


**CENTRE FOR
ECONOMIC
POLICY
RESEARCH**

CEPR PRESS



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

**ISSUE 40
30 JULY 2020**

**SOCIO-DEMOGRAPHIC FACTORS AND
SELF-PROTECTING**

Nicholas W. Papageorge, Matthew V. Zahn,
Michèle Belot, Eline van den Broek-Altenburg,
Syngjoo Choi, Julian C. Jamison and
Egon Tripodi

ASSESSING ALTERNATIVE POLICIES

Michael Gapen, Jonathan Millar, Blerina Uruçi
and Pooja Sriram

**WHO SUFFERS FROM COVID IN
JAPAN?**

Shinnosuke Kikuchi, Sagiri Kitao
and Minamo Mikoshiba

OPTIMAL SOCIAL DISTANCING

Yoseph Y. Getachew

THE SHAPE OF RECOVERY

Dhruv Sharma, Jean-Philippe Bouchaud,
Stanislao Gualdi, Marco Tarzia
and Francesco Zamponi

**CULTURE AND CONTAGION IN
GERMANY**

Ioannis Laliotisi and Dimitrios Minos

GROWTH FORECASTS

Javier G. Gómez-Pineda

**SCHOOL DISRUPTION DURING FOOT
AND MOUTH DISEASE**

Will Cook

Covid Economics

Vetted and Real-Time Papers

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

Founder: Beatrice Weder di Mauro, President of CEPR

Editor: Charles Wyplosz, Graduate Institute Geneva and CEPR

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

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

The Centre for Economic Policy Research (CEPR)

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

Chair of the Board

Sir Charlie Bean

Founder and Honorary President

Richard Portes

President

Beatrice Weder di Mauro

Vice Presidents

Maristella Botticini

Ugo Panizza

Philippe Martin

Hélène Rey

Chief Executive Officer

Tessa Ogden

Editorial Board

Beatrice Weder di Mauro, CEPR

Charles Wyplosz, Graduate Institute Geneva and CEPR

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

Abi Adams-Prassl, University of Oxford and CEPR

Guido Alfani, Bocconi University and CEPR

Franklin Allen, Imperial College Business School and CEPR

Michele Belot, European University Institute and CEPR

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

Nick Bloom, Stanford University and CEPR

Tito Boeri, Bocconi University and CEPR

Alison Booth, University of Essex and CEPR

Markus K Brunnermeier, Princeton University and CEPR

Michael C Burda, Humboldt Universitaet zu Berlin and CEPR

Aline Bütikofer, Norwegian School of Economics

Luis Cabral, New York University and CEPR

Paola Conconi, ECARES, Universite Libre de Bruxelles and CEPR

Giancarlo Corsetti, University of Cambridge and CEPR

Fiorella De Fiore, Bank for International Settlements and CEPR

Mathias Dewatripont, ECARES, Universite Libre de Bruxelles and CEPR

Jonathan Dingel, University of Chicago Booth School and CEPR

Barry Eichengreen, University of California, Berkeley and CEPR

Simon J Evenett, University of St Gallen and CEPR

Maryam Farboodi, MIT and CEPR

Antonio Fatás, INSEAD Singapore and CEPR

Francesco Giavazzi, Bocconi University and CEPR

Christian Gollier, Toulouse School of Economics and CEPR

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

Miklós Koren, Central European University and CEPR

Anton Korinek, University of Virginia and CEPR

Philippe Martin, Sciences Po and CEPR

Warwick McKibbin, ANU College of Asia and the Pacific

Kevin Hjortshøj O'Rourke, NYU Abu Dhabi and CEPR

Evi Pappa, European University Institute and CEPR

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

Richard Portes, London Business School and CEPR

Carol Propper, Imperial College London and CEPR

Lucrezia Reichlin, London Business School and CEPR

Ricardo Reis, London School of Economics and CEPR

Hélène Rey, London Business School and CEPR

Dominic Rohner, University of Lausanne and CEPR

Paola Sapienza, Northwestern University and CEPR

Moritz Schularick, University of Bonn and CEPR

Paul Seabright, Toulouse School of Economics and CEPR

Flavio Toxvaerd, University of Cambridge

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

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

Jan C. van Ours, Erasmus University Rotterdam and CEPR

Thierry Verdier, Paris School of Economics and CEPR

Ethics

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

Submission to professional journals

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

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

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

Covid Economics

Vetted and Real-Time Papers

Issue 40, 30 July 2020

Contents

Socio-demographic factors associated with self-protecting behavior during the Covid-19 pandemic	1
<i>Nicholas W. Papageorge, Matthew V. Zahn, Michèle Belot, Eline van den Broek-Altenburg, Syngjoo Choi, Julian C. Jamison and Egon Tripodi</i>	
Assessing the effectiveness of alternative measures to slow the spread of COVID-19 in the United States	46
<i>Michael Gapen, Jonathan Millar, Blerina Uruçi and Pooja Sriram</i>	
Who suffers from the COVID-19 shocks? Labor market heterogeneity and welfare consequences in Japan	76
<i>Shinnosuke Kikuchi, Sagiri Kitao and Minamo Mikoshiba</i>	
Optimal social distancing in SIR-based macroeconomic models	115
<i>Yoseph Y. Getachew</i>	
V-, U-, L- or W-shaped economic recovery after COVID-19? Insights from an agent-based model	164
<i>Dhruv Sharma, Jean-Philippe Bouchaud, Stanislao Gualdi, Marco Tarzia and Francesco Zamponi</i>	
Spreading the disease: The role of culture	180
<i>Ioannis Laliotisi and Dimitrios Minos</i>	
Growth forecasts and the Covid-19 recession they convey	211
<i>Javier G. Gómez-Pineda</i>	
School disruption and pupil academic outcomes – evidence from the 2001 foot and mouth disease epidemic in England	229
<i>Will Cook</i>	

Socio-demographic factors associated with self-protecting behavior during the Covid-19 pandemic¹

Nicholas W. Papageorge,² Matthew V. Zahn,³ Michèle Belot,⁴
Eline van den Broek-Altenburg,⁵ Syngjoo Choi,⁶
Julian C. Jamison⁷ and Egon Tripodi⁸

Date submitted: 24 July 2020; Date accepted: 24 July 2020

Disease spread is in part a function of individual behavior. We examine the factors predicting individual behavior during the Covid-19 pandemic in the United States using novel data collected by Belot et al. (2020). Among other factors, we show that people with lower income, less flexible work arrangements (e.g., an inability to tele-work) and lack of outside space at home are less likely to engage in behaviors, such as social distancing, that limit the spread of disease. We also find that individuals in Florida and Texas (versus California and New York), men (versus women) and people who perceive fewer benefits of protective behaviors (versus those who perceive more benefits) report lower levels of self-protecting behaviors. Broadly, our findings align with many typical relationships between health and socio-economic status. Moreover, we show that the burden of measures designed to stem the pandemic are unevenly distributed across socio-demographic groups in ways that affect behavior and thus potentially the spread of illness. Policies that assume otherwise are unlikely to be effective or sustainable.

- 1 We are grateful for helpful comments from Stefanie DeLuca, Barton Hamilton and Emma Kalish. Regarding author ordering, Papageorge and Zahn led this particular effort and are thus listed first. The remaining authors contributed equally and are listed in alphabetical order. Research funding from the Creative-Pioneering Researchers Program at Seoul National University, and from the European University Institute are gratefully acknowledged.
- 2 Broadus Mitchell Associate Professor, Johns Hopkins University, IZA, and NBER.
- 3 PhD Candidate, Department of Economics, Johns Hopkins University.
- 4 Professor of Economics, European University Institute and IZA.
- 5 Assistant Professor, University of Vermont.
- 6 Professor of Economics, Seoul National University.
- 7 Professor of Economics, University of Exeter.
- 8 PhD Candidate, European University Institute.

