2 for the full list of pandemic events. Standard errors in parentheses are clustered at the country level. *** p<0.01, ** p<0.05, * p<0.1. Country and time fixed effects included but not reported.

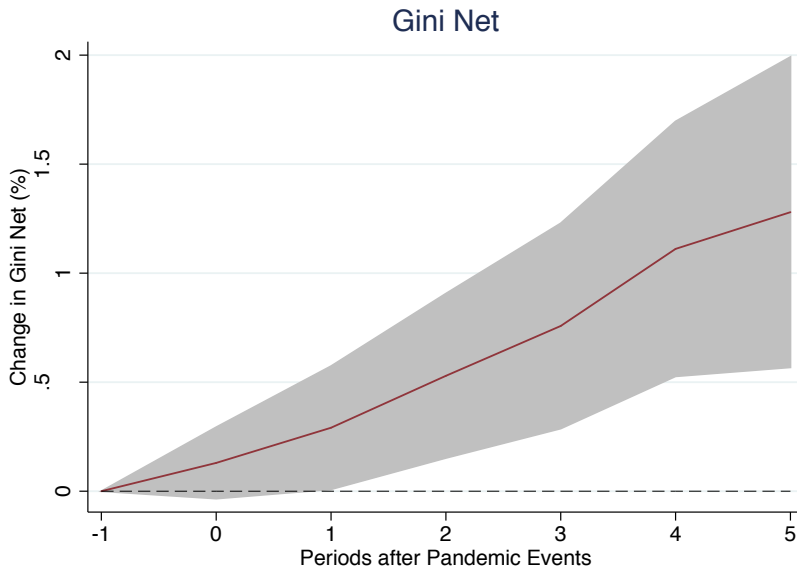
Pandemics lead to a persistent increase in inequality, with the impact being stronger in the case of the net Gini. Five years after the pandemic, both the market and net Gini are above the pre-shock trends by about 0.75% and 1.25%, respectively. Given that the Ginis are very slow-moving variables, these are quantitatively important effects, particularly since the Gini coefficient changes slowly over time—the effect corresponds to approximately $\frac{1}{2}$ standard deviation of the average change of the Gini in the sample.

The fact that the impact on the net Gini is larger than that on the market Gini is somewhat surprising and suggests that policies undertaken to address previous pandemics may actually have been regressive, especially in the medium term, though further analysis would be needed to confirm such a conclusion.

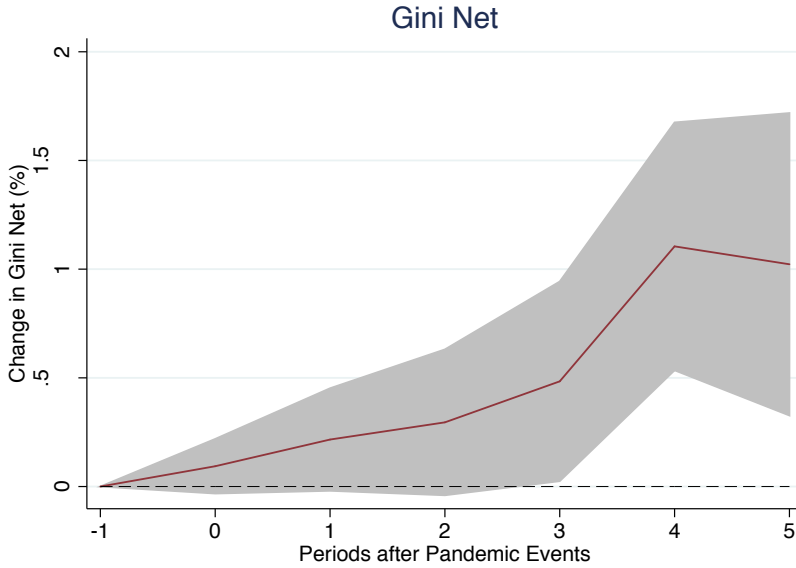
We have carried out several robustness checks of these findings. Here, we report the main three. First, we used as an alternative regression strategy the autoregressive distributed lag model (ADL), as in Romer and Romer (2010) and Furceri et al. (2019). The results in Figure 2 for the net Gini are very similar to those obtained in the baseline using the local projection method.

The second robustness check is to include several control variables in the regression—such as proxies for the level of economic development, demographics, and measures of trade and financial globalization. The results are reported in Figure 3 and are very similar to, and not statistically different from, the baseline.


Finally, since the episodes we considered have occurred in the post 2000 period, we replicated the analysis for this restricted sample. The results presented in Figure 4 are fairly similar to that for the full sample period, except that there is some attenuation in the impact.

Figure 2. Impact of pandemics on net Gini coefficients (%)—ADL

Notes: Impulse response functions are estimated using a sample of 175 countries over the period 1961-2017. The graph shows the response and 90 percent confidence bands. The x-axis shows years (k) after pandemic events; $t = 0$ is the year of the pandemic event. Estimates based on $\Delta y_{i,t} = \alpha_i + \gamma_t + \beta_k(l)D_{i,t} + \varepsilon_{i,t}$. $y_{i,t}$ is the log of the Gini coefficient for country i in year t ; α_i are country fixed effects; γ_t are time fixed effects; $D_{i,t}$ is a dummy variable indicating a pandemic event that affects country i in year t . See Table A2 for the full list of pandemic events. Standard errors in parentheses are clustered at the country level.

Figure 3. Impact of pandemics on net Gini coefficients (%)—Additional controls

Notes: Impulse response functions are estimated using a sample of 175 countries over the period 1961-2017. The graph shows the response and 90 percent confidence bands. The x-axis shows years (k) after pandemic events; $t = 0$ is the year of the pandemic event. Estimates based on $y_{i,t+k} - y_{i,t-1} = \alpha_i^k + \gamma_t^k + \beta^k D_{i,t} + \theta^k X_{i,t} + \varepsilon_{i,t+k}$. $y_{i,t}$ is the log of the Gini coefficient for country i in year t ; α_i are country fixed effects; γ_t are time fixed effects; $D_{i,t}$ is a dummy variable indicating a pandemic event that affects country i in year t . $X_{i,t}$ is a vector that includes two lags of the dependent variable, the pandemic dummy, the level of GDP, the level of GDP square, population density, the share of population in urban area, the KOF index of trade globalization and the KOF index of financial globalization. See Table A2 for the full list of pandemic events. Standard errors in parentheses are clustered at the country level.

Figure 4. Impact of pandemics on net Gini coefficients (%)—Restricted sample (2000-17)

Notes: Impulse response functions are estimated using a sample of 175 countries over the period 2001-2017. The graph shows the response and 90 percent confidence bands. The x-axis shows years (k) after pandemic events; $t = 0$ is the year of the pandemic event. Estimates based on $y_{i,t+k} - y_{i,t-1} = \alpha_i^k + \gamma_t^k + \beta^k D_{i,t} + \theta^k X_{i,t} + \varepsilon_{i,t+k}$. $y_{i,t}$ is the log of the Gini coefficient for country i in year t ; α_i are country fixed effects; γ_t are time fixed effects; $D_{i,t}$ is a dummy variable indicating a pandemic event that affects country i in year t . $X_{i,t}$ is a vector that includes two lags of the dependent variable and the pandemic dummy. See Table A2 for the full list of pandemic events. Standard errors in parentheses are clustered at the country level.

Heterogeneity across episodes depend on the economic impact of pandemics

As shown by Ma et al. (2020), the impact of pandemic events on economic activity is likely to vary both across episodes and countries. To examine how this heterogeneity in the economic effects affects the distributional consequences of pandemics, we estimated the following equation:

$$y_{i,t+k} - y_{i,t-1} = \alpha_i^k + \gamma_t^k + F(z_{it})[\beta_L^k D_{i,t} + \theta_L^k X_{i,t}] + (1 - F(z_{it}))[\beta_H^k D_{i,t} + \theta_H^k X_{i,t}] + \varepsilon_{i,t+k}$$


$$\text{with } F(z_{it}) = \frac{\exp^{-\gamma z_{it}}}{(\exp^{-\gamma z_{it}} + 1)}, \quad \gamma = 3.5 \quad (2)$$

where z is an indicator of the state of the economy normalized to have zero mean and a unit variance. The indicator of the state of the economy considered in the analysis is GDP growth. 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 β_L^k and β_H^k capture the distributional impact of a pandemic event at each horizon k in cases of pandemics associated with extreme recessions ($F(z_{it}) \approx 1$ when z goes to minus infinity) and booms ($1 - F(z_{it}) \approx 1$ when z goes to plus infinity), respectively.⁶ We choose $\gamma = 3.5$, following Tenreyro and Thwaites (2016).

The results in Figure 5 show that the distributional effect of pandemic events varies with their impact on economic activity. In particular, for episodes associated with significant economic contractions, the effect is statistically significant and larger than the average effect (the medium-term effect on Gini increases from 1.25 to about 2 percent), while it is not statistically significantly different from zero for episodes associated with high growth.

⁶ $F(z_{it})=0.5$ is the cutoff between weak and strong economic activity.

Figure 5. Impact of pandemics on net Gini coefficients (%)—The role of economic conditions associated with pandemic events



Notes: Impulse response functions are estimated using a sample of 175 countries over the period 1961-2017. The graph shows the response and 90 percent confidence bands. The x-axis shows years (k) after pandemic events; $t = 0$ is the year of the pandemic event. The dotted green line denotes the average (unconditional) effect reported in Figure 1. The red lines denote the estimates for pandemic events associated with very low and high growth. Estimates based on $y_{i,t+k} - y_{i,t-1} = \alpha_i^k + \gamma_t^k + F(z_{it})[\beta_L^k D_{i,t} + \theta_L^k X_{i,t}] + (1 - F(z_{it}))[\beta_H^k D_{i,t} + \theta_H^k X_{i,t}] + \varepsilon_{i,t+k}$. $y_{i,t}$ is the log of the Gini coefficient for country i in year t ; α_i are country fixed effects; γ_t are time fixed effects; $D_{i,t}$ is a dummy variable indicating a pandemic event that affects country i in year t . $X_{i,t}$ is a vector that includes two lags of the dependent variable and the pandemic dummy. $F(z_{it})$ is an indicator function of the state of the economy. The coefficients β_L^k and β_H^k capture the distributional impact of a pandemic event at each horizon k in cases of pandemics associated with extreme recessions ($F(z_{it}) \approx 1$ when z goes to minus infinity) and booms ($1 - F(z_{it}) \approx 1$ when z goes to plus infinity), respectively. See Table A2 for the full list of pandemic events. Standard errors in parentheses are clustered at the country level

Impact on other indicators of distribution

To shed some light on the channels through which pandemics affect inequality, we explored the impact of pandemic events on income shares and employment outcomes for various educational groups. Due to data limitations, these results are for a much smaller set of countries than those for the Gini results.

The results for the impact of pandemics on the shares of income by decile are shown in Figure 6. It is evident that the impact is to raise the shares of the upper-income deciles and reduce those of the lower-income deciles. The impacts are quantitatively significant. For instance, in our sample, the share of income going to the top two deciles is 46 percent on average; five years after the pandemic, this share increases to nearly 48 percent. The share of income going to the bottom two deciles is only 6 percent; five years after the pandemic, this share falls further to 5.5 percent.

Figure 7 shows the vastly disparate impact that pandemics have on the employment of people with different levels of educational attainment. Those with advanced or intermediate levels of education are scarcely affected, whereas the employment to population ratio of those with basic levels of education falls significantly, by more than 5 percent in the medium term.

IV. CONCLUSION

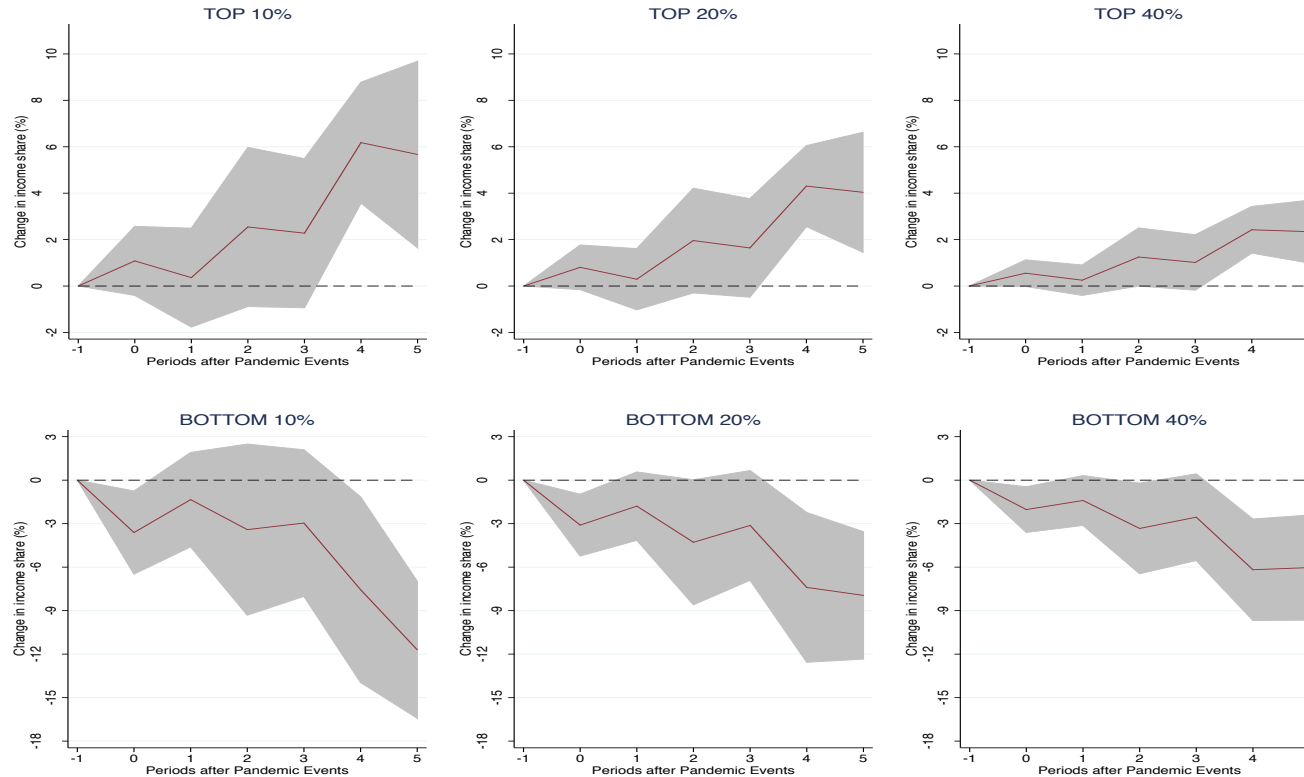
The Covid-19 crisis is showing how the more vulnerable socio-economic groups suffer from a greater risk of financial exposure, and also from greater health risks, and worse housing conditions during the lockdown period. These factors may potentially exacerbate inequalities.

This paper explores this possibility by providing evidence on the impact of pandemics and major epidemics from the past two decades on income distribution. Our results justify the concern that, in the absence of policies aimed at protecting the most vulnerable, the pandemic could end up exerting a significant adverse impact on inequality: past events of this kind, even though much smaller in scale, have led to increases in the Gini coefficient, raised the income shares accruing to the higher deciles of the income distribution, and lowered the employment-to-population ratio for those with basic education compared to those with higher education. In addition, the result that the inequality effect increases with the negative effect of pandemic events on economic activity

suggests that the distributional consequences of Covid-19 may be larger than those in previous pandemic episodes, all else equal.

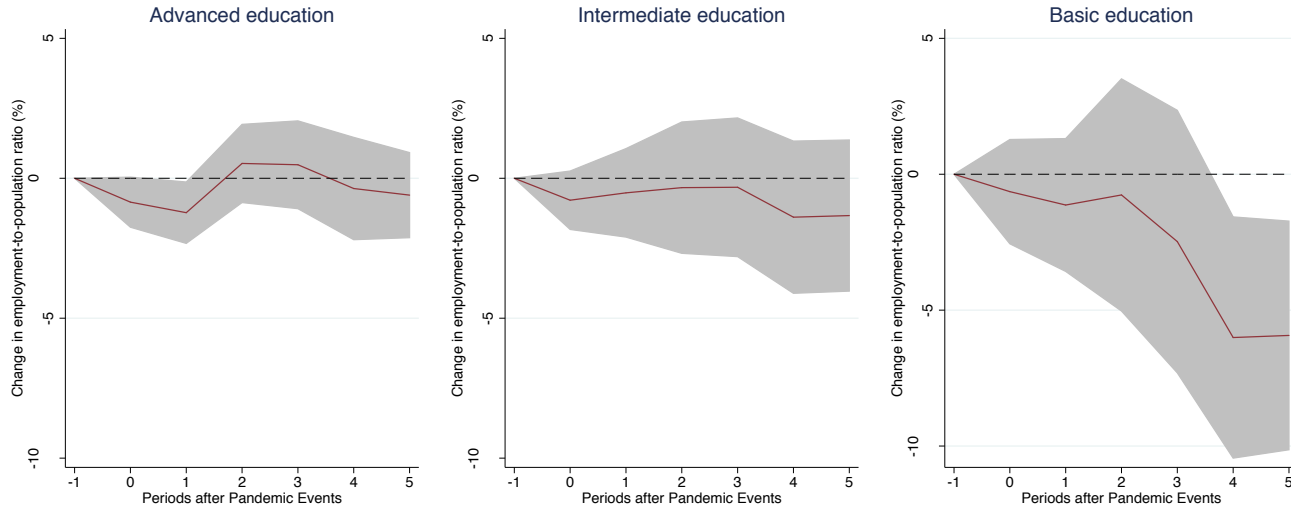
Our results leave several questions open for future research. First, the distributional effects of pandemic events are likely to vary considerably across countries, depending on country-specific characteristics, initial income distribution conditions as well as the stringency of containment measures. Second, there is growing evidence that the economic effects of Covid-19 may also vary between different segments of the population in terms of race, age, and gender. Third, the human cost of pandemics is also sadly higher in low-income groups, which are more prone to diseases and have often more limited access to health services.

Figure 6. Impact of pandemics on shares of income, by deciles



Notes: Impulse response functions are estimated using a sample of 64 countries over the period 1981-2017. The graph shows the response and 90 percent confidence bands. The x-axis shows years (k) after pandemic events; $t = 0$ is the year of the pandemic event. Estimates are based on $y_{i,t+k} - y_{i,t-1} = \alpha_i^k + \gamma_t^k + \beta^k D_{i,t} + \theta^k X_{i,t} + \varepsilon_{i,t+k}$. $y_{i,t}$ is, in turn, the log of the income share held by the top (bottom) 20% (10%) for country i in year t ; α_i are country fixed effects; γ_t are time fixed effects; $D_{i,t}$ is a dummy variable indicating a pandemic event that affects country i in year t . $X_{i,t}$ is a vector that includes two lags of the dependent variable and the pandemic dummy. See Table A2 for the full list of pandemic events. Standard errors in parentheses are clustered at the country level.

Figure 7. Impact of pandemics on employment-to-population ratio, by education level



Notes: Impulse response functions are estimated using a sample of 76 countries over the period 1990-2017. The graph shows the response and 90 percent confidence bands. The x-axis shows years (k) after pandemic events; $t = 0$ is the year of the pandemic event. Estimates are based on $y_{i,t+k} - y_{i,t-1} = \alpha_i^k + \gamma_t^k + \beta^k D_{i,t} + \theta^k X_{i,t} + \varepsilon_{i,t+k}$. $y_{i,t}$ is, in turn, the log of employment-to-population ratio by education level for country i in year t ; α_i are country fixed effects; γ_t are time fixed effects; $D_{i,t}$ is a dummy variable indicating a pandemic event that affects country i in year t . $X_{i,t}$ is a vector that includes two lags of the dependent variable and the pandemic dummy. See Table A2 for the full list of pandemic events. Standard errors in parentheses are clustered at the country level.

Covid Economics 12, 1 May 2020: 138-157

References

- Atkeson, A. (2020): “What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios,” NBER Working Paper No. 26867.
- Barro, R. J., J. F. Ursua, and J. Weng (2020). “The Coronavirus and the Great Influenza Pandemic: Lessons from the “Spanish Flu” for the Coronavirus’s Potential Effects on Mortality and Economic Activity,” NBER Working Paper No. 26866.
- Coibion, O. Y. Gorodnichenko and M. Weber (2020). “Labour markets during the Covid-19 crisis: A preliminary view. NBER Working Paper No. 27017
- De Haan, Jakob, and Jan-Egert Sturm (2017). “Finance and Income Inequality: A Review and New Evidence.” *European Journal of Political Economy*, 50, 171-195.
- Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020): “The Macroeconomics of Epidemics,” NBER Working Paper No. 26882.
- Furceri, D., P. Loungani and J.O. Ostry, 2019. The Aggregate and Distributional Effects of Financial Globalization: Evidence from Macro and Sectoral Data, *Journal of Money, Credit and Banking*, 51, 163-198.
- ILO (2020). ILO Monitor 2nd edition: COVID-19 and the world of work.
- Jordà, O. (2005): “Estimation and inference of impulse responses by local projections,” *American Economic Review*, 95, 161–182.
- Jordà, O. S.R. Singh and A.M. Taylor (2020). “Pandemics: Long-Run Effects”, *Covid Economics*, 1, 1-15.
- Ma, C., J. Rogers and S. Zhou (2020) “Global Financial Effects”. *Covid Economics*, 5, 56-78.
- Schmitt-Grohé, S. K. Teoh and M. Uribe (2020). “COVID-19: Testing Inequality in New York City”, *Covid Economics*, 8, 27-43.
- Romer, C. D. and D.H. Romer (2010). “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks.” *American Economic Review*, 100(3), 763-801.
- Tenreyro, S. and G. Thwaites (2016). “Pushing on a String: US Monetary Policy Is Less Powerful in Recessions.” *American Economic Journal: Macroeconomics*, 8 (4): 43-74.
- World Bank (2020). Poverty and Distributional Impacts of COVID-19: Potential Channels of Impact and Mitigating Policies BRIEF APRIL 16, 2020, World Bank.

APPENDIX

Table A1. Data Sources and Descriptive Statistics

Variable	Source	Obs	Mean	Std. Dev.	No. of Countries
Gini Market	SWIID 8.2	5,305	45.28	6.59	175
Gini Net	SWIID 8.2	5,305	38.33	8.76	175
Top 40% Income Share	WDI	1,444	67.77	6.65	64
Top 20% Income Share	WDI	1,444	46.28	7.83	64
Top 10% Income Share	WDI	1,444	30.85	7.31	64
Bottom 40% Income Share	WDI	1,444	17.12	4.56	64
Bottom 20% Income Share	WDI	1,444	6.31	2.19	64
Bottom 10% Income Share	WDI	1,443	2.44	1.02	64
<i>Employment/Population (E/P) ratios</i>					
E/P ratio – Basic Education	ILO	1,340	42.51	16.22	76
E/P ratio – Intermediate Education	ILO	1,333	61.03	9.23	76
E/P ratio – Advanced Education	ILO	1,338	75.14	7.60	76

Table A2. List of Pandemic and Epidemic Episodes

Starting year	Event Name	Affected Countries	Number of countries
2003	SARS	AUS, CAN, CHE, CHN, DEU, ESP, FRA, GBR, HKG, IDN, IND, IRL, ITA, KOR, MNG, MYS, NZL, PHL, ROU, RUS, SGP, SWE, THA, TWN, USA, VNM, ZAF	27
2009	N1H1	AFG, AGO, ALB, ARG, ARM, AUS, AUT, BDI, BEL, BGD, BGR, BHS, BIH, BLR, BLZ, BOL, BRA, BRB, BTN, BWA, CAN, CHE, CHL, CHN,CIV, CMR, COD, COG, COL, CPV, CRI, CYP, CZE, DEU, DJI, DMA, DNK, DOM, DZA, ECU, EGY, ESP, EST, ETH, FIN, FJI, FRA, FSM, GAB, GBR, GEO, GHA, GRC, GTM, HND, HRV, HTI, HUN, IDN, IND, IRL, IRN, IRQ, ISL, ISR, ITA, JAM, JOR, JPN, KAZ, KEN, KHM, KNA, KOR, LAO, LBN, LCA, LKA, LSO, LTU, LUX, LVA, MAR, MDA,MDG, MDV, MEX, MKD, MLI, MLT, MNE, MNG, MOZ, MUS, MWI, MYS, NAM, NGA, NIC, NLD, NOR, NPL, NZL, PAK,PAN, PER, PHL, PLW, PNG, POL, PRI, PRT, PRY, QAT, ROU, RUS, RWA, SAU, SDN, SGP, SLB, SLV, STP, SVK, SVN, SWE, SWZ, SYC, TCD, THA, TJK, TON, TUN, TUR, TUV, TZA, UGA, UKR, URY, USA, VEN, VNM, VUT, WSM, YEM, ZAF, ZMB, ZWE	148
2012	MERS	AUT, CHN, DEU, EGY, FRA, GBR, GRC, IRN, ITA, JOR, KOR, LBN, MYS, NLD, PHL, QAT, SAU, THA, TUN, TUR, USA, YEM	22
2014	Ebola	ESP, GBR, ITA, LBR, USA	5
2016	Zika	ARG, BOL, BRA, CAN, CHL, COL, CRI, DOM, ECU, HND, LCA, PAN, PER, PRI, PRY, SLV, URY, USA	18
Total Pandemic and Epidemic Events			220

Source: Based on Ma and others (2020).

The global impact of COVID-19 on fintech adoption¹

Jonathan Fu² and Mrinal Mishra³

Date submitted: 22 April 2020; Date accepted: 26 April 2020

We draw on mobile application data from 74 countries to document the effects of the COVID-19 pandemic on the adoption of digital finance and fintech. We estimate that the spread of COVID-19 and related government lockdowns have led to between a 24 and 32 percent increase in the relative rate of daily downloads of finance mobile applications in the sample countries. In absolute terms, this equates to an average daily increase of roughly 5.2 to 6.3 million application downloads and an aggregate increase of about 316 million app downloads since the pandemic's outbreak, taking into account prior trends. Most regions across the world exhibit notable increases in absolute, relative, and per capita terms. Preliminary analysis of country-level characteristics suggest that market size and demographics, rather than level of economic development and ex-ante adoption rates, drive differential trends.

- 1 We thank Annette Krauss for helpful comments and suggestions. Mishra gratefully acknowledges financial support from the European Research Council under the European Union's Horizon 2020 research and innovation programme. ERC ADG 2016 (No. 740272).
- 2 Postdoctoral researcher, Department of Banking & Finance, University of Zurich and The Center for Sustainable Finance and Private Wealth.
- 3 Ph.D. candidate, Department of Banking & Finance, University of Zurich and The Swiss Finance Institute.

1 Introduction

Shocks of various kinds can drive technological adoption in unanticipated ways. Moreover, these shocks can also result in longer term changes to societies and economies. The COVID-19 pandemic started out as a shock to public health and healthcare systems. However, the nature of the pandemic coupled with the speed of transmission have required societies to adopt “social-distancing” or more stringent lockdown measures imposed by the government. This has been the case both in areas affected and unaffected by the actual spread of COVID-19. Despite the high human and economic costs following its spread,¹ COVID-19 has resulted in some beneficiaries and possible silver linings. There is increasing anecdotal evidence that the technology sector and, in particular, companies enabling communication and exchange of goods and services over distance have seen notable increases in adoption and usage.² Such services have proven essential in helping many households and other firms to mitigate some of the health risks and adverse socioeconomic effects of the pandemic and allowing some critical aspects of day-to-day life to continue as normal.

In this paper, we provide evidence on the acceleration of digital transformation in the financial sector by documenting the early impact of COVID-19’s spread on fintech adoption worldwide. To do so, we draw on historical and real-time data on mobile application downloads, which provides us with a high dimensional measure that allows us to detect changes on the extensive margin within a small time frame. The use of such mobile application-based services intuitively provides an attractive option for accessing financial services particularly during a pandemic—given self- or government-imposed restrictions to movement and social distancing measures, and risk of contamination via physically handling cash (Arner et al. (2020)). Moreover, the nature of the fintech market generally gives providers greater flexibility to quickly deploy new products and services or adapt or scale existing ones, in response to shocks (Fuster et al. (2019)). We extract daily information on all finance mobile app downloads from January 1st, 2019 to present³ for 74 countries available and covering both the Android and iOS markets.⁴ The countries in our sample cover all global regions and account for approximately 80 percent of the global population and over 90 percent of the world economy in terms of nominal GDP.

¹The global death toll as of April 20th is at over 165K. Due to the curbs on gathering in public places and travel, a gamut of sectors such as transportation, hospitality, and leisure have been particularly hard hit. <https://www.brookings.edu/blog/the-avenue/2020/03/17/the-places-a-covid-19-recession-will-likely-hit-hardest/>

²See for example: <https://www.economist.com/briefing/2020/04/11/the-changes-covid-19-is-forcing-on-to-business>.

³Unless otherwise noted, the analysis in this paper uses data up until April 20th, 2020.

⁴While not comprehensive of all fintech activity, we believe this provides a reasonable proxy for capturing its spread, since an increasing share of fintech providers (both newer startups and traditional incumbents) establish a customer-provider relationship and deliver their services via mobile applications. This is particularly the case given the increasing pervasiveness of smartphones and mobile internet coverage, including in lower- and middle-income economies. At the end of 2019, 90 percent of the world’s population was covered by mobile networks, 67 percent of the global population (5.2 billion people) had a subscription to mobile services, and 49 percent of the global population (3.8 billion people) were mobile internet users (GSMA (2020)).