Copyright: Nicholas W. Papageorge, Matthew V. Zahn, Michèle Belot,
Eline van den Broek-Altenburg, Syngjoo Choi, Julian C. Jamison and Egon Tripodi

1 Introduction

The spread of illness is not solely a biological phenomenon. It is also a social one, driven in part by human behavior. In the presence of strong externalities, a concern is that individual behavior may not align with socially optimal outcomes (Posner and Philipson (1993)). This is especially salient in contexts where the costs of protective behaviors, i.e., behaviors that limit the spread of illness, are unevenly distributed across socio-demographic groups (Pampel, Krueger, and Denney (2010)). For instance, in the Covid-19 pandemic, individuals who face a relatively low risk of serious illness, but who are economically vulnerable (e.g., lacking comfortable housing, the ability to work from home, and so on) may not follow recommendations or directives to engage in protective behaviors, such as wearing a mask or social-distancing. This potentially puts high-risk groups in danger of infection and prolongs the pandemic.

The socially optimal amount of protective behaviors—the levels that balance public health concerns with individual burdens and aggregate economic costs—are not yet fully understood, and are unlikely to be for some time due to uncertainty about the virus and about which behaviors most effectively prevent its spread (Manski (1999)). Compounding this uncertainty, we do not yet understand the full extent of the economic and social costs of the pandemic itself and measures taken to avoid it, ranging from job losses, shuttered businesses, gaps in schooling, violence, and addiction, among others (Fairlie, Couch, and Xu (2020), Alon et al. (2020), Mongey, Pilossoph, and Weinberg (2020), Coibion, Gorodnichenko, and Weber (2020), Viner et al. (2020)). Never-the-less, understanding what factors affect individuals' incentives to engage in protective behaviors will be of critical importance as we develop effective and humane policy, evaluate the current epidemic, make plans to emerge from it, and begin to prepare for future pandemics.

This paper examines factors predicting individual self-protecting behaviors during the Covid-19 pandemic in the United States. The factors we study include income, socio-demographic variables, pre-pandemic health characteristics, job and income losses, work arrangements (e.g., the ability to tele-work) and housing, along with beliefs and perceptions about the pandemic (e.g., whether individuals perceive social-distancing to be an effective measure and the consequences of infection). We study how these factors relate to several behavior changes, including social-distancing, mask-wearing and hand-washing. We focus on individuals living in the United States

using unique survey data collected during the third week of April 2020 and detailed in Belot et al. (2020). The data set follows roughly 6,000 individuals in 6 different countries and includes about 1,000 individuals from four different states in the U.S.: California, Florida, New York, and Texas.

We begin by documenting a striking and robust pattern apparent in the data: higher income is associated with higher levels of self-protective behaviors. Figure 1 illustrates this relationship. The figure considers three measures: (i) whether the respondent changed any behavior at all in response to the pandemic; (ii) whether the respondent increased social-distancing, which includes avoiding public spaces, running fewer errands, and visiting friends and family less often; and (iii) whether the respondent increased hand-washing or mask-wearing. The data have information on a host of additional behaviors (and similar income gradients emerge when we examine them). We choose to focus on these three measures because they illustrate the wide range of possible self-protective behaviors: the first is very broad, including any change in behavior at all; the second is a relatively high-cost activity; and the third is a relatively low-cost activity. Using these measures, we show that on average individuals in the fifth income quintile (quintile mean \$233,895) are between 13 and 19 percentage points (16–54%) more likely to engage in protective behaviors compared to individuals in the first income quintile (quintile mean \$13,775).¹ For each of these behaviors, the difference between the first and fifth income quintile is statistically significant at the 1% level.

¹Quintile means come from the [Tax Policy Center](#), administered by the Urban Institute and Brookings Institution.




Figure 1: Probability of Actions by Income Quintile: Proportion of respondents within an income quintile that changed behaviors, increased social distancing, or increased handwashing-mask wearing.

Our main analysis explores the relationships between income, the pandemic, and self-protective behaviors in two ways. First, we show how income relates to initial consequences of the pandemic along with other factors that could affect social-distancing and other self-protective behaviors. We show that lower-income respondents are more likely to lose their job or some portion of their income due to the pandemic (or expect to soon). They are also less likely to be able to work from home or to have access to open space where they reside.

Next, we examine which socio-demographic characteristics predict self-protecting behaviors. In particular, we estimate a series of linear probability models in which a self-protecting behavior is the outcome and the predictor variables include income along with different sets of additional variables, culminating in a final specification including income along with all variables. Several

patterns emerge. Work arrangements and housing characteristics, particularly transitioning to tele-working and access to outside space at home, are associated with adopting more self-protecting behaviors. We also find that income losses due to the pandemic are associated with increases in the behaviors that we examine. Surprisingly, we find no meaningful patterns between pre-existing health conditions and increases in self-protective actions. We also show that males and respondents from Florida and Texas are less likely to engage in self protective behaviors compared to females and respondents from California and New York, respectively. Moreover, beliefs about the effectiveness of social distancing along with perceived benefits (e.g., more time spent with family) correlate with increases in self-protective behaviors. Finally, we demonstrate that the income gradient is only partially explained by the inclusion of these variables. The size of the income coefficient estimates are fairly stable across specifications where different sets of controls are included.

Broadly, our findings are consistent with two key ideas. One, the initial economic consequences of the pandemic are particularly harmful to low-income individuals. Two, behaviors to stem the pandemic could place relatively large burdens on individuals with lower incomes. For example, higher-income individuals are more likely to report being able to work from home and more likely to have transitioned to tele-working instead of losing their job. As a result, self-protective behaviors, such social-distancing, are more practical, comfortable, and feasible for people with more income, which is evident in Figure 1.

This paper relates to a number of ongoing research efforts that are providing new information on the current pandemic and people's responses to it on a near-daily basis.² To the degree our questions and findings overlap, we provide crucial replication in an era of rapid-response, hastily-completed research. When questions or focus differs, we provide new information that other efforts do not. Finally, if we provide contradictory answers to similar questions, this is also important since it highlights where further research is needed. Existing research ranges from other commissioned surveys to formal economic models. For example, Adams-Prassl et al. (2020a) conducted a nationally representative survey of the UK, US, and German population on the labor market impacts of Covid-19. Ashraf (2020) explores the relationship between socio-economic factors, government policy, and Covid-19 health outcomes using a rich panel data set covering 80 countries.

²Because new research continues to emerge, we will augment our list of citations in later drafts of this paper.

The author finds a strong negative association between Covid-19 cases and socio-economic conditions, which can be alleviated by government policy. Adams-Prassl et al. (2020b) examines the impacts of Covid-19 lockdowns on mental health. Other studies have focused on the impacts of the Covid-19 pandemic on personal economic outcomes. Borjas (2020) examines how demographic factors influence testing and infections in the New York City area. Another example is Mongey, Pilossoph, and Weinberg (2020). The authors analyze which workers and jobs are most affected by social distancing measures. Other studies have focused on work arrangements such as work from home. Examples include Saltiel (2020), Okubo (2020), and Rahman (2020), which examine the types of occupations and workers that have access to work from home and how they have been affected by the Covid-19 pandemic. Closely related to this paper, Wozniak (2020) uses a unique U.S. data set, also publicly available, to show declines in well-being due to the pandemic and patterns between disease exposure and the decision to work or take protective measures. She finds people with Covid-19 exposure continue to work at similar rates as the non-exposed, and people with elevated risk for contracting the disease do not reduce work hours or take protective measures. We place greater attention on the associations between an individual's protective behaviors and their characteristics and beliefs, while she focuses more on protective behaviors through the lens of risks factors (i.e., an individual's susceptibility and potential spreading) and protective behaviors.