Using these data, we apply panel data regression models controlling for seasonality at the aggregate- or country-level to estimate the change in fintech app adoption pre- and post-COVID-19. We estimate that the spread of COVID-19 and related government lockdowns have led to between a 24 and 32 percent increase in the relative rate of daily downloads. In absolute terms, and accounting for the average number of days since confirmed cases or lockdowns in our sample, this equates to an aggregate increase of about 316 million app downloads since the pandemic's outbreak. We then test for signs of differential trends across regions. While certain regions do exhibit larger increases than others, we find that all global regions exhibit significantly higher demand compared to pre-COVID-19 trends. The one notable exception is Europe, where there has been a flat or slight decrease in demand. Finally, we further conduct some exploratory analyses on country-level characteristics that may drive differential trends. Our results generally do not suggest that countries' economic development level or ex-ante rates of adoption are particularly predictive of post-COVID 19 adoption. Rather, and in line with some existing literature, it suggests factors related to market size (Foster & Rosenzweig (2010)) and demographics (Carlin et al. (2017)), which may point to the importance of network effects.

Our paper relates primarily to literature on fintech adoption and the role of shocks in driving acceleration. Existing studies have highlighted that idiosyncratic shocks *within countries*—affecting either the demand or supply-side—can have immediate and longer-term effects on local rates of fintech adoption (Higgins (2019), Crouzet et al. (2019), Chava et al. (2018)).⁵ These studies tend to emphasize how network effects may either amplify or lessen rates of adoption during a shock. For example, as more individuals and businesses use digital payments and other digital products, the benefits to all users increase, creating externalities that can further accelerate adoption. A main implication noted by Crouzet et al. (2019) is that whether a shock leads to a persistently higher or only transitory increase in adoption is thus dependent on both the magnitude of the shock and the ex-ante state of adoption.

To date, however, there are no studies to our knowledge that look at how global-level shocks occurring *across countries* may similarly play out and either lead to convergence in their fintech adoption rates or exacerbate existing differences. This is a particular concern due to an existing “digital gap” that exists between many lower and middle-income economies and advanced economies (GSMA (2020)). Given the broad range of potential benefits (and risks)

⁵“Fintech” can be broadly defined as technologically-enabled financial innovation that results in new business models, applications, processes, products, or services with an associated material effect on financial markets and institutions and the provision of financial services (Schindler (2017)). Typically, it refers to innovators and disruptors in the financial sector that make use of the availability of information communication technology—particularly, Internet-related technologies (e.g., cloud computing, mobile Internet)—and data processing technologies (e.g., predictive analytics, automation, and machine learning). As noted by Gomber et al. (2017), this may come either from “disruptive fintech”—new start-ups and larger technology companies (“BigTech”) offering new products and services that challenge incumbent traditional providers, or from “sustaining fintech”—established financial services providers that try to protect their market position by similarly adopting and innovating in use of information technologies.

from more widespread adoption of digital finance and fintech—particularly for mitigating effects from major shocks such as COVID-19—this may have important implications both in the short- and long-term. For example, consumers and retail firms generally benefit from fintech adoption by gaining access to more affordable, responsive, and tailored banking services (Das (2019)); particularly those who have been traditionally unbanked or underbanked (Demirguc-Kunt et al. (2018), GPMI (2016)).⁶ Financial service providers typically reduce operational costs significantly and can use novel processes to better assess credit risk ex-ante and reduce delinquency ex-post.⁷ Governments, monetary authorities, and financial regulators are also expected to benefit from increased efficiency of tax collection and subsidy payment, combating inflation via reduced circulation of excess or bad money, and gaining more diverse tools for reviving credit markets (Manyika et al. (2016), BIS (2018)). Finally, at the macroeconomic-level, large-scale aggregate increases in financial access may generate productivity gains across the economy and help spur economic growth (Manyika et al. (2016)), or in the case of COVID-19, mitigate the expected freefall.

More generally, comparable cross-country data on many aspects of fintech adoption are still relatively limited and, hence, comparative research on drivers of adoption across countries is still scarce. Among the few comparative studies in this area, Claessens et al. (2018) use a static estimate of the volume of fintech credit per capita in the economy for 63 countries and run cross-sectional regressions to find that the size of an economy’s fintech credit market is positively related to its income level, and negatively related to the competitiveness of its banking system and the stringency of its banking regulation. Frost (2020), meanwhile, similarly draws on cross-sectional evidence from a household survey of 38 countries by EY (2019) to note the important role of country demographics on predicting broader “digital finance” adoption—defined as active use of multiple categories of fintech products and services. Both sets of authors (Claessens et al. (2018) and Frost (2020)) note that measuring fintech activity and, in particular, having a sense of real-time trends, can be difficult given the diversity and constantly evolving landscape of fintech providers, the small size of many of the platforms, and because many of the providers still lie outside of prudential regulatory reporting requirements. Cross-country demand-side surveys such as the Global Findex (Demirguc-Kunt et al. (2018)) may offer important insights to fill this gap, but are relatively infrequent given their high cost. Thus, particularly for measuring outcomes from a crisis on fintech activity, there is an important need for obtaining high-frequency data

⁶A large body of impact evaluations have provided consistent evidence that digital payments, remittances, and insurance are particularly useful at mitigating effects from exogenous shocks by expanding households networks for safety-nets and risk-pooling (Jack et al. (2013), Jack & Suri (2014), and Bharadwaj et al. (2019)). (Albeit, the implications of such studies are less clear in the case of a shock such as COVID-19, which has affected most of the world, both across- and within-countries.)

⁷For example, many fintech providers use digital trails to better assess risk, collect digital repayments on an automated basis, and reduce loan delinquency via SMS or mobile app “nudges” to prompt borrowers when they have missed a payment. This has allowed them to underwrite loans and insurance policies for clients deemed too risky for traditional providers.

that is comprehensive of all types of providers.

Building on the extant literature, we focus our analysis on testing the following preliminary hypotheses:

- *H₁: The COVID-19 pandemic has increased adoption of finance mobile applications on the extensive margin.*
- *H₂: The effect on adoption is smaller in countries with higher levels of economic development and with higher rates of ex-ante adoption and larger in countries with larger market size (population and demographics).*
- *H₃: The effect on adoption is primarily driven by government restrictions rather than the spread of COVID-19 itself.*

The remainder of the paper proceeds as follows. Section 2 describes the study's data sources. Section 3 presents the paper's main analyses and results. Section 4 discusses broader implications and concludes.

2 Data

2.1 Data Sources

To obtain a reasonable estimate of fintech adoption driven by COVID-19, we need a high-frequency measure that allows us to detect changes within a small time frame. To do so, we draw on mobile app download data from the AppTweak platform.⁸ The platform contains historical and real-time data on mobile app downloads at the aggregate- or country-level for all app categories in the Android and iOS markets. We extract daily information on all finance category mobile apps downloads from January 1st, 2019 to present for the 74 countries available through the platform.⁹ (See Appendix Table A.1 for a full country list and additional details.) These data are pulled at the country-level, which allows us to aggregate up to broader regions, as required for a given analysis. Note that preliminary data investigation revealed some day-to-day volatility in downloads—particularly, for some of the smaller countries or apps. Thus, we calculate a leading and equally-weighted 14-day moving average for the number of downloads to smooth

⁸<https://www.apptweak.com/>

⁹The categorization of finance apps are currently taken directly from the platform as listed by the mobile applications' developers. The advantages of drawing on the self-categorization is that it provides a way to capture real-time and comprehensive data on all providers in the market, of any size. The disadvantage is that it is cumbersome to control whether all applications in the entire platform are necessarily correctly categorized, given its comprehensive scale. However, as the platform allows us to also target individual financial service providers of interest and their particular apps, we plan to use this more granular level data in our next iteration to better understand what specific provider and product types are driving results.

these day-to-day fluctuations. This is used for our descriptive figures and serves as the main dependent variable for our empirical analysis.¹⁰

Our main explanatory variables are proxies for the spread of COVID-19 and related government policy. For the former, we draw on data from the Oxford COVID-19 Government Response Tracker (OxCGRT).¹¹ OxCGRT provides daily updates on the number of confirmed cases and confirmed deaths, as well as various different types of policies being implemented at the country-daily level (e.g., school closures, workplace closures, travel restrictions, etc.). These policies are used by OxCGRT to construct a policy “stringency index”. For the latter, we draw on data from Kaggle, which provides the date of each country’s lockdowns implemented due COVID-19, if applicable.¹² Finally, we also merge data from the World Bank’s World Development Indicator database and the 2017 Global Findex database for some of our exploratory analyses on predictors of differential post-COVID-19 trends in adoption.

2.2 Descriptive Statistics and Trends

Figure 1 presents the average daily downloads for fintech apps across our full country sample. We note a clear uptrend in downloads starting from around mid-February. To see which mobile operating system dominates the results, we split the total downloads to Android and iOS apps and observe that the uptick seems to be mainly driven by the Android market with a muted or almost no effect for the iOS market. The figure shows a heightened adoption of fintech apps and a greater push towards digitalization during the pandemic. This affirms the general notion that exogenous shocks to the economy or society end up accelerating trends which would have otherwise played out in a more protracted fashion.

Figure 2 depicts a scatter plot of daily downloads and COVID-19 cases across countries. We remove the time dimension from our data by aggregating all the data from 2020. Subsequently, we also plot the best fit line and observe a general upward trend. The figure gives us a high-level understanding of how the spread of COVID-19 is associated with fintech adoption. Initial evidence, as depicted by Figure 2 suggests there is a strong positive correlation between both. A crude analysis suggests that countries above the best fit line show increased fintech adoption whereas those below the line show lower adoption owing to the effects of the pandemic. A myriad of reasons like the intensity of enforced lockdowns, extent of previous adoption of fintech, or access to mobile internet and smartphones could be instrumental in driving these effects. However, as Figure 6 in the Appendix shows, the correlation between lockdown intensity and

¹⁰We conduct sensitivity tests on our main results using the original raw download numbers or applying 10-, 20-, and 30-day moving averages in lieu of the 14-day window. Results are provided in Appendix Tables A.2 and A.3 and help demonstrate that the main results and takeaways are largely unchanged.

¹¹<https://www.bsg.ox.ac.uk/research/research-projects/coronavirus-government-response-tracker>

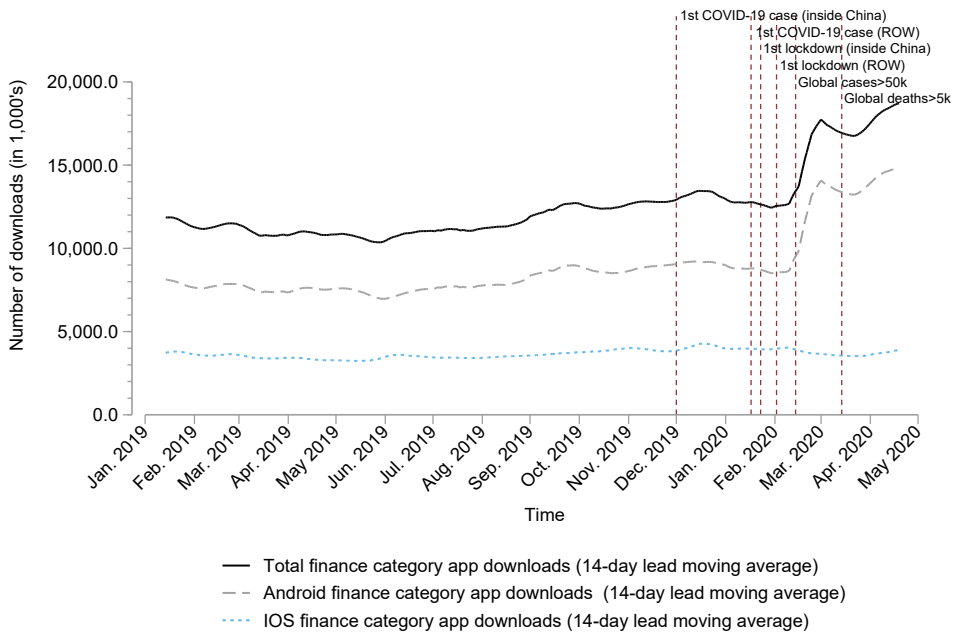
¹²<https://www.kaggle.com/jcyzag/covid19-lockdown-dates-by-country>.

app downloads is relatively weaker than the relationship in Figure 2, which may suggest that adoption is not primarily driven by the exogenously-imposed lockdowns instituted in countries.

Further descriptive statistics at the aggregate- and country-level are provided in Tables 1 and 2, respectively.

Figure 1: The Impact of COVID-19 on the Adoption of Fintech Mobile Apps

This figure depicts the daily number of downloads for finance category mobile applications across the iOS and Android platforms from 74 countries. The sample data covers the period from January 1st, 2019 to April 20th, 2020.



Covid Economics 12, 1 May 2020: 158-191

3 Results

3.1 What is the Estimated Impact of the Spread of COVID-19 on Fintech Adoption?

To provide a concise way of estimating the effect of COVID-19 on the magnitude of fintech adoption during our study period, we use an empirical specification of the following form:


$$y_{it} = \beta_0 + \beta_1 COVID19_{it} + \theta_i + \gamma_m + u_{it} \tag{1}$$

where the dependent variable y is either 1) the absolute number of daily downloads or 2) the relative (logarithmic) growth in daily downloads for location i and at time (day) t .¹³ Note that

¹³We also calculate and run comparable results using daily downloads per capita. The results are provided in

Figure 2: Scatterplot of Finance Category Mobile App Downloads vs. Confirmed COVID-19 Cases

This figure depicts the relationship between countries' average daily number of downloads for finance category mobile applications and average daily confirmed new cases of COVID-19. We roughly restrict to the general period after the first signs of the COVID-19 outbreak (i.e., since January 1st, 2020). The underlying data includes mobile apps from both the iOS and Android platforms for 74 countries. Country codes are listed in Appendix Table A.1.



location i can refer to either an individual country or a region, as required for a given analysis. $COVID19$ denotes a dummy that is set to one after time t , either when location i had its first confirmed cases (*Post-Confirmed case*) or entered into lockdown (*Post-Lockdown*), respectively. We include location-level dummies (θ_i) to control for country or region fixed effects, and month of the year dummies (denoted by γ_m) to adjust for aggregate seasonal trends in the finance app market. We alternatively interact the location and month dummies in our preferred specification to control for more disaggregated seasonality in downloads (e.g., country- or region-level time trends, which would substitute $(\theta_i \times \gamma_m)$ in lieu of the location and month dummies.¹⁴

Tables 3 and 4 present results for the country-level model estimating the impact of COVID-19 on fintech adoption using the *Post-Confirmed case* or *Post-Lockdown* variables, respectively.¹⁵

Appendix Tables A.4 and A.5.

¹⁴For example, preliminary data exploration reveals that the United States experiences a significant uptick annually in finance-related app downloads every February and March prior to its tax-filing deadlines in April. As this happens to coincide with the timing of the spread of COVID-19 cases and lockdown in the United States, applying the country-month interactions helps mitigate confounding these seasonality trends and overestimating the change from COVID-19.

¹⁵In Appendix Tables A.2 and A.3, we present results from alternate constructions of the main explanatory variables by altering the definition of app downloads. For our baseline, we use a 14-day moving average of app

The results across both the tables are generally quite consistent and show that the COVID-19 outbreak led to an absolute and relative increase in fintech adoption, even after controlling for country fixed characteristics and aggregate- or country-level seasonality in app downloads. The coefficients depicted in Tables 3 and 4 are in terms of average per country. In absolute terms, we estimate that, on average, a given country saw a daily increase of about 88,000 and 76,000 finance mobile app downloads following the date of its first confirmed case and lockdown, respectively, taking into account prior trends. (These figures are based on our preferred specifications in columns 4, which control for country-level seasonality.)

If we aggregate the effects across the 74 countries in our full sample, of which 72 have had a confirmed COVID-19 case and 68 have entered lockdown (see Table 1 or Appendix Table A.1), the economic size can be interpreted as an increase of approximately 5.2 to 6.3 million finance app downloads per day “globally” driven by COVID-19. Furthermore, accounting for the average number of days since confirmed COVID-19 cases or lockdowns in our sample (see Table 2), this equates to a aggregate increase of 316 and 135 million finance app downloads, respectively. In relative terms, this is an increase of around 32 percent since countries’ first confirmed COVID-19 case and around 24 percent since the start of COVID-19 related lockdowns, compared to prior trends. (Figures based on Tables 3 and 4, columns 8.)

3.2 Are All Regions Equally Affected?

The spread of COVID-19 has not been uniform across all continents and regions. Countries in Asia, Europe, and North America have generally been hit earlier and—thus far—harder than in Africa, South America, the Middle East, and Oceania. As a result, the adoption in fintech might also show some heterogeneity depending on the duration or intensity of the pandemic in certain regions. This analysis become important as some areas like Europe and USA-Canada are free movement zones for capital and labour. As a preliminary investigation into region-level effects, we re-construct our main variables at the region level and then re-run Equation (1). We interact the *COVID19* dummy with the respective region dummies to test for the differential effects of COVID-19 on fintech adoption rates.

Table 5 depicts results from this region-level analysis. Europe serves as the base category in the results. We can observe that the countries in our sample from North America and Asia exhibit the largest increase both in absolute and relative terms and are responsible for most of the adoption on the extensive margin. That having been said, the countries in our sample from Africa, Middle East, Oceania, and South America also see fairly large increases in either

downloads. However, to ensure that our results are not overly sensitive to the choice of the number of days in the moving average calculation, we alter this number to 10, 20 and 30 days and observe that our main results still hold.

absolute or relative terms. In our preferred specifications that control for regional seasonality (in columns 6 and 8), their estimated relative increases fall in the range between 15 and 39 percent, post-confirmed cases, and between 33 and 44 percent, post-lockdown. The European countries in our sample are the notable exception, where we observe that, after an initial minor uptick after cases began to spread, there was a flat or slight decrease in rates of downloads following the spread of government-instituted lockdowns.

Diving a level deeper into our underlying data, we observe that the post-COVID-19 increases appear particularly driven by the United States, India, and Mexico in absolute terms, but that a range of other countries, such as Indonesia, Turkey, Russia, the Netherlands, and Brazil in very recent weeks also experience large relative increases. Apart from the major economies and those with larger populations, a similar aggregate trend exists for the remaining countries in our sample. Meanwhile, we can also observe that Germany and Italy experienced large decreases and play a key role in driving the negative result for Europe. See Appendix Figures 3 and 4 for further details on per capita trends for the countries that are the largest markets and most affected by the spread of COVID-19 cases, respectively.¹⁶

3.3 What Country-level Characteristics Predict Differential Adoption due to COVID-19?

We subsequently test whether a number of other country-level characteristics drawn from the extant literature are significant predictors of differential trends in digital finance adoption due to COVID-19. Specifically, we merge country-level characteristics on country development level (log of GDP per capita in \$PPP), market size (log of country population), demographics (percentage of population over 65), access to ICT (percentage of the population with a mobile phone subscription), and ex-ante adoption of digital banking (as a proxy, we use the percentage of a country's population that reported regularly receiving wages using their mobile phones). We re-run Equation (1) including these country-level factors as interaction terms to test whether any pre-existing trends are driving our end outcomes, i.e., fintech adoption.

Tables 6 and 7 depict the results from these model specifications for the absolute number of downloads and relative downloads respectively. In line with existing literature, we find that

¹⁶We also took a preliminary look at trends for some of the underlying top finance mobile applications in our country sample. Payment and money transfers apps—both from global “BigTech” companies and local providers—appear to have seen the most consistent increases in absolute terms following COVID-19's spread. However, most other product categories are represented as well. For example, across all regions, we also see increased uptake of general banking apps for both traditional incumbents and some neo-banks. Moreover, we observe that some of the top apps in the sample countries from Africa, Asia, and Oceania are personal loan or consumer credit apps, which have seen particularly high growth in very recent weeks. There are also more select examples of wealth management, investment/trading, and cryptocurrency apps seeing large increases in adoption. In our next iteration, we plan to set up empirical specifications to formally test hypotheses related to provider types and product categories and to provide more concrete estimates of any differential trends.

country economic development level is not necessarily a strong predictor of differential adoption rates of mobile finance apps due to COVID-19. There are some signs that countries with higher incomes and industrialization level (as proxied by GDP per capita) had lower rates of adoption, post-pandemic, but these results are insignificant. Moreover, we do not find particularly strong evidence that ICT access or ex-ante digital finance adoption is strongly associated with differential rates of adoption due to COVID-19. This may suggest that, from a country-level perspective and for the time being, COVID-19 does not appear to have exacerbated differences in digital finance adoption between those with higher or lower ex-ante adoption.

What we do note is that larger market size (proxied by country population) is a strong predictor of fintech adoption during the pandemic. Moreover, it is intuitively associated with increased adoption in both absolute and relative terms with the economic magnitude of this relationship being quite large. This may be related to positive adoption externalities and spillover effects for many network-based financial technologies, as noted by (Crouzet et al., 2019; Higgins, 2019). In line with prior studies, we also find evidence that demographics plays a key role in the dissemination of fintech as aging populations are seen to be negatively associated with decreased adoption.

3.4 Does the Spread of COVID-19 or Government Policy Drive Adoption?

Our initial results are consistent in showing that a positive and significant relationship exists between the spread of COVID-19 and associated fintech app downloads. However, the shock due to COVID-19 is indeed unique with many different effects overlapping each other. Our results and empirical setting do not allow us to readily distinguish what primarily drives this uptick in fintech adoption. The effects we observe may be due to the direct spread of COVID-19 (first-order effects), the lockdowns instituted by various governments (second-order effects), or more protracted economic shocks from job losses, business closures, etc. arising from the lockdowns over time (third-order effects).

Table 8 depicts preliminary results for country-level specifications where we try to disentangle the first- and second-order effects, i.e., the spread of the pandemic and the government lockdowns. There is considerable overlap between the outcomes due to COVID-19 and government lockdowns and hence, any evidence presented, is suggestive at best. Our results depict that the adoption in fintech, in absolute and relative terms, seems to be driven mainly by the spread of the pandemic even when we control for the stringency of government policies to mitigate COVID-19 (columns 1 and 6). When we flip our dummy to capture post lockdown effects instead, and control for the growth in cases (columns 3 and 8), the results are again synonymous with depicting a strong positive correlation with the actual spread of COVID-19 itself. However, when we use the dummy and variables for confirmed deaths instead of confirmed cases, the lockdown effects dominate the

outcomes (column 7). The evidence, while somewhat mixed, seems more consistent in suggesting that a larger portion of the results are driven by the initial spread of COVID-19 and to a lesser extent the government lockdowns. The spread of the pandemic appears to have driven more individuals to adopt new digital finance and fintech platforms at an earlier stage of the pandemic propelled by the desire to follow *social distancing* norms and practice physical isolation.

4 Conclusion

We draw on mobile application data from 74 countries to document the effects of the COVID-19 pandemic on the adoption of digital finance and fintech, as proxied by finance category mobile application downloads on the extensive margin. We find that the spread of COVID-19 and related government lockdowns have led to statistically and economically significant increases in downloads of such mobile application, whether in absolute, relative, or per capita terms. These early effects have been fairly widespread across global regions, with the exception of Europe. Furthermore, at the country-level, we do not see strong initial signs that countries are exhibiting divergent trends based on economic development level or ex-ante rates of adoption.

There are several limitations to our work in its current iteration. First, we have not yet begun to address the considerable diversity of the fintech market, and analyze whether there are differential trends that may exist across different types of providers and products. For example, [BIS \(2018\)](#) identify several stylized scenarios of how competition between incumbent, fintech, and bigtech players may play out as adoption of financial technology increases, with divergent implications risks and opportunities for banks and banking systems. Moreover, it would be logical to better use the granular app-level data to assess whether the new adoption driven by COVID-19 is disproportionately stemming from certain product types—for example, payments and money-transfer, consumer or business lending, or other. Secondly, we currently focus solely on effects on the extensive margin (indication of changes in consumer demand and access) and do not directly touch on intensive margin effects for providers (e.g., transaction volumes, changes to financial performance, etc.) nor further welfare effects on consumers. On the one hand, we believe that focusing on short-term effects on the extensive margin makes sense since it serves as an important leading indicator. On the other hand, as time goes on, it will be increasingly important to also analyze these additional dimensions to see whether the documented large-scale increase in adoption has generally led to more opportunities or risks to consumers, firms, and governments. In a later iteration, we plan to complement the current analysis by drawing on available intensive margin data on mobile application usage and retention, as well as up-to-date data on provider financials to expand our analysis to cover these areas.

References

- Arner, D. W., Barberis, J. N., Walker, J., Buckley, R. P., Dahdal, A. M., & Zetsche, D. A. (2020). Digital finance & the COVID-19 crisis. *Available at SSRN 3558889*.
- Bharadwaj, P., Jack, W., & Suri, T. (2019). *Fintech and household resilience to shocks: Evidence from digital loans in Kenya* (Tech. Rep.). National Bureau of Economic Research.
- BIS. (2018). *Sound Practices: Implications of fintech developments for banks and bank supervisors* (Tech. Rep.). Bank for International Settlements.
- Carlin, B., Olafsson, A., & Pagel, M. (2017). *Fintech adoption across generations: Financial fitness in the information age* (Tech. Rep.). National Bureau of Economic Research.
- Chava, S., Ganduri, R., Paradkar, N., & Zeng, L. (2018). Shocked by bank funding shocks: Evidence from 500 million consumer credit cards. *Available at SSRN 3509356*.
- Claessens, S., Frost, J., Turner, G., & Zhu, F. (2018). Fintech credit markets around the world: Size, drivers and policy issues. *BIS Quarterly Review September*.
- Crouzet, N., Gupta, A., & Mezzanotti, F. (2019). *Shocks and technology adoption: Evidence from electronic payment systems* (Tech. Rep.). Northwestern University Working Paper.
- Das, S. R. (2019). The future of fintech. *Financial Management*, 48(4), 981–1007.
- Demirguc-Kunt, A., Klapper, L., Singer, D., Ansar, S., & Hess, J. (2018). *The global index database 2017: Measuring financial inclusion and the fintech revolution*. The World Bank.
- EY. (2019). *Global fintech adoption index 2019*. Ernst & Young.
- Foster, A. D., & Rosenzweig, M. R. (2010). Microeconomics of technology adoption. *Annual Review of Economics*, 2(1), 395–424.
- Frost, J. (2020). *The economic forces driving fintech adoption across countries* (Tech. Rep. No. 838). Bank for International Settlements.
- Fuster, A., Plosser, M., Schnabl, P., & Vickery, J. (2019). The role of technology in mortgage lending. *The Review of Financial Studies*, 32(5), 1854–1899.
- Gomber, P., Koch, J.-A., & Siering, M. (2017). Digital finance and fintech: Current research and future research directions. *Journal of Business Economics*, 87(5), 537–580.
- GPFI. (2016). *G20 high-level principles for digital financial inclusion*. Global Partnership for Financial Inclusion (GPFI).
- GSMA. (2020). The mobile economy 2020. *London: GSM Association*.

Higgins, S. (2019). Financial technology adoption. *JMP Berkeley*.

Jack, W., Ray, A., & Suri, T. (2013). Transaction networks: Evidence from mobile money in Kenya. *American Economic Review*, 103(3), 356–61.

Jack, W., & Suri, T. (2014). Risk sharing and transactions costs: Evidence from Kenya's mobile money revolution. *American Economic Review*, 104(1), 183–223.

Manyika, J., Lund, S., Singer, M., White, O., & Berry, C. (2016). Digital finance for all: Powering inclusive growth in emerging economies. *McKinsey Global Institute*.

Schindler, J. W. (2017). Fintech and financial innovation: Drivers and depth.

Table 1: “Global” finance mobile app market and COVID-19 statistics at a glance

This table provides aggregate-level descriptive statistics on the total and average number of finance category mobile app downloads, COVID-19 cases and deaths, and related government lockdowns or policies across our sample and study period. The app data covers the iOS and Android mobile app markets for 74 countries from January 1st, 2019 to present. (‡ denotes figures as of April 20th, 2020.) The sample countries collectively cover all global regions and account for approximately 80 percent of the global population and over 90 percent of the world economy in terms of nominal GDP. Data sources: AppTweak, OxCGRT, and Kaggle.

Variable	#
Total # of finance category mobile app downloads across sample countries (in 1,000’s)‡	5,999,545
Mean # of <i>daily</i> app downloads (in 1,000’s)	12,578
Total # of COVID-19 cases across sample countries (in 1,000’s)‡	2,147
Total # of COVID-19 deaths across sample countries (in 1,000’s)‡	153
Total # of sample countries with confirmed cases‡	72
Total # of sample countries with lockdowns‡	67
Total population across sample countries (in million’s)	5,970

Table 2: Country-level descriptive statistics

This table provides country-level descriptive statistics on the average number of daily finance mobile app downloads, average current number of COVID-19 cases and deaths, and related government lockdowns or policies. We also provide summary statistics on some additional country economic and demographic characteristics. The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to present. (‡ denotes figures as of April 20th, 2020.) The countries in our sample collectively cover all global regions and account for approximately 80 percent of the global population and over 90 percent of the world economy in terms of nominal GDP. Data sources: AppTweak, OxCGRT, Kaggle, WDI, Global Findex.