More broadly, this paper relates to a vast literature studying how socio-demographic characteristics associate with health and health behaviors. While our contribution reports evidence in a very specific context, the Covid-19 pandemic, it is noteworthy that many of the same relationships found in other health contexts are evident here.³ In other words, well-documented socio-demographic differences in health behaviors—and resulting health disparities—extend to the current pandemic (Yancy (2020)). Understanding this could help to inform optimal policy during a pandemic. Moreover, what we learn about behavior during a pandemic, a period when stakes are high and shifts in behavior are swift and large, can help us to understand health behavior differences more generally. As a concrete example, if we learn that certain types of work arrangements prevent social-distancing, such arrangements may prevent a host of other healthy behaviors that are unrelated to the pandemic. In this way, the current pandemic can provide useful directions for

³Cutler, Lleras-Muney, and Vogl (2011) provide an excellent summary and overview of the socioeconomic status-health gradient literature.

future research on health behaviors and health disparities.

Finally, this paper relates to a literature examining the tension between individual behavior and public health in the presence of externalities. A key historical analogy is the HIV epidemic.⁴ In that context, reduction of risky sex behavior not only protected individuals, but also slowed the spread of the virus, which is socially beneficial. The social benefit means there is a positive externality and thus potentially a sub-optimally low level of safer sex. In the current context, the tension between private behavior and public health is exacerbated by the fact that many of the people asked to incur the most brutal economic and social costs of protective behaviors face relatively low personal risk of serious health problems. This opens up a host of broad and general ethical questions about who should bear the greatest costs to protect public health. It also casts doubt on the sustainability of policies that presume full compliance.⁵

The remainder of this paper proceeds as follows. First, we discuss our data source and its preliminary analysis, highlighting patterns in the data. From there, we then discuss the results from our main analysis, which quantify the associations between individual characteristics and behavior changes. The final section concludes and describes further work in this area.

2 Data and Summary Statistics

We rely on recently-collected survey data from six different countries. In their accompanying paper (Belot et al. (2020)), the authors describe their sampling procedure and key features of the data. The United States sample, which is the focus of this paper, contains information on approximately one thousand individuals, roughly 250 from each of the following states: California, Florida, New York and Texas. The survey was constructed to be nationally representative along gender, age, household income, and race.⁶ While the survey is representative along these characteristics, it is not a random sample of people in the U.S., which means estimates must be interpreted with care. For example, it would be inappropriate at this stage to interpret links between economic factors

⁴Papageorge (2016) and Chan, Hamilton, and Papageorge (2016) are two studies which examine the history of the HIV epidemic to make broader points about health. More generally, Cawley and Ruhm (2011) provide a summary of the literature that examines cases where individuals make a trade off between controllable behaviors and health outcomes.

⁵Many of these ideas were explored in a recent blog post by DeLuca, Papageorge, and Kalish (2020). The current study builds on this piece, in part by using data to test hypotheses it put forth.

⁶The survey did not collect information about educational attainment.

and behavior as *causal*.

The first panel of Table 1 summarizes the socio-demographic characteristics and outcomes we study from the survey sample. 15% of respondents are non-white, 44% are male and about 39% are 56 years or older. 45% of respondents report at least one pre-existing health condition. Approximately 70% of respondents reported being employed. Among those that reported working, nearly 54% work full time, 14% work part time and 82% are able to work from home at the time of the survey when the pandemic was well underway. Also among these respondents, 34% report shifting to tele-work, 38% are no longer working, and 20% reported no change in their employment situation. On average, respondents lost \$770 in household income due to the pandemic.⁷

⁷Some respondents reported very large income losses. Given the difficulty in determining whether these are actual losses or survey errors, we use a dummy variable to flag these values. There are 173 of these observations. Our results are robust to the inclusion of these observations. The only difference is the magnitude on the lost income coefficient. The number in the table reflects the average of the remaining values.

Table 1: Summary Statistics

Characteristic	US Sample		P(Low Income)				P(Increased SD)			
	Prob	N	No	Yes	Δ	p	No	Yes	Δ	p
Socio-Demographics:										
Non-White	0.15	1006	0.33	0.47	-0.14	0.00	0.48	0.40	0.08	0.07
Male	0.44	1006	0.44	0.25	0.19	0.00	0.54	0.37	0.16	0.00
New York-California	0.54	1006	0.41	0.31	0.10	0.00	0.41	0.51	-0.1	0.00
Elderly	0.39	1006	0.31	0.42	-0.11	0.00	0.43	0.51	-0.07	0.03
Health:										
Pre-existing Condition	0.45	1006	0.37	0.34	0.03	0.40	0.47	0.45	0.02	0.58
Housing:										
Urban	0.41	1006	0.38	0.32	0.06	0.05	0.49	0.42	0.06	0.05
Home w/o Open Air Access	0.15	1006	0.32	0.58	-0.26	0.00	0.49	0.28	0.22	0.00
Elderly Exposure	0.46	1006	0.33	0.38	-0.05	0.11	0.44	0.49	-0.05	0.14
Work Arrangements and Losses:										
Working	0.70	990	0.49	0.29	0.20	0.00	0.51	0.44	0.07	0.06
Full-Time Working	0.54	701	0.49	0.13	0.36	0.00	0.45	0.44	0.02	0.64
Part-Time Working	0.14	701	0.27	0.43	-0.15	0.00	0.44	0.48	-0.04	0.46
Can Work From Home Working	0.82	701	0.39	0.27	0.11	0.01	0.45	0.44	0.01	0.92
Stopped Working Working	0.38	701	0.20	0.44	-0.24	0.00	0.44	0.45	-0.01	0.82
Tele-Working Working	0.34	701	0.40	0.09	0.31	0.00	0.42	0.49	-0.07	0.07
Still Working Working	0.20	701	0.28	0.36	-0.08	0.07	0.48	0.31	0.16	0.00
Mean Income Quintile	3.11	1006					0.43	0.51	-0.08	0.01
Mean Lost HH. Income (\$1,000)	0.77	475	0.40	0.30	0.10	0.00	0.45	0.48	-0.02	0.44
Beliefs and Perceptions:										
Social Distancing Effectiveness	3.99	1006	0.40	0.34	0.06	0.10	0.29	0.53	-0.24	0.00
Local Infection Rate	0.24	971	0.39	0.31	0.08	0.01	0.51	0.40	0.11	0.00
Benefits from Pandemic	0.84	1006	0.54	0.32	0.22	0.00	0.37	0.48	-0.11	0.01
Outcomes:										
Changed Behavior	0.88	1006	0.56	0.33	0.24	0.00	0.21	0.50	-0.29	0.00
Increased SD	0.46	967	0.39	0.30	0.09	0.00				
Increased Hand Washing-Mask Wearing	0.71	933	0.46	0.32	0.14	0.00	0.14	0.63	-0.48	0.00
N		1006	1006				967			

Values in the first panel represent the proportion of respondents with the listed characteristic for the entire US sample. The other panels are results from difference of means tests for the proportion of people that are low income or increased social distancing behaviors.

The survey contains two variables about labor status. The first asks about the current work arrangement and the second asks about changes due to the pandemic. Using these two variables, we construct a single measure that captures possible ways that the pandemic has affected individuals with different work arrangements. There are five possibilities: (i) “Never Worked” refers to individuals who were not working prior to or during the pandemic (e.g., retirees); (ii) “Stopped Working” refers to people who were working (full-time, part-time or self-employed) and stopped working due to the pandemic; (iii) “Began Tele-Working” refers to people who were employed prior

to the pandemic and are still employed, but transitioned to working from home due to the pandemic; (iv) “Still Working” describes individuals who were working prior to the pandemic and whose work status has not changed; and (v) “Other” includes people who report working prior to the pandemic and also report “other” when asked how their work status has changed. We should note that roughly 17 individuals do not fit into any of the categories above due to contradictory answers. As an example, some respondents report not working before the pandemic and also transitioning to tele-working, which is difficult to categorize.⁸

Figure A.1 summarizes this work status variable by age group and income quintile. As expected, we find that younger respondents tend to be more likely to begin tele-working whereas older workers are more likely to not be working before or during the pandemic. We also see that higher income people are more likely to transition to tele-working, while lower income people are more likely to either not be working before the pandemic or to have stopped working due to the pandemic. The share of respondents that were still working without any change was relatively stable across income quintiles. A deeper look into the survey finds an intuitive pattern. Those that were in lower income quintiles and reported still working belong to professions such as cashiers, packers and packagers, among others. Those in higher quintiles that reported still working included lawyers, computer programmers and managers. These professions differ substantially in what they pay, but include people who have been deemed “essential workers.” This explains the lack of an income gradient for this particular category.

The survey also includes information on beliefs and perceptions. Approximately 73% believe that social-distancing is either very effective or extremely effective, versus 24% who believe it is slightly or moderately effective. Only 3% of the sample believed social distancing was not effective at all. Almost 39% of respondents believed it to be extremely effective. On average, respondents believe that 24% of the people in their locality are infected. This high number is driven by a mass of respondents reporting implausibly high numbers (including some saying over 90%), which will be discussed below.⁹ Finally, about 84% perceive some benefits from the pandemic (e.g., getting to spend more time with family or reductions in pollution). As we show below, these perceptions

⁸If we include these observations in the “Other” category, our results remain unchanged.

⁹In some specifications, we experiment with dropping some extreme values under the assumption that they reflect respondent confusion. However, doing so does not alter our results.

predict behavior.

The survey measures behavior change in two ways. First, respondents are asked directly if they engaged in any self-protective behavior in response to the pandemic. About 88% report having done so. Second, the survey asks respondents to report how frequently they engaged in 15 different activities before the pandemic, at the start of the pandemic, and a few weeks after the pandemic began. These behaviors ranged from hand washing and eating healthy to visiting large open or closed spaces and visiting friends and family. Given this data structure, we observe how the respondents changed their behavior over time. For each behavior, respondents could answer (1) never, (2) rarely, (3) sometimes, (4) very often, or (5) always. Figure A.2 in the Appendix summarizes the average frequency with which respondents engaged in some of these behaviors at each time period. Consistent with Figure 2, we see an increase in the average frequency of self protective behaviors from before the pandemic to weeks after the pandemic started. To get a sense of behavior change, we construct a count variable for each individual for the number of changes towards (or away from) self-protection from before the pandemic to a few weeks after it had begun, when the data were collected. Figure 2 plots the distribution of the resulting variable. The changes are normalized such that self-protecting changes are positive, while reducing such behaviors is recorded as negative.¹⁰ The median number of changes towards self-protection is 9. We highlight two main takeaways from Figure 2. First, very few people exhibit a net decline in self-protective behaviors. Second, most people either do not change their behavior much at all (note the spike in the distribution at zero) or make a fairly large number of changes. Figure 2 also presents the densities for the lowest and highest income quintiles. We find that a greater mass of low income people are not increasing their self protective behaviors. We see the opposite pattern for high income people, with a greater share of high income people increasing their self protective behaviors.

¹⁰To fix ideas, suppose an individual answered 3 (sometimes) for hand washing and 3 (sometimes) for taking public transport before the pandemic and 5 and 4, respectively, during the pandemic. The increase in hand washing accounts for 2 increases in self protective behaviors and the increase in taking public transportation is recorded as a 1 increment *decrease* in self protective behaviors. Taken together this results in a net effect of $2 - 1 = 1$.




Figure 2: Distribution of Behavior Changes: Distribution of behavior changes from before the pandemic to a few weeks after the pandemic started. Changes are normalized such that self-protecting changes are positive and reductions in these behaviors are negative. The densities for the lowest and highest income quintiles are broken out separately.

In our subsequent analysis, we focus on three measures of behavior change. The first is from the aforementioned question asking respondents directly if they changed their behavior in response to the epidemic. The second is a composite behavior variable for social distancing. We focus on: visiting large open spaces, visiting large closed spaces, attending large social gatherings, and visiting with friends or family. As we are mainly interested in how behaviors change, we look at how this social distancing composite changes over time.¹¹ An increase in social distancing is defined as an above-median increase in the number of self-protecting improvements for the four behaviors of interest. The third measure focuses on changes in hand-washing and in wearing a mask.

¹¹Note that we drop respondents that were never engaging in these activities (i.e., were always social distancing) prior to the pandemic since they would have no way of increasing their social distancing measures in response to the pandemic. This means we lose 40–73 individuals, depending on the specification.

In these analyses, we make four types of drops from the full U.S. survey sample for our analysis. First, we drop respondents who did not report an income quintile or gender (15 observations). Second, as discussed previously, our composite behavior dependent variables examine increases in behaviors. We drop respondents that were already engaging in these behaviors and had no ability to increase further (40 observations for social distancing and 73 for hand washing-mask wearing). The third set of drops are the outliers in local infection rate beliefs (36 observations). Finally, we drop about 17 respondents that did not map into our composite work status variable. For each dependent variable, we hold the analysis sample constant across specifications to facilitate comparisons. Thus for each outcome, we take the maximum number of observations that remain after these sets of drops.

3 Results

3.1 Connecting Socio-demographic Characteristics to Income

The second panel of Table 1 summarizes the difference in means of several characteristics broken-out by income groups. Here, we have defined high income as the top three quintiles and low income as the bottom two quintiles. We find significant differences for most of these characteristics between high income and low income respondents. For example, non-white respondents were more likely to be low income than white respondents. We find no significant income differences for pre-existing health conditions, exposure to a person 56 years or older, and beliefs in the effectiveness of social distancing. We also see that lower income people have a significantly smaller probability of increasing self-protective behaviors.

Figure A.3 in the Appendix explores expected losses to labor and household income by labor status and income quintile. Expected losses to labor income are a much larger share of income for low income respondents. For example, people in the first income quintile reported expected labor income losses of over 10% while respondents in the fifth income quintile expected losses of no more than 5%. We observe a similar pattern when looking at expected household income losses. First quintile expected losses range from nearly 20% to about 25% while fifth quintile losses range from 10% to just under 15%. We also find that the difference between the mean expected labor income

loss for the first and fifth income quintile is statistically significant. Figure 3 assesses losses that have already occurred. We observe a similar relationship between income quintile and the magnitude of income losses. The second panel examines changes in work status. We see that transitions to tele-working rise with income. We observe the reverse pattern for low income people, who were most likely to have stopped working altogether. The third panel digs into job losses due to the pandemic. We find that the lowest-income respondents had the lowest amount of job security and the highest probability of temporary unemployment. An interesting pattern is that higher income individuals were the most likely to permanently lose their jobs, despite also having the highest level of job security across all income quintiles. This may reflect selection: higher-income jobs are more secure in general so a job loss reflects a large and permanent shift, e.g., a bankruptcy. Similar to expected income losses, the difference between mean household income losses for the first and fifth income quintile was statistically significant.