Variable	Obs.	Mean	Std. Dev.	Min	Max
<i>Panel A. Country-level statistics</i>					
# of daily downloads (in 1,000's)	34,272	168.1	418.1	0	4,464.1
# of COVID-19 cases (in 1,000's)‡	74	24.0	75.1	0.0	583.0
# of COVID-19 deaths (in 1,000's)‡	74	1.6	4.6	0.0	23.7
OxCGRT government stringency index (0 to 100)‡	74	64.5	41.3	0.0	100.0
# of days since 1st COVID-19 case‡	74	50.0	20.2	0.0	104.0
# of days since lockdown startdate‡	74	26.1	11.5	0.0	82.0
<i>Panel B. General sample characteristics</i>					
GDP per capita, \$PPP	70	32,482	22,333	2,688	112,600
Population (in millions)	74	83.0	229.0	0.6	1,390.0
% of population aged 65 and up	72	13.1	6.6	1.1	27.6
% of population with mobile phone	71	124.1	30.6	64.5	270.0
% of population using mobile phone to pay utilities	70	13.3	11.4	0.0	43.6

Table 3: The spread of COVID-19 cases and change in finance mobile app adoption

This table presents coefficient estimates for a panel regression model estimating the country-level relationship between the spread of COVID-19 on absolute number of daily downloads and relative (logarithmic) number of daily downloads for finance category mobile apps. We use a 14-day leading moving average to calculate downloads to smooth day-to-day fluctuations. *Post-Confirmed case* and *Post-Lockdown* denote dummy indicators that are turned on after a given country saw its first confirmed COVID-19 case or its government-initiated lockdown, respectively. The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to April 20th, 2020. Depending on the model specification, we include country fixed effects, month dummies to control for aggregate level seasonality in finance app downloads, and country X month interaction dummies to control for country-level seasonality. Standard errors are clustered at the country level. *p*-values in parentheses *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

	DV=Absolute # of daily app downloads (in 1,000's)				DV=Relative (ln) of # daily app downloads			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Post-Confirmed case</i>	142.138*	81.173*	81.566*	87.745*	0.712***	0.337***	0.286***	0.316***
	(0.013)	(0.040)	(0.044)	(0.033)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100
R^2	0.013	0.891	0.892	0.923	0.016	0.874	0.892	0.899
<i>Additional controls:</i>								
Country fixed effects	No	Yes	Yes	No	No	Yes	Yes	No
Month dummies (aggregate seasonality)	No	No	Yes	No	No	No	Yes	No
Country X month interactions (local seasonality)	No	No	No	Yes	No	No	No	Yes

Table 4: The spread of COVID-19 government lockdown and change in finance mobile app adoption

This table presents coefficient estimates for a panel regression model estimating the country-level relationship between the spread of COVID-19 on absolute number of daily downloads and relative (logarithmic) number of daily downloads for finance category mobile apps. We use a 14-day leading moving average to calculate downloads to smooth day-to-day fluctuations. *Post-Confirmed case* and *Post-Lockdown* denote dummy indicators that are turned on after a given country saw its first confirmed COVID-19 case or its government-initiated lockdown, respectively. The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to April 20th, 2020. Depending on the model specification, we include country fixed effects, month dummies to control for aggregate level seasonality in finance app downloads, and country X month interaction dummies to control for country-level seasonality. Standard errors are clustered at the country level. *p*-values in parentheses *** $p < 0.001$, ** $p < 0.01$, * $p < 0.5$. Data sources: AppTweak, OxCGRT, and Kaggle.

	DV=Absolute # of daily app downloads (in 1,000's)				DV=Relative (ln) of # daily app downloads			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Post-Lockdown</i>	103.903*	78.295*	72.197*	75.994*	0.488***	0.328***	0.228***	0.243***
	(0.033)	(0.032)	(0.037)	(0.032)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100
R^2	0.004	0.889	0.891	0.921	0.004	0.872	0.891	0.898
<i>Additional controls:</i>								
Country fixed effects	No	Yes	Yes	No	No	Yes	Yes	No
Month dummies (aggregate seasonality)	No	No	Yes	No	No	No	Yes	No
Country X month interactions (local seasonality)	No	No	No	Yes	No	No	No	Yes

Table 5: The spread of COVID-19 cases or lockdowns and change in finance mobile app adoption, by global region

This table presents coefficient estimates for a panel regression model estimating the region-level relationship between the spread of COVID-19 on absolute number of downloads and relative (logarithmic) number of downloads for finance category mobile apps. We use a 14-day leading moving average to calculate downloads to smooth day-to-day fluctuations. *Post-Confirmed case* and *Post-Lockdown* denote dummy indicators that are turned on after a given region saw its first confirmed COVID-19 case or its government-initiated lockdown, respectively. The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to April 20th, 2020. Europe serves as the base category. Depending on the model specification, we include region fixed effects, month dummies to control for aggregate level seasonality in finance app downloads, and region X month interaction dummies to control for regional seasonality. Standard errors are clustered at the region level. *p*-values in parentheses *** *p*<0.001, ** *p*<0.01, * *p*<0.5. Data sources: AppTweak, OxCGRT, and Kaggle.

	DV=Absolute # of daily app downloads (in 1,000's)				DV=Relative (ln) of # daily app downloads			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Post-Confirmed case</i>	55.20*	-22.44			0.23**	0.16*		
	(0.018)	(0.245)			(0.002)	(0.049)		
<i>Post-Confirmed case</i> X Africa	-55.13*	112.13***			0.18*	0.39***		
	(0.041)	(0.000)			(0.036)	(0.000)		
<i>Post-Confirmed case</i> X Asia	1,871.80***	2,604.00***			0.69***	1.10***		
	(0.000)	(0.000)			(0.000)	(0.000)		
<i>Post-Confirmed case</i> X Middle East	75.99**	162.87***			0.22*	0.31**		
	(0.003)	(0.000)			(0.019)	(0.001)		
<i>Post-Confirmed case</i> X North America	1,858.95***	1,778.11***			0.85***	0.85***		
	(0.000)	(0.000)			(0.000)	(0.000)		
<i>Post-Confirmed case</i> X Oceania	221.28***	292.28***			0.29**	0.35***		
	(0.000)	(0.000)			(0.002)	(0.001)		
<i>Post-Confirmed case</i> X South America	307.93***	593.70***			0.06	0.15		
	(0.000)	(0.000)			(0.547)	(0.084)		
<i>Post-Lockdown</i>			3.79	-36.21***			0.04	-0.02***
			(0.868)	(0.000)			(0.455)	(0.000)
<i>Post-Lockdown</i> X Africa			9.45	100.73***			0.25***	0.43***
			(0.620)	(0.000)			(0.000)	(0.000)
<i>Post-Lockdown</i> X Asia			1,896.57***	2,326.09***			0.44***	0.65***
			(0.000)	(0.000)			(0.000)	(0.000)
<i>Post-Lockdown</i> X Middle East			99.91***	195.56***			0.32***	0.44***
			(0.000)	(0.000)			(0.000)	(0.000)
<i>Post-Lockdown</i> X North America			2,369.55***	2,178.99***			0.97***	0.91***
			(0.000)	(0.000)			(0.000)	(0.000)
<i>Post-Lockdown</i> X Oceania			256.05***	295.97***			0.35***	0.36***
			(0.000)	(0.000)			(0.000)	(0.000)
<i>Post-Lockdown</i> X South America			434.24***	611.74***			0.22**	0.33***
			(0.000)	(0.000)			(0.008)	(0.000)
Observations	3,325	3,325	3,325	3,325	3,325	3,325	3,325	3,325
R ²	0.941	0.961	0.932	0.947	0.595	0.602	0.583	0.587
<i>Additional controls:</i>								
Region fixed effects	Yes	No	Yes	No	Yes	No	Yes	No
Month dummies (aggregate seasonality)	Yes	No	Yes	No	Yes	No	Yes	No
Region X month interactions (local seasonality)	No	Yes	No	Yes	No	Yes	No	Yes

Covid Economics 12, 1 May 2020: 158-191

Table 6: COVID-19 cases and change in fintech mobile app adoption, country-level characteristics predicting differential results

This table presents coefficient estimates for a panel regression model estimating the country-level relationship between the spread of COVID-19 on absolute number of daily downloads and relative (logarithmic) number of daily downloads for finance category mobile apps. We use a 14-day leading moving average to calculate downloads to smooth day-to-day fluctuations. *Post-Confirmed case* and *Post-Lockdown* denote dummy indicators that are turned on after a given country saw its first confirmed COVID-19 case or its government-initiated lockdown, respectively. The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to April 20th, 2020. All models include country X month interaction dummies to control for country-level seasonality. Standard errors are clustered at the country level. *p*-values in parentheses *** $p < 0.001$, ** $p < 0.01$, * $p < 0.5$. Data sources: AppTweak, OxCGRT, Kaggle, and WDI.

	DV=Absolute # of daily app downloads (in 1,000's)					DV=Relative (ln) of # daily app downloads				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>Post-Confirmed case</i>	930.370 (0.296)	-1,662.417* (0.039)	192.626 (0.079)	269.632 (0.161)	103.992 (0.073)	1.065 (0.074)	-1.332*** (0.000)	0.514*** (0.000)	0.247 (0.111)	0.285*** (0.000)
<i>Post-Confirmed case</i> X ln(GDP per capita)	-82.075 (0.336)					-0.072 (0.217)				
<i>Post-Confirmed case</i> X ln(Country population)		102.518* (0.036)					0.097*** (0.000)			
<i>Post-Confirmed case</i> X % Pop. aged 65 and up			-7.434 (0.210)					-0.014 (0.051)		
<i>Post-Confirmed case</i> X % Pop. w/ mobile subscription				-1.388 (0.269)					0.001 (0.593)	
<i>Post-Confirmed case</i> X % Pop. ex-ante receiving wages via mobile phone					-1.246 (0.693)					0.006 (0.312)
Observations	32,775	32,775	33,250	32,775	32,300	33,250	33,725	34,200	33,725	33,250
R^2	0.924	0.932	0.923	0.923	0.922	0.891	0.891	0.889	0.890	0.886
<i>Additional controls:</i>										
Country X month interactions (local seasonality)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 7: COVID-19 lockdowns and change in fintech mobile app adoption, country-level characteristics predicting differential results

This table presents coefficient estimates for a panel regression model estimating the country-level relationship between the spread of COVID-19 on absolute number of daily downloads and relative (logarithmic) number of daily downloads for finance category mobile apps. We use a 14-day leading moving average to calculate downloads to smooth day-to-day fluctuations. *Post-Confirmed case* and *Post-Lockdown* denote dummy indicators that are turned on after a given country saw its first confirmed COVID-19 case or its government-initiated lockdown, respectively. The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to April 20th, 2020. All models include country X month interaction dummies to control for country-level seasonality. Standard errors are clustered at the country level. *p*-values in parentheses *** $p < 0.001$, ** $p < 0.01$, * $p < 0.5$. Data sources: AppTweak, OxCGRT, Kaggle, and WDI.

	DV=Absolute # of daily app downloads (in 1,000's)					DV=Relative (ln) of # daily app downloads				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>Post-Lockdown</i>	716.814 (0.329)	-1,335.540 (0.066)	166.995 (0.063)	273.522 (0.230)	100.859* (0.047)	1.017 (0.098)	-0.978* (0.025)	0.493*** (0.000)	0.216 (0.344)	0.256*** (0.000)
<i>Post-Lockdown</i> X ln(GDP per capita)	-62.427 (0.378)					-0.074 (0.216)				
<i>Post-Lockdown</i> X ln(Country population)		83.432 (0.062)					0.073** (0.006)			
<i>Post-Lockdown</i> X % Pop. aged 65 and up			-6.404 (0.187)					-0.018** (0.005)		
<i>Post-Lockdown</i> X % Pop. w/ mobile subscription				-1.563 (0.346)					0.000 (0.874)	
<i>Post-Lockdown</i> X % Pop. ex-ante receiving wages via mobile phone					-2.404 (0.296)					-0.002 (0.705)
Observations	32,775	32,775	33,250	32,775	32,300	33,250	33,725	34,200	33,725	33,250
R^2	0.921	0.925	0.921	0.921	0.920	0.890	0.889	0.888	0.889	0.884
<i>Additional controls:</i>										
Country X month interactions (local seasonality)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 8: COVID-19's spread and change in finance mobile app adoption, public health versus government policy shock

This table presents coefficient estimates for a panel regression model estimating the country-level relationship between the spread of COVID-19 on absolute number of daily downloads and relative (logarithmic) number of daily downloads for finance category mobile apps. We use a 14-day leading moving average to calculate downloads to smooth day-to-day fluctuations. *Post-Confirmed case*, *Post-Confirmed death* and *Post-Lockdown* denote dummy indicators that are turned on after the first confirmed COVID-19 case, death, and government-initiated lockdown, respectively. Meanwhile, we estimate whether the intensity of COVID-19 cases, deaths or government policy also factors into rates of adoption through variables measuring the number of confirmed cases, deaths, and the OxCGRT index score for government policy stringency (transformed to logarithmic form). The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to April 20th, 2020. Standard errors are clustered at the country level. *p*-values in parentheses *** $p < 0.001$, ** $p < 0.01$, * $p < 0.5$. Data sources: AppTweak, OxCGRT, and Kaggle.

	DV=Absolute # of daily app downloads (in 1,000's)					DV=Relative (log) of # daily app downloads				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>Post-Confirmed case</i>	55.914*					0.173**				
	(0.057)					(0.014)				
<i>Post-Confirmed deaths</i>		35.008					-0.005			
		(0.126)					(0.916)			
ln(Government stringency index)	9.627	16.136*			14.593	0.043**	0.080***			0.075***
	(0.129)	(0.066)			(0.140)	(0.010)	(0.000)			(0.000)
<i>Post-Lockdown</i>			-58.315	-6.233				-0.094	0.159***	
			(0.228)	(0.840)				(0.139)	(0.002)	
ln(Number of confirmed cases)			19.648*		-0.902			0.049***		0.006
			(0.078)		(0.868)			(0.000)		(0.722)
ln(Number of confirmed deaths)				24.635	15.086				0.025*	-0.008
				(0.116)	(0.151)				(0.064)	(0.708)
Observations	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100
R ²	0.923	0.923	0.923	0.922	0.923	0.899	0.899	0.898	0.898	0.899
<i>Additional controls:</i>										
Country X month interactions (local seasonality)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

A Appendix

A.1 Data sample

Table A.1: Overview of countries in data sample

This table lists the 74 countries for which we have mobile app download data along with their region, date of first confirmed case of COVID-19, and date of lockdown (if applicable). Data sources: AppTweak, OxGCRT, and Kaggle.