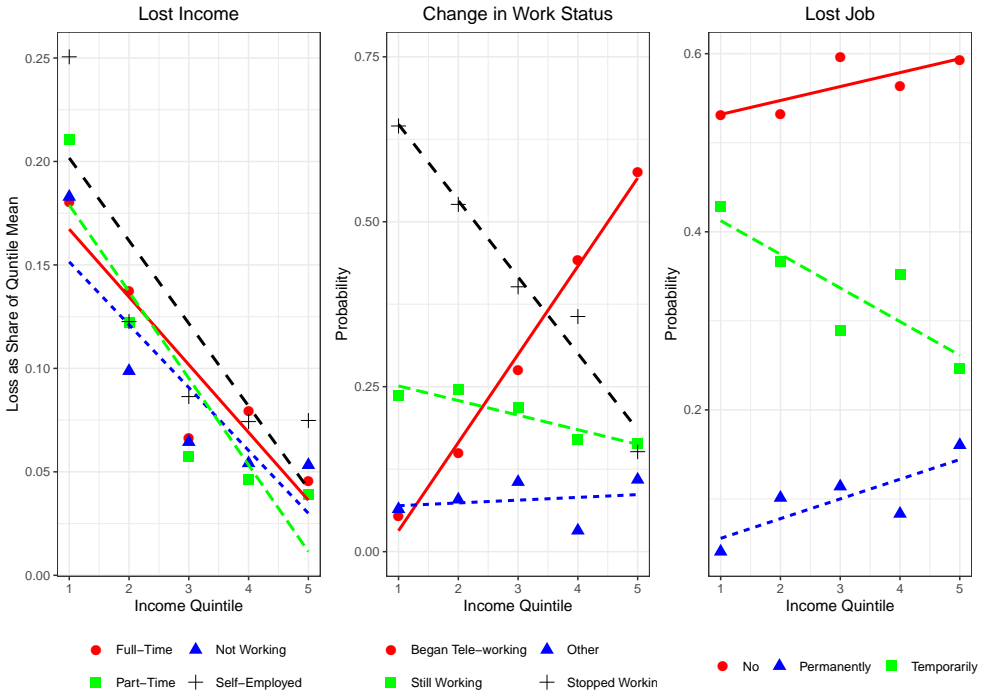


Figure 3: Realized Losses and Labor Market Changes by Income: Realized income losses by income quintile. Income losses are normalized by the mean income of the respondent’s quintile. Changes in work status are the proportion of respondents within an income quintile. Job losses are proportions of each type of labor transition within an income quintile.

Next, we consider work arrangements by income. Figure A.4 consists of two panels, which plot labor status (full-time, part-time, self-employed or not working) and how well they can work from home, respectively. According to the figure, full-time employment and the ability to work from home rise with income. Lower-income people are more likely to report either that they stopped working or that they experienced no change in work status (which includes not having switched to tele-working). For example, nearly 75% of respondents in the fifth income quintile are working full-time. About the same percentage of respondents in the first quintile are not working. From the center panel of Figure 3, we see that more than 50% of respondents in the fifth income quintile reported transitioning to tele-work post-pandemic, whereas only 10% of those in the first quintile did so. This pattern appears consistent with the fact that nearly 40% of low income respondents

report being unable to work from home.¹²

The other broad categories of socio-demographic characteristics that are potentially associated with income and behavior changes are pre-existing health conditions, household features, and beliefs related to the pandemic. From Table 1 we know that 45% of survey respondents reported having a pre-existing health condition such as diabetes, high blood pressure, heart disease, or asthma. Given strong income-health gradients found in other literature, this pattern is surprising and may reflect that the sample is not representative of the U.S. in terms of the relationship between income and health. In contrast, housing is strongly related to income. Higher-income respondents are far more likely to live in homes (versus apartments) and to have access to open air where they reside compared to lower-income respondents. As we noted earlier, we do not observe a pattern between income and beliefs in the effectiveness of social distancing.¹³ Across all income quintiles, 70%–78% of respondents believe social distancing is either very effective or extremely effective.¹⁴

The survey data also contain respondents' beliefs about various rates related to the disease such as infection rates, likelihood of contracting the disease, and so on. However, it is difficult to interpret responses. For example, consider beliefs about the local infection rate. The distribution of these beliefs is presented in Figure A.5. We can see that many respondents report implausibly small and large numbers. This could reflect several factors, including misinformation about the spread of illness, difficulties with probabilistic thinking, which is well-documented in the literature (see e.g., Barth, Papageorge, and Thom (2020), Lillard and Willis (2001), Delavande, Perry, and Willis (2006), etc.), fatalistic beliefs (e.g., Akesson et al. (2020)), optimism about herd immunity, etc.¹⁵ In any case, these interpretational difficulties will limit conclusions we can draw using some of the beliefs variables.

¹²We also ran a multinomial logit model where the outcome variable is our work status variable to study which factors predict which types of work changes and find that these patterns hold. These results are available upon request.

¹³We measure a respondent's belief in social distancing as the average of their beliefs about the effectiveness of shutting down non-essential businesses, limiting mobility outside the home, and forbidding mass gatherings.

¹⁴These patterns are explored further in Tables A.1, A.2, A.3, and A.4.

¹⁵Indeed, we must leave open the possibility that, due to our lack of knowledge about the illness, implausibly-high answers are correct.

3.2 Factors Associated with Behavior Change

Our analysis until now shows that several factors are related to income. These factors could potentially help to explain differences across income groups in self-protective behaviors depicted in Figure 1. In this section, we explore which socio-demographic characteristics are associated with behavior changes.

To begin, the third panel of Table 1 summarizes the difference in means across different characteristics for those who increased social distancing behavior according to our metric. We find significant differences between males and females and between those who believe in the effectiveness of social distancing. These findings suggest that women are more likely than men to increase social distancing as are those who believe strongly in the effectiveness of social distancing. There are significant differences between those who have not stopped working and those that are still working. We do not find any other significant differences across individual characteristics. We also see that people increasing social distancing behaviors have a significantly larger probability of increasing other self protective behaviors.

For our main analysis, we examine three outcomes: any behavior change, social distancing, and mask-wearing or hand-washing. For each outcome, we estimate linear probability models as a function of income and different sets of explanatory variables. Our main findings are summarized in Table 2. In general, we find that income, work arrangements such as tele-working, lost income, and beliefs about the effectiveness of social distancing are significantly associated with the self-protective measures we examine. Detailed results are presented in Tables A.5, A.6, and A.7 in the Appendix. In each table, all columns include income quintiles as explanatory variables. Column (1) includes only income, Column(2) adds in socio-demographic characteristics, Column (3) adds pre-existing health conditions, Column (4) brings in housing characteristics, Column (5) introduces work arrangements and economic loss characteristics, Column (6) adds in beliefs about social distancing and local infection rates and perceived benefits from the pandemic. Finally, in Column (7) we include all of these sets of controls in a single specification. We will discuss each of these columns in the following subsections.¹⁶

¹⁶We examined other characteristic associations beyond what is presented in the paper. These results are available upon request.

Table 2: Summary of Factors Associated with Self-Protective Behaviors

Variable	Metric	CB	SD	WM
Income	Significant	✓	✓	✓
	Sign	+	+	+
Not Working	Significant	✓	✗	✓
	Sign	+	NA	+
Tele-Working	Significant	✓	✓	✓
	Sign	+	+	+
Not Working Pre-Covid	Significant	✗	✓	✗
	Sign	NA	+	NA
Lost Income	Significant	✓	✓	✓
	Sign	+	+	+
Effectiveness of SD	Significant	✓	✓	✓
	Sign	+	+	+
Area Infection Rate	Significant	✗	✓	✓
	Sign	NA	-	-
Benefits	Significant	✓	✓	✓
	Sign	+	+	+

This table summarizes the core findings from linear probability models examining the association between socio-demographic characteristics and three different self-protective behaviors. Detailed tables with co-efficient estimates, standard errors, significance levels, and other controls can be found in the Appendix.

3.2.1 Income

Across all three of our dependent variables we find strong, statistically significant associations with income. Higher income individuals are more likely to engage in the behaviors we examine. To fix ideas, relative to the first income quintile, a member of the fifth income quintile is 10%–15% more likely to change their behaviors, 14%–28% more likely to increase social distancing behaviors, and 18%–28% more likely to increase hand washing or mask wearing. We find that these income effects are fairly robust to the inclusion of controls. From the baseline to the case where we include all of our controls, the size of the coefficient estimates remain fairly stable as we add additional variables, which means that these other factors do not fully explain the income gradient.

3.2.2 Gender, Age, Race and Location

The next set of controls we include are gender, age, race and state. We do not find many significant associations between these factors and the change behaviors outcome. In the baseline case we see negative associations between males, people 56 years or older, and some regional effects. However, most of these relationships disappear when other variables are added to the analysis. We find more robust significance for the increased social distancing behavior. We find strong negative associations between males, and respondents from Florida and Texas, which maintain significance once other controls are added. To fix ideas, we find that males are 15% less likely than females to increase social distancing. Similarly, we find that relative to respondents from California, people in Texas and Florida tended to be 16% and 12% less likely to increase social distancing, respectively. These results may presage the surges in Covid-19 cases that happened in these two states that began toward the end of June 2020.¹⁷ Finally, we find positive significant associations between race and people 56 years or older for the hand washing-mask wearing outcome. Specifically, we find that Black respondents are 13% more likely than white respondents to increase hand washing or mask wearing.¹⁸ We find a similarly sized relationship for those 56 years or older. It is interesting that we pick up these effects for increased hand washing or mask wearing. This may be reflect the fact that of the three activities we examine, this one is a relatively low-cost way to self protect for people who face risks, but are unable to engage in higher-cost, less practical activities, such as social-distancing.

3.2.3 Health

We also examine various pre-existing health conditions, including diabetes, high blood pressure, heart disease, asthma, allergies, and other conditions. Overall, and surprisingly, these variables are not strongly correlated to behavior change.¹⁹ Yet, it is surprising that health conditions more strongly associated with serious illness (e.g., diabetes, asthma, or high blood pressure) are not

¹⁷This finding complements recent work by Makridis and Rothwell (2020), which found a significant relationship between political affiliation and beliefs about the Covid-19 pandemic and the adoption of social distancing behaviors.

¹⁸A data set that over-samples non-white individuals would potentially reveal other differences by race or ethnicity in the likelihood of engaging in self-protective behaviors.

¹⁹Oddly, we find a strong *negative* association between heart disease and increased social distancing, which may reflect that people with heart disease are generally unhealthy and thus less likely to engage in self-protective behaviors.

associated with behavior change. An exception is that we find a robust significant association between allergies and increases in hand-washing and mask-wearing. A number of factors could explain this, including the possibility that people with allergies could feel they are becoming ill even if they are not and thus be more willing to take precautions.

3.2.4 Housing

Next we examine housing characteristics. We find a negative significant relationship between respondents in the countryside and changing behaviors but this association becomes indistinguishable from zero as other controls are added. We find a robust negative association for having no access to open air and increased social distancing behavior. In our full control case, we find that respondents that live in homes without open air access are 12% less likely to increase social distancing behaviors. We find this to be an intuitive result. People who are more comfortable sheltering-in-place are more likely to do it. Policies aiming to slow the pandemic should take these factors into account as they suggest cramped and uncomfortable housing can potentially undermine efforts to “flatten the curve.” We find similar patterns for the two other outcome variables, but they lose significance in the final specification, where we add additional variables.

3.2.5 Work Arrangements and Losses

We also consider work arrangements and economic losses. In general we find fairly consistent results across all three of our outcome variables. People who transitioned into tele-working are more likely to change behaviors, increase social distancing, and increase hand washing-mask wearing. This association ranges from roughly 9%–12% relative to somebody who continued to work. This effect is robust to the inclusion of other controls. We find a similarly-sized effect for those that stopped working or never worked but significance was retained with less consistency. This result is intuitive. People who can work from home are more likely to abide by stay-at-home orders. Factors related to work arrangements, which vary across socio-demographic groups, can determine the sustainability and effectiveness of policies aiming to prevent the spread of illness.²⁰ We also find that realized

²⁰These estimates can be converted into moments perhaps usable in an epidemiological model. For example, across income quintiles, the probability that an average respondent would increase social distancing behavior if they began tele-working ranged from about 33% to 57%.

household income losses have a significant positive association with each of these behaviors. After controlling for extreme lost income values, we find that for every \$1,000 lost a respondent is 1%–4% more likely to adjust each of the behaviors we examined. People who have experienced these losses have already been harmed by the pandemic. As a result, they may be more careful than others and view contracting the disease as a higher risk. Another possibility is that these people have fewer monetary resources and may not have money to cover medical expenses if they were to contract the disease.

3.2.6 Beliefs and Perceptions

The final set of variables we examined were beliefs and perceptions about the pandemic. Reassuringly, we find a fairly consistent effect for beliefs in the effectiveness of social distancing across the three behaviors we included. These findings are strongest for the changed behaviors and increase social distancing variables. We find similar results but with weaker significance for increased hand washing-mask wearing. As we previewed earlier, we find a significant negative association between beliefs about the local infection rate and each of the three outcome variables. We discussed possible reasons why respondents may have reported these implausibly high beliefs earlier. Our findings are robust to removing this variable as well as using a dummy variable to control for extremely high beliefs.²¹ We also find some positive associations between perceived benefits from the pandemic and increases in our behaviors of interest. Less pollution and more family time were two that came out as significant and tended to retain significance as other controls were added. In Appendix Table A.8, we present cross-tabulations of the survey data which indicate most of the people identifying these benefits belonged to higher income quintiles and were non-white.

One finding that surprised us was the negative association between beliefs about local infection rates and increases in self protecting behaviors. As discussed previously, the distribution of respondent beliefs about local infection rates has significant mass at the low and high end, which are difficult to reconcile with reality. In Figures A.6, A.7, and A.8 we present lowess smoother results for three behavioral outcomes of interest and this belief. In each case, people who reported an infection rate of 20% or fewer exhibit the expected response: a rise in perceived infection rates is

²¹We defined the cutoff for “extremely high” as any belief above 20%.

associated with more protective behavior. Thus, negative coefficient estimates are driven by people with implausibly high perceptions of infection rates. This could reflect respondent confusion. It could also reflect a sort of fatalism, i.e., people believe infection rates are so high that they are bound to become infected, too, and thus don't bother to engage in protective behaviors. Fatalism is a well documented phenomenon in several fields, see e.g., Akesson et al. (2020), Ferrer and Klein (2015), Shapiro and Wu (2011), etc.²²

3.3 Robustness Checks

We conducted a series of robustness checks to these specifications. The results for these analyses are available upon request. We considered whether a respondent was engaging in these behaviors at all following the start of the pandemic. We also looked at whether there were distinct behavior differences between those that had experienced a loss due to the pandemic and the pooled sample. Another analysis examined each state individually and pooled groups of states. In general, these analyses aligned with our main results or were statistically indistinguishable from zero. We also considered the intensive margin for increased social distancing and hand washing or mask wearing behaviors. A respondent's income and beliefs about the effectiveness of social distancing did not have a significant association with larger increases in either of these self-protecting behaviors. Other effects are consistent with our main analysis. Finally, as mentioned previously, the data collected by Belot et al. (2020) do not contain information on educational attainment. We use information on a respondent's profession to construct a proxy for whether they have a college degree and incorporate it into our specifications.²³ We find that education is positively associated with increases in self-protective behaviors. The inclusion of this variable does not appreciably alter our other findings.

4 Conclusion

While many of the questions raised and discussed in this paper focus on a specific point in time, the Covid-19 pandemic will eventually run its course. However, it would be shortsighted and naive to think that another virus, perhaps an even more damaging one, will not come about in the future.

²²Figures A.9, A.10, and A.11 present similar figures using the lpolynomial smoother.

²³According to this variable, 30% of the sample is college educated and the average income is 3.88 for the college educated and 2.78 for people without a college education.

Indeed, many specialists believe that this virus will be cyclical, returning annually. If so, the questions we are addressing now will be important not only as we move through the current crisis, but also as we begin to prepare for the next one. Social scientists who study behavior—and the policies that affect it—must play a critical role in these efforts. One way is through the collection and analysis of new survey data, which shed light on what behavior can be expected of different segments of the population during a pandemic given heterogeneity in the incentives, constraints and circumstances people face. These data could be used not only to describe behavior, but also in more targeted research projects, such as: examining how information is transmitted and belief about the pandemic are formed and affect behavior, analyzing location-specific policy responses and their relative merit, and calibrating epidemiologically-grounded models relating variation in individual behavior to the spread of illness, among many others.