Country	Country Code	Region	1st COVID-19 case	Lockdown startdate
Algeria	DZA	Africa	2 Mar 2020	24 Mar 2020
Argentina	ARG	South America	6 Mar 2020	20 Mar 2020
Australia	AUS	Oceania	26 Jan 2020	24 Mar 2020
Austria	AUT	Europe	26 Feb 2020	16 Mar 2020
Belarus	BLR	Europe	–	–
Belgium	BEL	Europe	2 Mar 2020	17 Mar 2020
Brazil	BRA	South America	1 Mar 2020	–
Bulgaria	BGR	Europe	8 Mar 2020	13 Mar 2020
Canada	CAN	North America	28 Jan 2020	12 Mar 2020
Chile	CHL	South America	5 Mar 2020	26 Mar 2020
China	CHN	Asia	1 Dec 2019	23 Jan 2020
Colombia	COL	South America	10 Mar 2020	25 Mar 2020
Croatia	HRV	Europe	27 Feb 2020	22 Mar 2020
Czech Republic	CZE	Europe	2 Mar 2020	16 Mar 2020
Denmark	DNK	Europe	29 Feb 2020	11 Mar 2020
Ecuador	ECU	South America	2 Mar 2020	24 Mar 2020
Egypt	EGY	Middle East	2 Mar 2020	24 Mar 2020
Estonia	EST	Europe	4 Mar 2020	13 Mar 2020
Finland	FIN	Europe	27 Feb 2020	27 Mar 2020
France	FRA	Europe	25 Jan 2020	16 Mar 2020
Germany	DEU	Europe	29 Jan 2020	20 Mar 2020
Greece	GRC	Europe	28 Feb 2020	18 Mar 2020
Hong Kong	HKG	Asia	23 Jan 2020	–
Hungary	HUN	Europe	5 Mar 2020	16 Mar 2020
Iceland	ISL	Europe	2 Mar 2020	16 Mar 2020
India	IND	Asia	2 Feb 2020	24 Mar 2020
Indonesia	IDN	Oceania	2 Mar 2020	15 Mar 2020
Ireland	IRL	Europe	4 Mar 2020	12 Mar 2020
Israel	ISR	Middle East	24 Feb 2020	12 Mar 2020
Italy	ITA	Europe	31 Jan 2020	11 Mar 2020
Japan	JPN	Asia	24 Jan 2020	27 Feb 2020
Jordan	JOR	Middle East	16 Mar 2020	21 Mar 2020
Kazakhstan	KAZ	Asia	15 Mar 2020	15 Mar 2020
Kuwait	KWT	Middle East	24 Feb 2020	16 Mar 2020
Lebanon	LBN	Middle East	27 Feb 2020	16 Mar 2020

Continued on next page


Continued from previous page

Country	Country Code	Region	1st COVID-19 case	Lockdown startdate
Lithuania	LTU	Europe	–	16 Mar 2020
Luxembourg	LUX	Europe	7 Mar 2020	16 Mar 2020
Macao	MAC	Asia	23 Jan 2020	–
Malaysia	MYS	Asia	25 Jan 2020	18 Mar 2020
Mexico	MEX	North America	29 Feb 2020	14 Mar 2020
Netherlands	NLD	Europe	29 Feb 2020	16 Mar 2020
New Zealand	NZL	Oceania	4 Mar 2020	25 Mar 2020
Nigeria	NGA	Africa	10 Mar 2020	25 Mar 2020
Norway	NOR	Europe	28 Feb 2020	12 Mar 2020
Oman	OMN	Middle East	25 Feb 2020	29 Mar 2020
Pakistan	PAK	Asia	27 Feb 2020	26 Mar 2020
Peru	PER	South America	9 Mar 2020	10 Mar 2020
Philippines	PHL	Asia	2 Feb 2020	16 Mar 2020
Poland	POL	Europe	7 Mar 2020	12 Mar 2020
Portugal	PRT	Europe	3 Mar 2020	18 Mar 2020
Qatar	QAT	Middle East	2 Mar 2020	11 Mar 2020
Romania	ROU	Europe	29 Feb 2020	25 Mar 2020
Russia	RUS	Europe	1 Feb 2020	27 Mar 2020
Saudi Arabia	SAU	Middle East	6 Mar 2020	15 Mar 2020
Serbia	SRB	Europe	11 Mar 2020	15 Mar 2020
Singapore	SGP	Asia	24 Jan 2020	26 Mar 2020
Slovak Republic	SVK	Europe	8 Mar 2020	–
Slovenia	SVN	Europe	6 Mar 2020	11 Mar 2020
South Africa	ZAF	Africa	8 Mar 2020	26 Mar 2020
South Korea	KOR	Asia	24 Jan 2020	23 Feb 2020
Spain	ESP	Europe	10 Feb 2020	14 Mar 2020
Sweden	SWE	Europe	27 Feb 2020	–
Switzerland	CHE	Europe	28 Feb 2020	18 Mar 2020
Thailand	THA	Asia	17 Jan 2020	22 Mar 2020
Tunisia	TUN	Middle East	10 Mar 2020	20 Mar 2020
Turkey	TUR	Middle East	13 Mar 2020	23 Mar 2020
Ukraine	UKR	Europe	13 Mar 2020	17 Mar 2020
United Arab Emirates	ARE	Middle East	30 Jan 2020	24 Mar 2020
United Kingdom	GBR	Europe	31 Jan 2020	18 Mar 2020
United States	USA	North America	25 Jan 2020	13 Mar 2020
Uruguay	URY	South America	15 Mar 2020	16 Mar 2020
Venezuela	VEN	South America	15 Mar 2020	16 Mar 2020
Vietnam	VNM	Asia	24 Jan 2020	19 Mar 2020
Zimbabwe	ZWE	Africa	22 Mar 2020	27 Mar 2020

A.2 Additional figures

Figure 3: Daily rate of downloads per capita for finance category mobile apps across 10 largest country markets

This figure depicts the daily rate of downloads per capita for finance category mobile applications from the top 10 country markets (Brazil, China, Germany, India, Indonesia, Japan, Mexico, Russia, South Korea, and the United States) and the “rest of the world” (ROW). Country codes are listed in Appendix Table A.1. The data includes mobile apps from both the iOS and Android platforms from October 1st, 2019 to April 20th, 2020. We calculate a 14-day leading moving average on downloads to smooth day-to-day fluctuations.



Covid Economics 12, 1 May 2020: 158-191

Figure 4: Daily rate of downloads per capita for finance category mobile apps across countries most affected by COVID-19

This figure depicts the daily rate of downloads per capita for finance category mobile applications for countries with the most confirmed cases of COVID-19 (Belgium, China, France, Germany, Italy, the Netherlands, Spain, Switzerland, United Kingdom, and the United States). Country codes are listed in Appendix Table A.1. The data includes mobile apps from both the iOS and Android platforms from October 1st, 2019 to April 20th, 2020. We calculate a 14-day leading moving average on downloads to smooth day-to-day fluctuations.


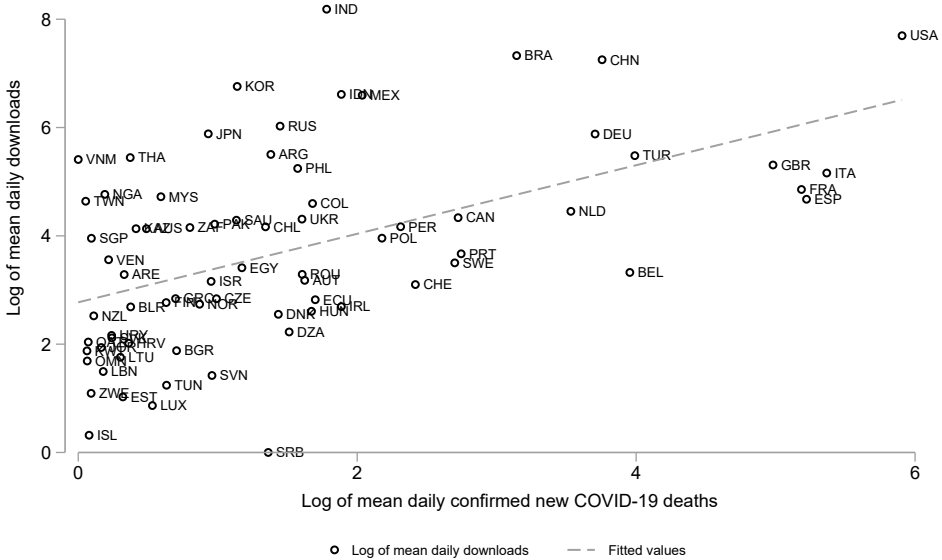


Figure 5: Scatterplot of Finance Category Mobile App Downloads vs. Confirmed COVID-19 Deaths


This figure depicts the relationship between average daily number of downloads for finance category mobile applications and average daily confirmed new deaths of COVID-19. We roughly restrict to the general period after the first signs of the COVID-19 outbreak (i.e., since January 1st, 2020). The underlying data includes mobile apps from both the iOS and Android platforms for 74 countries. Country codes are listed in Appendix Table A.1. Data sources: AppTweak and OxCGRT.



Covid Economics 12, 1 May 2020: 158-191

Figure 6: Scatterplot of Finance Category Mobile App Downloads vs. Government Policy Stringency Index

This figure depicts the relationship between average daily number of downloads for finance category mobile applications and average daily government stringency index score. Specifically, the Oxford COVID-19 Government Response Tracker (OxCGRT) systematically collects information on several different common policy responses governments have taken to mitigate COVID-19, scores the stringency of such measures, and aggregates these scores into a common Stringency Index. We roughly restrict to the general period after the first signs of the COVID-19 outbreak (i.e., since January 1st, 2020). The underlying data includes mobile apps from both the iOS and Android platforms for 74 countries. Country codes are listed in Appendix Table A.1. Data sources: AppTweak and OxCGRT.



A.3 Sensitivity tests on use of moving averages for estimating change in downloads

We test how sensitive our main results are to the use of different windows for calculating moving averages for downloads. We rerun main results for specifications 4 and 8 from Tables 3 and 4 using the original number of downloads versus a 10-day moving average, 14-day moving average, 20-day moving average, and 30-day moving average.

Tables A.2 and A.3 present results from this sensitivity analysis. In general, we do observe that the economic size and statistical significance of the effects tend to grow incrementally as we increase the window for the moving average. This is most notable when using the *Post-Confirmed case* variable for models estimating relative change in downloads. That having been said, the main trends and results are fairly consistent regardless of explanatory variable used or dependent variable construction. We imagine the original and 30-day moving average results provide an lower and upper-bound, respectively. Thus, in our main analysis, we use the 14-day moving average results to provide a middle ground estimate.

Table A.2: Sensitivity tests on construction of main DV, post-confirmed case

This table presents coefficient estimates for a panel regression model estimating the country-level relationship between the spread of COVID-19 on absolute number of daily downloads and relative (logarithmic) number of daily downloads for finance category mobile apps. We rerun specifications 4 and 8 from Table 3 using original download numbers and then when calculating the moving average using a 10-day (10DMA), 14-day (14DMA), 20-day (20DMA), and 30-day (30DMA) leading window. *Post-Confirmed case* and *Post-Lockdown* denote dummy indicators that are turned on after a given country saw its first confirmed COVID-19 case or its government-initiated lockdown, respectively. The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to April 20th, 2020. Standard errors are clustered at the country level. *p*-values in parentheses *** $p < 0.001$, ** $p < 0.01$, * $p < 0.5$. Data sources: AppTweak, OxCGRT, and Kaggle.

	DV=Absolute # of daily app downloads (in 1,000's)					DV=Relative (ln) of # daily app downloads				
	Original	10DMA	14DMA	20DMA	30DMA	Original	10DMA	14DMA	20DMA	30DMA
<i>Post-Confirmed case</i>	87.418*	87.625*	87.745*	88.106*	89.473*	0.251***	0.297***	0.316***	0.345***	0.395***
	(0.047)	(0.037)	(0.033)	(0.028)	(0.020)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100
R^2	0.937	0.927	0.923	0.918	0.914	0.993	0.921	0.899	0.876	0.861
<i>Additional controls:</i>										
Country X month interactions (local seasonality)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A.3: Sensitivity tests on construction of main DV, post-lockdown

This table presents coefficient estimates for a panel regression model estimating the country-level relationship between the spread of COVID-19 on absolute number of daily downloads and relative (logarithmic) number of daily downloads for finance category mobile apps. We rerun specifications 4 and 8 from Table 4 using original download numbers and then when calculating the moving average using a 10-day (10DMA), 14-day (14DMA), 20-day (20DMA), and 30-day (30DMA) leading window. *Post-Confirmed case* and *Post-Lockdown* denote dummy indicators that are turned on after a given country saw its first confirmed COVID-19 case or its government-initiated lockdown, respectively. The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to April 20th, 2020. Standard errors are clustered at the country level. *p*-values in parentheses *** $p < 0.001$, ** $p < 0.01$, * $p < 0.5$. Data sources: AppTweak, OxCGRT, and Kaggle.

	DV=Absolute # of daily app downloads (in 1,000's)					DV=Relative (ln) of # daily app downloads				
	Original	10DMA	14DMA	20DMA	30DMA	Original	10DMA	14DMA	20DMA	30DMA
<i>Post-Lockdown</i>	79.079*	76.937*	75.994*	75.258*	76.454*	0.249***	0.244***	0.243***	0.244***	0.249***
	(0.034)	(0.033)	(0.032)	(0.030)	(0.028)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100
R^2	0.936	0.925	0.921	0.916	0.912	0.993	0.920	0.898	0.874	0.859
<i>Additional controls:</i>										
Country X month interactions (local seasonality)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes



A.4 Results in terms of change in rate of daily downloads per capita

We alternatively calculate and present results in terms of the daily rate of downloads per capita. Across the 74 countries in our dataset, we observe a daily average of roughly 0.00195 downloads per capita in the period roughly prior to the advance of COVID-19 (i.e., in 2019). For the roughly 6 billion individuals covered in our country sample, this translates into an average of 11.7 million finance category app downloads daily from the Android and iOS platforms in the pre-COVID-19 period.

Table A.4 presents the comparable per capita-adjusted results for Tables 3 and 4. We observe that following the onset of confirmed COVID-19 cases or related lockdowns, the daily rate of downloads per capita increased by between 0.00066 and 0.00061—i.e., from roughly 0.00195 to between 0.00256 and 0.00261. Again taking into consideration the full country sample population, this equates to increasing from 11.7 to between 15.2 and 15.6 million downloads a day.

Similarly, Table A.5 presents the comparable per capita-adjusted results for the region-level analysis in Table 5. At a high-level, we observe similar patterns to the results from our specifications using logarithmic form to measure our DV. However, whereas results from Table 5 show Asia and North America as seeing the largest increases in terms of absolute terms, the per capita results emphasize that the magnitude of increase in North America and South America was particularly large, adjusting for countries' population size. European countries again show signs of a decrease in daily downloads per capita.

Table A.4: The spread of COVID-19 government lockdown and daily finance mobile app downloads per capita

This table presents coefficient estimates for a panel regression model estimating the country-level relationship between the spread of COVID-19 on the daily number of finance category mobile app downloads per capita. We use a 14-day leading moving average to calculate downloads to smooth some day-to-day fluctuations. *Post-Confirmed case* and *Post-Lockdown* denote dummy indicators that are turned on after a given country saw its first confirmed COVID-19 case or its government-initiated lockdown, respectively. The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to April 20th, 2020. Depending on the model specification, we include country fixed effects, month dummies to control for aggregate level seasonality in finance app downloads, and country X month interaction dummies to control for country-level seasonality. Standard errors are clustered at the country level. *p*-values in parentheses *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$. Data sources: AppTweak, OxCGRT, and Kaggle.

	DV=# of daily app downloads per capita							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Post-Confirmed case</i>	0.00100*** (0.000)	0.00065*** (0.000)	0.00061*** (0.000)	0.00066*** (0.000)				
<i>Post-Lockdown</i>					0.00077** (0.002)	0.00068*** (0.000)	0.00057*** (0.001)	0.00061*** (0.000)
Observations	36,100	36,100	36,100	36,100	36,100	36,100	36,100	36,100
R^2	0.020	0.895	0.900	0.924	0.007	0.892	0.897	0.921
<i>Additional controls:</i>								
Country fixed effects	No	Yes	Yes	No	No	Yes	Yes	No
Month dummies (aggregate seasonality)	No	No	Yes	No	No	No	Yes	No
Country X month interactions (local seasonality)	No	No	No	Yes	No	No	No	Yes

Table A.5: The spread of COVID-19 cases or lockdowns and daily finance mobile app downloads per capita, by global region

This table presents coefficient estimates for a panel regression model estimating the region-level relationship between the spread of COVID-19 on the daily number of download per capita for finance category mobile apps. We use a 14-day leading moving average to calculate downloads to smooth day-to-day fluctuations. *Post-Confirmed case* and *Post-Lockdown* denote dummy indicators that are turned on after a given region saw its first confirmed COVID-19 case or its government-initiated lockdown, respectively. The data sample covers the full iOS and Android mobile app markets for 74 countries daily from January 1st, 2019 to April 20th, 2020. Depending on the model specification, we include region fixed effects, month dummies to control for aggregate level seasonality in finance app downloads, and region X month interaction dummies to control for regional seasonality. Standard errors are clustered at the region level. Europe serves as the base region. *p*-values in parentheses *** *p*<0.001, ** *p*<0.01, * *p*<0.5. Data sources: AppTweak, OxCGRT, and Kaggle.