References

- Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh (Apr. 2020a). *Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys*. Working Paper 2018. Cambridge-INET.
- (May 2020b). *The Impact of the Coronavirus Lockdown on Mental Health: Evidence from the US*. Working Paper 2020-030. Human Capital and Economic Opportunity Working Group.
- Akesson, Jesper, Sam Ashworth-Hayes, Robert Hahn, Robert D Metcalfe, and Itzhak Rasooly (May 2020). *Fatalism, Beliefs, and Behaviors During the COVID-19 Pandemic*. Working Paper 27245. National Bureau of Economic Research.
- Alon, Titan M, Matthias Doepke, Jane Olmstead-Rumsey, and Michele Tertilt (Apr. 2020). “The Impact of COVID-19 on Gender Equality”. In: *Covid Economics* 4, pp. 62–85.
- Ashraf, Badar Nadeem (July 2020). “Socioeconomic Conditions, Government Interventions and Health Outcomes During COVID-19”. In: *Covid Economics* 37, pp. 141–162.
- Barth, Daniel, Nicholas W Papageorge, and Kevin Thom (2020). “Genetic Endowments and Wealth Inequality”. In: *Journal of Political Economy* 128.4, pp. 1474–1522.
- Belot, Michele, Syngjoo Choi, Julian Jamison, Nicholas W Papageorge, Egon Tripodi, and Eline van den Broeck (May 2020). “Six-country Survey on Covid-19”. In: *Covid Economics* 17, pp. 205–219.
- Borjas, George J (Apr. 2020). “Demographic Determinants of Testing Incidence and COVID-19 Infections in New York City Neighborhoods”. In: *Covid Economics* 3, pp. 12–39.
- Cawley, John and Christopher J Ruhm (2011). “The Economics of Risky Health Behaviors”. In: *Handbook of Health Economics*. Vol. 2. Elsevier, pp. 95–199.
- Chan, Tat Y, Barton H Hamilton, and Nicholas W Papageorge (2016). “Health, Risky Behaviour and the Value of Medical Innovation for Infectious Disease”. In: *The Review of Economic Studies* 83.4, pp. 1465–1510.
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber (June 2020). “The Cost of the COVID-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending”. In: *Covid Economics* 29, pp. 1–49.
- Cutler, David M., Adriana Lleras-Muney, and Tom Vogl (Apr. 2011). “Socioeconomic Status and Health: Dimensions and Mechanisms”. In: *The Oxford Handbook of Health Economics*. Oxford University Press.
- Delavande, Adeline, Michael Perry, and Robert J Willis (2006). *Probabilistic Thinking and Social Security Claiming*. Working Paper 129. Michigan Retirement Research Center.
- DeLuca, Stephanie, Nicholas W Papageorge, and Emma Kalish (Apr. 2020). *The Unequal Cost of Social Distancing*. Available at <https://coronavirus.jhu.edu/from-our-experts/the-unequal-cost-of-social-distancing>.
- Fairlie, Robert W, Kenneth Couch, and Huanan Xu (May 2020). *The Impacts of COVID-19 on Minority Unemployment: First Evidence from April 2020 CPS Microdata*. Working Paper 27246. National Bureau of Economic Research.
- Ferrer, Rebecca A and William MP Klein (2015). “Risk Perceptions and Health Behavior”. In: *Current Opinion in Psychology* 5, pp. 85–89.
- Lillard, Lee A and Robert J Willis (2001). *Cognition and Wealth: The Importance of Probabilistic Thinking*. Working Paper 7. Michigan Retirement Research Center.
- Makridis, Christos A. and Jonathan T. Rothwell (July 2020). “The Real Cost of Political Polarization: Evidence from the COVID-19 Pandemic”. In: *Covid Economics* 34, pp. 50–87.
- Manski, Charles F (1999). *Identification Problems in the Social Sciences*. Harvard University Press.

- Mongey, Simon, Laura Pilossoph, and Alex Weinberg (May 2020). “Which Workers Bear the Burden of Social Distancing Policies?” In: *Covid Economics* 12, pp. 69–86.
- Okubo, Toshihiro (June 2020). “Spread of COVID-19 and Telework: Evidence from Japan”. In: *Covid Economics* 32, pp. 1–25.
- Pampel, Fred C, Patrick M Krueger, and Justin T Denney (2010). “Socioeconomic Disparities in Health Behaviors”. In: *Annual Review of Sociology* 36, pp. 349–370.
- Papageorge, Nicholas W (2016). “Why Medical Innovation is Valuable: Health, Human Capital, and the Labor Market”. In: *Quantitative Economics* 7.3, pp. 671–725.
- Posner, Richard A and Tomas J. Philipson (1993). *Private Choices and Public Health: The AIDS Epidemic in an Economic Perspective*. Harvard University Press.
- Rahman, Ahmed S. (July 2020). “Why Can’t Everybody Work Remotely? Blame the Robots.” In: *Covid Economics* 36, pp. 105–128.
- Saltiel, Fernando (Apr. 2020). “Who Can Work From Home in Developing Countries?” In: *Covid Economics* 6, pp. 104–118.
- Shapiro, Joel and Stephen Wu (2011). “Fatalism and Savings”. In: *The Journal of Socio-Economics* 40.5, pp. 645–651.
- Viner, Russell M, Simon J Russell, Helen Croker, Jessica Packer, Joseph Ward, Claire Stansfield, Oliver Mytton, Chris Bonell, and Robert Booy (2020). “School Closure and Management Practices During Coronavirus Outbreaks Including COVID-19: A Rapid Systematic Review”. In: *The Lancet Child & Adolescent Health*.
- Wozniak, Abigail (May 2020). *Disparities and Mitigation Behavior During COVID-19*. Institute Working Paper 32. Opportunity and Inclusive Growth Institute: Federal Reserve Bank of Minneapolis.
- Yancy, Clyde W. (May 2020). “COVID-19 and African Americans”. In: *Journal of the American Medical Association* 323.19, pp. 1891–1892.

A Additional Figures and Tables

A.1 Figures




Figure A.1: Probability of Work Arrangement by Age Group and Income Quintile: Proportion of respondents within an age group or income quintile that are classified into our work arrangements variable.




Figure A.2: Average Frequency of Select Behaviors by Income Quintile and Time Period: These tabulations are the average frequency respondents within an income group reported doing each listed activity. Frequencies ranged from 1 (Never) to 5 (Always). These are calculated for the time period before the pandemic, at the start of the pandemic, and a few weeks after the pandemic.




Figure A.3: Expected Income Losses by Income and Labor Status: Expected labor and household income losses across income quintiles. Income losses are normalized by the mean income of the respondent's quintile.