	DV=# of daily app downloads per capita			
	(1)	(2)	(3)	(4)
<i>Post-Confirmed case</i>	0.00006 (0.079)	-0.00003 (0.245)		
<i>Post-Confirmed case</i> X Africa	0.00005 (0.202)	0.00032*** (0.000)		
<i>Post-Confirmed case</i> X Asia	0.00060*** (0.000)	0.00077*** (0.000)		
<i>Post-Confirmed case</i> X Middle East	0.00039*** (0.000)	0.00055*** (0.000)		
<i>Post-Confirmed case</i> X North America	0.00380*** (0.000)	0.00361*** (0.000)		
<i>Post-Confirmed case</i> X Oceania	0.00082*** (0.000)	0.00094*** (0.000)		
<i>Post-Confirmed case</i> X South America	0.00090*** (0.000)	0.00146*** (0.000)		
<i>Post-Lockdown</i>			-0.00002 (0.573)	-0.00005*** (0.000)
<i>Post-Lockdown</i> X Africa			0.00014*** (0.000)	0.00026*** (0.000)
<i>Post-Lockdown</i> X Asia			0.00046*** (0.000)	0.00071*** (0.000)
<i>Post-Lockdown</i> X Middle East			0.00049*** (0.000)	0.00064*** (0.000)
<i>Post-Lockdown</i> X North America			0.00487*** (0.000)	0.00442*** (0.000)
<i>Post-Lockdown</i> X Oceania			0.00097*** (0.000)	0.00093*** (0.000)
<i>Post-Lockdown</i> X South America			0.00114*** (0.000)	0.00149*** (0.000)
Observations	3,325	3,325	3,325	3,325
R ²	0.870	0.893	0.835	0.864
<i>Additional controls:</i>				
Region fixed effects	Yes	No	Yes	No
Month dummies (aggregate seasonality)	Yes	No	Yes	No
Region X month interactions (local seasonality)	No	Yes	No	Yes

Sub-national allocation of COVID-19 tests: An efficiency criterion with an application to Italian regions¹

Christelle Baunez,² Mickael Degoulet,³ Stéphane Luchini,⁴
Patrick A. Pintus⁵ and Miriam Teschl⁶

Date submitted: 23 April 2020; Date accepted: 24 April 2020

Tests are crucial to know about the number of people who have fallen ill with COVID-19 and to understand in real-time whether the dynamics of the pandemic is accelerating or decelerating. But tests are a scarce resource in many countries. The key but still open question is thus how to allocate tests across sub-national levels. We provide a data-driven and operational criterion to allocate tests efficiently across regions or provinces, with the view to maximize detection of people who have been infected. We apply our criterion to Italian regions and compute the shares of tests that should go to each region, which are shown to differ significantly from the actual distribution.

- 1 This work was supported by French National Research Agency Grants ANR-17-EURE-0020. The authors gratefully thank the Editor and referee for suggestions that helped improve this paper.
- 2 Senior Research Fellow, Aix-Marseille University, CNRS, Institut Neuroscience Timone.
- 3 Research Fellow, Aix-Marseille University, CNRS, Institut Neuroscience Timone.
- 4 Research Fellow, Aix-Marseille University, CNRS, EHESS, Centrale Marseille, Aix-Marseille School of Economics..
- 5 Professor, Aix-Marseille University, CNRS, EHESS, Centrale Marseille, Aix-Marseille School of Economics.
- 6 Associate Professor, Aix-Marseille University, CNRS, EHESS, Centrale Marseille, Aix-Marseille School of Economics.

1 Introduction

Testing is crucial to detect people with SARS-CoV-2, the virus responsible for the COVID 19 pandemic (see e.g. [1], [2], [6], [3], [14]). In this vein, on March 16th 2020, the Director-General of the WHO strongly advised countries to “test, test, test” because “a fire cannot be fought blindfolded”.¹ But in many countries, tests are too scarce a resource to guarantee widespread and comprehensive testing (see e.g. [13]). Countries have set up guidelines spelling out the priority of who should get tested given that not all people can be tested. But this is only one side of the story and does not address the question of how to allocate a limited amount of tests across sub-national levels. We provide in this paper a data-driven and operational criterion that distributes tests according to how the virus spreads during the pandemic in different geographical units that either accelerate or decelerate.

The idea of the criterion is to base the allocation decision of a limited number of tests on data about the marginal benefit of an additional test. Such marginal benefit can be measured directly on the basis of a function representing the dynamics of the pandemic. Such a function is estimated using data on the numbers of both COVID-19 confirmed cases and total tests. Section 2 explains what information is necessary to calculate this function in real-time and how to derive the weights for the distribution of tests on a sub-national level. In Section 3 we apply this criterion to Italian regions, where certain regions struggle more than others with the pandemic. We compute the optimal shares of tests for each region and compare it with the actual distribution. Our main conclusion is that there is room for significant improvement in the geographical allocation of COVID-19 tests across Italian regions.


An important point of our approach is that it does not provide predictive model but rather relates to now-casting analysis. Forecasting models typically have many parameters, are highly uncertain at times of unknown and new events and can fall prey to estimation errors. Our data-driven criterion builds upon non-parametric estimation and as such constitutes an important real-time source of information for decision-makers, complements traditional epidemiological models and is amenable to public health analysis.

2 Characterization of the Testing Allocation Criterion

Our criterion builds upon a new approach that is designed to monitor and respond to a pandemic like COVID-19 (see [12]). The approach organises the data in real time in order to detect whether an ongoing pandemic is accelerating or decelerating. It is based on the idea that in times of thick uncertainty, when it is difficult to have precise estimates of probabilities and thus difficult to make reliable forecasting models, decisions should arguably put more weight on the information available in real time. Following [17] and [12], the question is how we can use available information

¹<https://www.who.int/dg/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19—16-march-2020-last> access April 14th 2020

Figure 1: Epidemic Dynamics in Two Hypothetical Regions with Different Dynamics



Note: Blue dots represent the Fraction of Total COVID-19 Cases (*y*-axis) against Fraction of COVID-19 Tests (*x*-axis) across Time, for Region *Q* in panel (a) and Region *S* in panel (b). Fraction of realised tests (resp. of cases) equals the cumulative number of tests realised (resp. cases) on day *t* divided by the cumulative number of tests (resp. cases) realised at the end date *T* - Red lines are the functional relationship between positive cases and total tests and depict the dynamic of the pandemic - Black Lines are the tangent lines of the functions in red at end point, the slopes of which measure the marginal benefits of testing at end point.

to detect acceleration or deceleration of harm in the current case of the pandemic. To answer this question the number of tests is crucial. Tests allow visualising the spread of the pandemic and are therefore a precondition for a reliable real-time analysis. If ever more cases of infected people are detected with an increasing number of total tests, the pandemic is accelerating. If, on the contrary, ever fewer infected people are detected with an increasing number of test, the pandemic is decelerating. This insight can be visualised by plotting the cumulated number of cases that have been tested positively against the cumulated number of total tests. The convexity of such functional relationship indicates acceleration, concavity deceleration. This is an important real-time guide for health policy decisions that aim to curb the pandemic: what health policy decision want to achieve and to keep is a concave functional relationship of positively tested subjects to total tests and eventually a flat relationship that indicates the end of the pandemic.²

We illustrate our approach by considering the following hypothetical example. A country has been subject to an epidemic like COVID-19 for several weeks. This country is composed of two regions labeled, say, N and S. While in region N, the virus spread rapidly from very early on, region S only saw a few cases at the beginning but experienced subsequently a rapid increase of infected people.

² Tests may also be an important health policy response to contain the pandemic, but this is not what we are discussing here. In this paper, testing allows to know about the pandemic spread whatever the health policy response put in place.

The pandemic has followed different dynamics across the two regions and the resulting heterogeneity is best visualized in Figure 1. Figure 1 depicts a scatter-plot of the number of COVID-19 tests against the number of COVID-19 cases, for hypothetical daily data from the two regions N and S, at a given date that we think of as today (the blue dots). Both the number of cases and the number of tests are normalized by dividing raw data over the last data point available today. The red lines in Figure 1 are the functional relationship between positive cases and total tests and thus represent the dynamics of the pandemic in each region. While in Figure 1 panel (a) this line is concave and indicates that tests find ever fewer infected people, in Figure 1 panel (b), the line is convex and points to the fact that tests find ever more cases of infected people. Said differently, what testing tells us is that while in region N the pandemic is decelerating, in region S it is accelerating.

With that interpretation in mind, deriving a criterion for allocating a given number of tests across the two regions is straightforward. The most important information is the marginal benefit of an additional test in each region. Knowing this, it is possible to compute the proportion of total available tests at the national level that each region should get. It should be obvious that a larger fraction of tests should go to the region where the marginal test is higher.

In Figure 1, the marginal benefit of a test is measured by the slope of the tangent of each curve at $x = 1$, which is the end point of each graph.³ Marginal benefits are depicted as the black lines, tangent to the red curve in Figure 1 panels (a) and (b). Given the steepness of the tangent in Figure 1 panel (b) due to the acceleration of the pandemic, it is obvious that more tests should be done in region S than in region N. Both curves in Figure 1 are the graphs of power function, so that in this example the slope at $x = 1$ is directly given by the exponent, which equals 0.2 for region N and 5 for region S. These numbers are interpreted as follows. In region N, the most recent 10% of the total number of test have helped uncover 2% (that is, 0.2×10) of total confirmed infections. This is much less than in region S, where the most recent 10% are associated with the detection of 50% (that is, 5×10) of the total number of COVID-19 cases. In that sense, more testing is needed in region S than in region N. Therefore, given that the marginal benefit of an additional test in region S is 25 (that is, $5/0.2$) times higher than in region N, the number of total tests available today should be divided as follows. Region S would get about 96% (that is, $5/5.2$) while region N would get the remainder 4% (that is, $0.2/5.2$).

It goes without saying that shares should be updated regularly, given the time-varying nature of the weights defined using our criterion, following changes in the marginal benefits/slopes of some regions relative to others. For example, the allocations derived from the early points in Figure 1 would imply that more tests should go to region N than to region S. Although the above weights do not take into account population in each region, it is straightforward to combine them with population weights, as we now show. We summarize below, for convenience, the steps to compute

³The marginal benefit of testing as defined here is reminiscent of distributional weights used in public policy analysis—see [7] for a discussion on distributional weights for health policies.

such weights and the corresponding share for each sub-national level, $i = 1, \dots, n$:

Criterion to allocate COVID-19 tests across sub-national levels

$i = 1, \dots, N$, with population n_i :

- for each sub-national level i , estimate a smooth curve from the scatter-plot of the number of confirmed COVID-19 cases against the number of COVID-19 tests across time. Estimate next at end point, that is, at the latest date for which data is available, the slope of such a curve for each i , which delivers a weight w_i
- each sub-national level i should receive a share of the total number of tests available at the national level that is defined by $A_i = \frac{w_i n_i}{\sum_i w_i n_i}$, taking into account population. Alternatively, such a share is defined as $a_i = \frac{w_i}{\sum_i w_i}$ if population is not accounted for.

3 An Application to Italian Regions


Italy is in a national lock-down period since March 9th 2020. About a month later, the number of newly detected cases are decreasing every day on a national level. However, across regions, the pandemic spread was and still is not uniform. As it is commonly known, regions such as Lombardia and Veneto have seen very high numbers of infected people as well as deaths, whereas more southern regions have had much fewer cases. Although testing has been deficient at the early stage of the pandemic (see [13]), Italy has successively implemented an aggressive testing strategy reaching a total of over a million tests by mid-April. These are many tests but still a low number for a population of about 60 million inhabitants. Given this limited number of possible tests for a large population, the crucial but still open question is how to distribute tests efficiently, in particular in view of the differences in pandemic dynamics across regions.

The data we use is made publicly available by the Italian Department of Civil Protection⁴ from which can be accessed the official dashboard that monitors the situation of the COVID-19 pandemic in Italy (“Mappa della situazione”). This map provides the link to data that contains information on all Italian regions.⁵ There are 20 regions, but the data has 21 entries because the region of Trentino-Alto Adige counts for two to take account of the two autonomous provinces (P.A.) that form that region, named *P.A. Bolzano* and *P.A. Trento* in the data set. From this

⁴<http://www.protezionecivile.gov.it/home>: last access April 14th 2020

⁵<https://github.com/pcm-dpc/COVID-19/tree/master/dati-regioni>: last access April 14th 2020

Figure 2: Optimal Allocations of tests in Italian regions as compared to actual allocation



Note: The figure presents the optimal allocations for each region as well as the mean actual allocation of tests in the last three days in the region ($x_i, i \in 1 \dots 21$). $n_i, k \in 1 \dots 21$, is the number of inhabitants in the region. $a_i, i \in 1 \dots 21$, is the optimal allocation that does not account for population size whereas $A_i, i \in 1 \dots 21$, is the optimal allocation that ponderates by population size n_i . Pink (resp. orange and blue) bar represents the optimal allocation (resp. the optimal allocation weighted by population size and the mean actual allocation).

dataset, we use the information on the total number of tests performed (“tamponi”) and the total number of positive cases (“totale casi”) for each region, which is updated on a daily basis.

We estimate both the functional relationship between the total number of tests and the total

number of positive case for each region and its first derivative non-parametrically by local polynomial regression fitting with Epanechnikov weights [15]. The method allows to fit a smooth curve relating tests and cases without imposing restrictive parametric restrictions on the functional form [9]. We report in Figure 2 the estimated (population-unweighted) optimal shares a_i , $i = 1 \dots 21$, and the (population-weighted) optimal shares A_i , $i = 1 \dots 21$, using the mean estimate of the first derivative on the last 2.5% of output design points (see Appendix A and B for the detailed results). We also report as a benchmark the mean actual allocation of tests in the last three days of the observed sample, for comparison purpose.

In Figure 2, each line represents an Italian region, whose population n_i appears in parentheses under the name of the region in the first column. In the second column we report the population-weighted and population-unweighted shares, respectively, A_i and a_i , as well as the actual share denoted x_i . The third column helps visualising those shares. As can be seen in Figure 2, differences between the optimal shares as computed by using our (population-weighted and unweighted) criterion and actual shares of tests in all regions are significant. A striking conclusion is that certain regions currently receive too many tests, compared to others who have received too little. In particular, a substantial share of tests conducted in regions such as Lombardia, Veneto and Emilia-Romagna, where the pandemic spread more rapidly at the beginning, should be allocated to other regions. This is particularly true if we consider the population-unweighted share a_i of tests but also if we take the population-weighted share A_i into account. However, what also becomes apparent is that Lombardia and Veneto have currently about the same share of total tests allocated (about 20% each) despite the fact that Veneto has only about half of the population of Lombardia. Our population-weighted criterion A_i takes this into account and allocates more tests to Lombardia than to Veneto given its population size. If we continue to focus on the population-weighted share A_i , another striking example are the regions of Lazio, but also Sicily, which should receive a much larger share of tests, given the number of inhabitants in those regions, than they currently do. Indeed, Lazio, the second most populated region, should receive twice as many tests as it currently does, that is about 13% of the total number of Italian tests. This is because both its marginal benefit of testing and its regional population are among the largest across Italy. Sicily is in a similar situation and should also get twice as many tests. Generally, in view of Figure 2, it is hard to conclude that tests have been allocated across Italian regions where they are the most efficient and mostly needed. Obviously, to achieve the goal of efficient distribution of tests, it is important that regional governments and administrations are team-players and a nation-wide coordination is possible.

4 Conclusion

In this paper, we provide a data-driven and operational criterion that allows to distribute efficiently a limited amount of tests across sub-national units, depending on the dynamic of the pandemic.

This allocation enhances the likelihood of detecting the largest possible number of infected persons. Knowing about the number of infected people is crucial to understand in real-time whether the dynamics of pandemic is accelerating or decelerating. If with an increasing number of tests ever more positive cases are found, the harm of the pandemic is accelerating. If with an increasing number of tests the number of detected positive cases declines on a daily basis, the pandemic starts to be contained and harm decelerates. We represented this in terms of a scatter-plot of total number of positive cases over total number of tests and pointed to the curvature of the underlying functional relationship. The pandemic ends when further tests do not find any new positive cases and the functional relationship becomes a flat line. The test-distribution criterion uses the information of the tangent to this functional form at endpoint, which we claim is a useful measure of the marginal benefit of testing. The policy implication of our criterion is that more tests will be needed in regions where the tangent is steeper and indicates an acceleration of the pandemic in comparison to regions where the tangent is flatter and thus means that the pandemic is spreading less quickly. Our criterion could be combined with group or pool testing (see [10] and [8]) and drive-in testing ([11] and [19]) so as to use the limited amount of test capacity in a best way in each region or province. Pool testing is currently developed and implemented in Germany ([16]) and Israel⁶. We show how the criterion can be applied to Italy and its regions where the pandemic had a very different dynamic. We find that the current actual allocation of tests across regions is not efficient with respect to the criterion we provide in this paper. In particular, there should be a shift in testing from certain regions where the pandemic at first spread very quickly but has subsequently been more and more controlled, to some of those regions where the pandemic started out slower, but seems to speed up more recently.

Although our application makes use of Italian data, a similar exercise could obviously be done for other countries where more fine-tuned regional data is available. Using such a criterion will be particularly important at a moment where certain countries start considering to remove at least some aspects of the lock-down. Removing the lock-down will certainly trigger a change in the dynamic of the pandemic and it is crucial to have real-time evaluations of the re-opening of the social and economic life as well as a tool to allocate tests efficiently within the country. It goes without saying that this criterion could also be used across nations and in particular for political unions such as the European Union to efficiently allocate tests across member countries. It has become clear quite soon that certain countries were better prepared for testing from very early on than others. In this context, it turns out that the often cited aspect of “European solidarity” may have most usefully been an economic principle of efficient testing.

⁶<https://www.timesofisrael.com/to-ease-global-virus-test-bottleneck-israeli-scientists-suggest-pooling-samples/>


References







- [1] Bedford J., Enria D., Giesecke J., et al. COVID-19 (2020): Towards Controlling of a Pandemic. *The Lancet* S0140-6736(20)30673-5. Available at [https://doi.org/10.1016/S0140-6736\(20\)30673-5](https://doi.org/10.1016/S0140-6736(20)30673-5)
- [2] Binnicker, M.J. (2020): Emergence of a Novel Coronavirus Disease (COVID-19) and the Importance of Diagnostic Testing: Why Partnership between Clinical Laboratories, Public Health Agencies, and Industry Is Essential to Control the Outbreak. *Clinical Chemistry* hvaa071 Available at <https://doi.org/10.1093/clinchem/hvaa071>
- [3] COVID-19 National Emergency Response Center, Epidemiology and Case Management Team, Korea Centers for Disease Control and Prevention (2020): Contact Transmission of COVID-19 in South Korea: Novel Investigation Techniques for Tracing Contacts. *Osong Public Health Res Perspect* 2020, 11:60-63. Available at <https://doi.org/10.24171/j.phrp.2020.11.1.09>
- [4] De Brabanter, K., and Liu, Y. (2015): Smoothed Nonparametric Derivative Estimation Based on Weighted Difference Sequences, In *Stochastic Models, Statistics and Their Applications*, A. Steland, E. Rafajlowicz, K. Szajowski (Eds.), Chapter 4 (pp. 31-38), Springer.
- [5] ECDC technical report (2020): Guidelines for the use of non-pharmaceutical measures to delay and mitigate the impact of 2019-nCoV. European Center for Disease Control. https://www.ecdc.europa.eu/sites/default/files/documents/novel-coronavirus-guidelines-non-pharmaceutical-measures_0.pdf
- [6] Fineberg, H.V. (2020): Ten Weeks to Crush the Curve. *New England Journal of Medicine* 10.1056/NEJMe2007263. Available at <https://www.nejm.org/doi/full/10.1056/NEJMe2007263>
- [7] Fleurbaey, M, Luchini, S., Muller, C. and E. Schokkaert (2013): Equivalent Income and fair evaluation of health care. *Health economics*, 22(6) 711-729. Available at <https://onlinelibrary.wiley.com/doi/abs/10.1002/hec.2859>
- [8] Gollier, C., and Gossner, O. (2020): Group Testing against Covid-19. *Covid Economics*, Issue 2. Available at <https://cepr.org/sites/default/files/news/CovidEcon1%20final.pdf>
- [9] Harrel, F.Z. Jr (2015). *Regression Modeling Strategies: With Applications to Linear Models, Logistic and Ordinal Regression, and Survival Analysis*. Springer.
- [10] Jain, B.N., and Jain, T. (2020). Accelerated Infection Testing at Scale: a Proposal for Inference with Single Test on Multiple Patients. Available at <https://arxiv.org/ftp/arxiv/papers/2003/2003.13282.pdf>







- [11] Kwon, K.T., Ko J.H., Shin H., Sung M., Kim J.Y. (2020): Drive-Through Screening Center for COVID-19: a Safe and Efficient Screening System against Massive Community Outbreak. *Journal of Korean Medical Science* 35:e123. Available at <https://doi.org/10.3346/jkms.2020.35.e123>
- [12] Luchini, S., Teschl, M., Pintus, P.A., Baunez, C., Degoulet, M., and Moatti, J.P. (2020): Urgently Needed for Policy Guidance: An Operational Tool for Monitoring the COVID-19 Pandemic. AMSE working paper 2020-09, available at <https://ssrn.com/abstract=3563688> and <https://www.amse-aixmarseille.fr/en/research/working-papers>
- [13] Rosenbaum, L. (2020): Facing Covid-19 in Italy - Ethics, Logistics, and Therapeutics on the Epidemic's Front Line. *New England Journal of Medicine* 10.1056/NEJMp2005492. Available at <https://doi.org/10.1056/NEJMp2005492>
- [14] Salathé M., Christian A., Richard N., et al. COVID-19 (2020): Epidemic in Switzerland: on the Importance of Testing, Contact Tracing and Isolation. Available at <https://doi.org/10.4414/smw.2020.20225>
- [15] Seifert, B., Brockmann, M., Engel, J. and T. Gasser (1994) Fast algorithms for nonparametric curve estimation. *J. Computational and Graphical Statistics* 3, 192-213. Available at <https://www.tandfonline.com/doi/abs/10.1080/10618600.1994.10474638>
- [16] Seifried, E., Ciesek, S., et al. (2020): Pool testing of SARS-CoV-2 samples increases worldwide test capacities many times over. Available at <https://aktuelles.uni-frankfurt.de/englisch/pool-testing-of-sars-cov-02-samples-increases-worldwide-test-capacities-many-times-over/>
- [17] Taleb, N. (2012). *Antifragile, things that gain from disorder*. Random House.
- [18] Watkins, J. (2020): Preventing a covid-19 pandemic. *BMJ* 368:m810. Available at <https://doi.org/10.1136/bmj.m810>
- [19] Won, J., Lee S., Park M., et al. (2020): Development of a Laboratory-safe and Low-cost Detection Protocol for SARS-CoV-2 of the Coronavirus Disease 2019 (COVID-19). *Experimental Neurobiology* 10.5607/en20009. Available at <https://doi.org/10.5607/en20009>




A Epidemic Dynamics in Italian Regions

In the following figures, we present for each region a scatter-plot of the raw data as well as of the non parametric estimation of the functional relationship between the fraction of COVID-19 tests and the fraction of COVID-19 cases. The fraction of COVID-19 tests is the daily observation of the cumulative sum of tests divided by the total number of tests observed at the end date. The fraction of COVID-19 cases is the daily observation of the cumulative sum of positive cases divided by the total number of cases observed at the end date. Blue dots correspond to raw data. Each blue dot corresponds to a daily observation. The lowest dot (down, left) on each Figure is the first date and the highest (up, right) dot corresponds to the end date. The red line is the non parametric estimation of the functional relationship between tests and cases using local polynomial regression fitting with Epanechnikov weights. The plain line in black is the tangent line derived from the end point from a non-parametric estimate of the first derivative of the functional relationship. The dashed line is the tangent line obtained from the mean first derivative estimate of the last 2.5% output design points and the dotted line is the tangent line derived from the mean first derivative estimate of the 10% output design points – see Appendix B for numerical results.









B Numerical Estimates

In the following tables, we present the estimated efficient allocations using different non-parametric estimates of the first derivative at the end date: (1) the end point estimate, (2) the mean estimate of the last 2.5% output design points and (3) the mean estimate of the last 10% output design points. The sensitivity analysis indicates that efficient allocation shares are robust to the choice of the first estimate derivative. We consider that the mean estimate of the last 2.5% output design points offers the most versatility for our purpose in our setting.

B.1 Efficient Allocations with End Point Estimate

Region	Inhabitants (n_i)	Weight (w_i)	Allocation (a_i)	Allocation (A_i)	Actual All. (x_i)
Abruzzo	1.311.580	1.354	9.27	4.19	1.970
Basilicata	562.869	1.147	7.85	1.52	0.411
Calabria	1.947.131	0.466	3.19	2.14	1.853
Campania	5.801.692	0.229	1.57	3.14	3.444
Emilia-Romagna	4.459.477	0.497	3.40	5.23	9.746
Friuli-Venezia Giulia	1.215.220	1.077	7.37	3.09	3.086
Lazio	5.879.082	0.949	6.50	13.17	6.702
Liguria	1.550.640	0.928	6.35	3.40	2.176
Lombardia	10.060.574	0.732	5.01	17.38	20.862
Marche	1.525.271	0.871	5.96	3.13	2.168
Molise	305.617	-0.010	0.07	-0.01	0.233
P.A. Bolzano	520.891	0.405	2.78	0.50	2.339
P.A. Trento	538.223	0.908	6.21	1.15	1.638
Piedmont	4.356.406	0.643	4.40	6.61	6.012
Apulia	4.029.053	0.679	4.65	6.46	3.052
Sardinia	1.639.591	0.651	4.46	2.52	1.038
Sicilia	4.999.891	0.730	5.00	8.61	3.353
Toscana	3.729.641	0.547	3.74	4.81	7.575
Umbria	882.015	0.418	2.86	0.87	1.761
Valle d'Aosta	125.666	0.351	2.40	0.10	0.366
Veneto	4.905.854	1.035	7.09	11.99	20.214

B.2 Efficient Allocations with Last 2.5% Output Design Points

Region	Inhabitants (n_i)	Weight (w_i)	Allocation (a_i)	Allocation (A_i)	Actual All. (x_i)
Abruzzo	1.311.580	1.248	8.85	3.99	1.970
Basilicata	562.869	1.071	7.59	1.47	0.411
Calabria	1.947.131	0.473	3.36	2.25	1.853
Campania	5.801.692	0.267	1.89	3.78	3.444
Emilia-Romagna	4.459.477	0.530	3.76	5.76	9.746
Friuli-Venezia Giulia	1.215.220	1.019	7.22	3.02	3.086
Lazio	5.879.082	0.898	6.36	12.86	6.702
Liguria	1.550.640	0.890	6.31	3.36	2.176
Lombardia	10.060.574	0.696	4.93	17.06	20.862
Marche	1.525.271	0.831	5.89	3.09	2.168
Molise	305.617	0.020	0.15	0.02	0.233
P.A. Bolzano	520.891	0.417	2.96	0.53	2.339
P.A. Trento	538.223	0.862	6.11	1.13	1.638
Piedmont	4.356.406	0.634	4.50	6.73	6.012
Apulia	4.029.053	0.663	4.70	6.51	3.052
Sardinia	1.639.591	0.641	4.55	2.56	1.038
Sicilia	4.999.891	0.681	4.83	8.30	3.353
Toscana	3.729.641	0.534	3.79	4.85	7.575
Umbria	882.015	0.396	2.80	0.85	1.761
Valle d'Aosta	125.666	0.347	2.46	0.11	0.366
Veneto	4.905.854	0.986	6.99	11.78	20.214

B.3 Efficient Allocations with Last 10% Output Design Points

Region	Inhabitants (n_i)	Weight (w_i)	Allocation (a_i)	Allocation (A_i)	Actual All. (x_i)
Abruzzo	1.311.580	0.988	7.58	3.39	1.970
Basilicata	562.869	0.885	6.79	1.30	0.411
Calabria	1.947.131	0.500	3.84	2.55	1.853
Campania	5.801.692	0.370	2.84	5.61	3.444
Emilia-Romagna	4.459.477	0.618	4.74	7.20	9.746
Friuli-Venezia Giulia	1.215.220	0.880	6.75	2.80	3.086
Lazio	5.879.082	0.780	5.99	11.99	6.702
Liguria	1.550.640	0.801	6.14	3.25	2.176
Lombardia	10.060.574	0.619	4.75	16.29	20.862
Marche	1.525.271	0.738	5.66	2.94	2.168
Molise	305.617	0.120	0.92	0.10	0.233
P.A. Bolzano	520.891	0.457	3.51	0.62	2.339
P.A. Trento	538.223	0.752	5.77	1.06	1.638
Piedmont	4.356.406	0.621	4.77	7.08	6.012
Apulia	4.029.053	0.632	4.85	6.66	3.052
Sardinia	1.639.591	0.625	4.79	2.68	1.038
Sicilia	4.999.891	0.571	4.38	7.46	3.353
Toscana	3.729.641	0.509	3.90	4.96	7.575
Umbria	882.015	0.347	2.66	0.80	1.761
Valle d'Aosta	125.666	0.352	2.70	0.12	0.366
Veneto	4.905.854	0.869	6.67	11.15	20.214