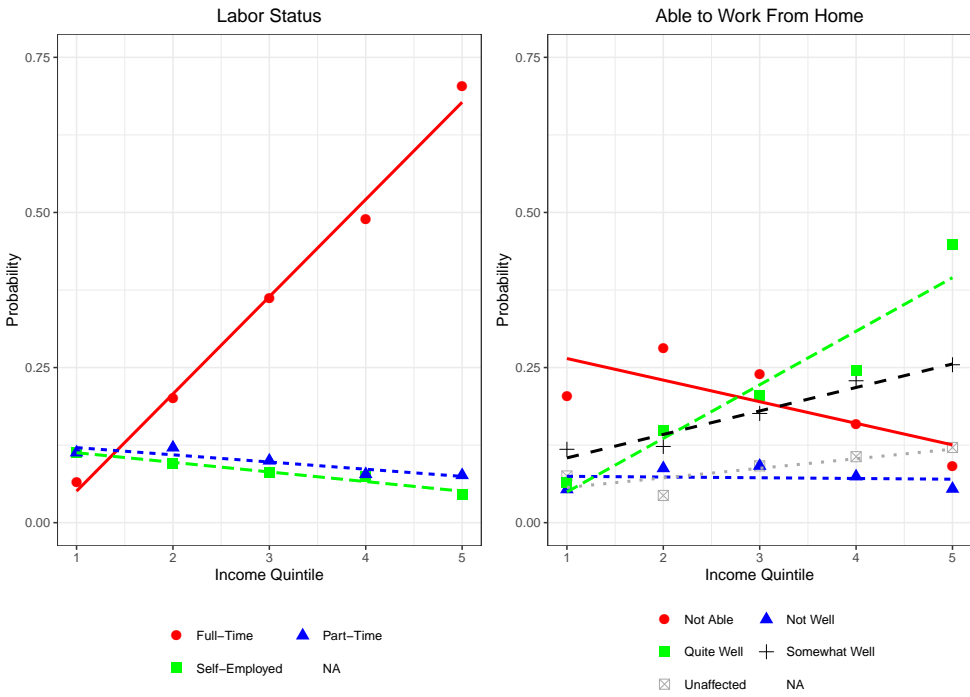


Figure A.4: Probability of Labor Characteristic by Income Quintile: Proportion of respondents within an income quintile with various labor characteristics.




Figure A.5: Distribution of Beliefs About Local Infection Rates: Distribution of the reported beliefs of respondents about the rate of infections within their local community.




Figure A.6: Lowess Smoother Changed Behaviors: Lowess smoother applied to the binary changed behaviors outcome to beliefs about the local infection rate.




Figure A.7: Lowess Smoother Increased Social Distancing: Lowess smoother applied to the binary increased social distancing variable to beliefs about the local infection rate. Increased social distancing is defined as an above median increase in self protective behaviors (e.g., visiting public spaces, visiting friends and family, etc.) from before the pandemic to a few weeks after it started.




Figure A.8: Lowess Smoother Increased Hand Washing-Masking Wearing: Lowess smoother applied to the binary increased hand washing-mask wearing variable to beliefs about the local infection rate. Increased hand washing-mask wearing is defined as an above median increase in hand washing or mask wearing behavior from before the pandemic to a few weeks after it started.




Figure A.9: Lpolyci Smoother Changed Behaviors: Lpolyci smoother applied to the binary changed behaviors outcome to beliefs about the local infection rate.

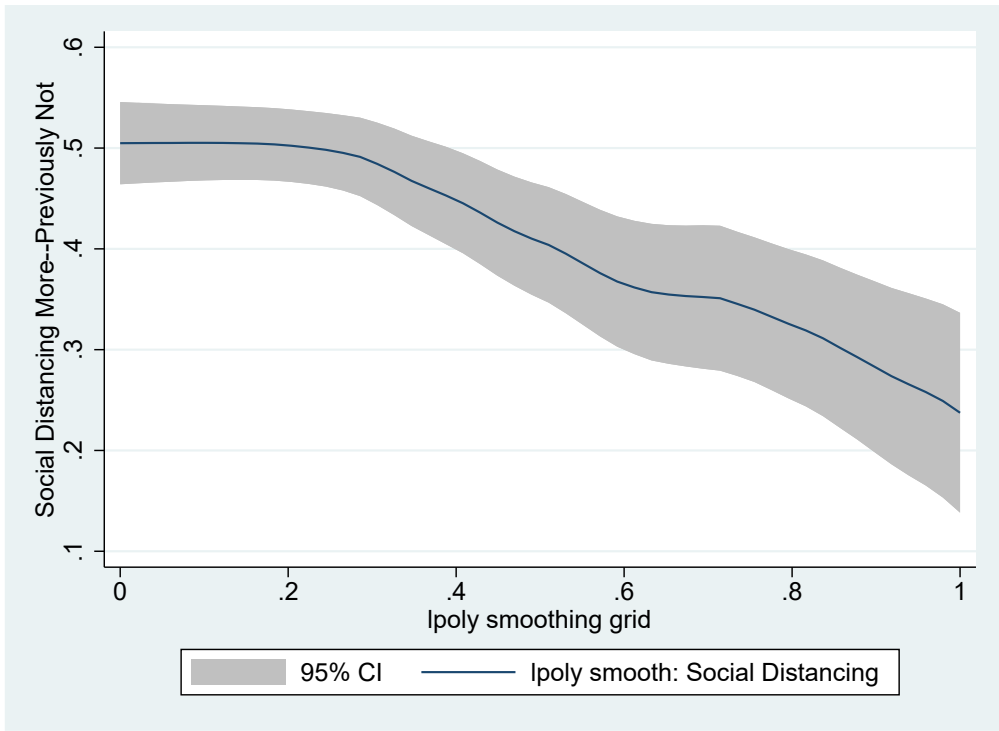


Figure A.10: Lpolyci Smoother Increased Social Distancing: Lpolyci smoother applied to the binary increased social distancing outcome to beliefs about the local infection rate. Increased social distancing is defined as an above median increase in self protective behaviors (e.g., visiting public spaces, visiting friends and family, etc.) from before the pandemic to a few weeks after it started.




Figure A.11: Lpolyci Smoother Increased Hand Washing-Mask Wearing: Lpolyci smoother applied to the binary increased hand washing-mask wearing outcome to beliefs about the local infection rate. Increased hand washing-mask wearing is defined as an above median increase in hand washing or mask wearing behavior from before the pandemic to a few weeks after it started.

A.2 Tables

Table A.1: Pre-Existing Health Conditions by Income

Income	Diabetes	High Blood Pressure	Heart Disease	Asthma
First quintile	0.14	0.3	0.06	0.13
Second quintile	0.14	0.32	0.05	0.12
Third quintile	0.19	0.32	0.05	0.14
Fourth quintile	0.16	0.28	0.05	0.15
Fifth quintile	0.25	0.21	0.05	0.14

Values are share of respondents within each income quintile that reported having the listed pre-existing health condition. Note respondents may have multiple conditions.

Table A.2: Current Living Area by Income

Current Living Area	First	Second	Third	Fourth	Fifth
Urban	0.41	0.33	0.29	0.41	0.62
Semi-urban / residential	0.41	0.52	0.6	0.49	0.35
Countryside	0.18	0.15	0.11	0.1	0.03
Total	168	190	210	243	196

Cross tabulation of US sample for current living area by income quintiles.

Table A.3: Home Characteristics by Income

Variable	First	Second	Third	Fourth	Fifth
Current Home:					
House	0.43	0.56	0.63	0.78	0.77
Apartment	0.39	0.29	0.24	0.15	0.18
Condominium	0.03	0.07	0.07	0.04	0.05
Trailer	0.1	0.07	0.04	0.01	NA
Shelter	0.02	NA	NA	NA	NA
Other	0.03	0.02	0.01	0.02	0.01
Total	168	190	210	243	196
No Open Air Access:					
Not selected	0.67	0.83	0.88	0.92	0.9
Selected	0.33	0.17	0.12	0.08	0.1
Total	168	190	210	243	196

Cross tabulation of US sample for home characteristics by race.

Table A.4: Belief in Social Distancing by Income

Belief in Social Distancing	First	Second	Third	Fourth	Fifth
Not effective at all	0.03	0.04	0.03	0.02	0.03
Slightly effective	0.1	0.11	0.07	0.06	0.04
Moderately effective	0.17	0.16	0.17	0.19	0.15
Very effective	0.31	0.32	0.35	0.35	0.39
Extremely effective	0.39	0.38	0.38	0.38	0.39
Total	168	190	210	243	196

Cross tabulation of US sample for belief in social distancing by income quintiles.

Table A.5: Factors Associated with Changing Behaviors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Income:							
Second quintile	-0.00 (0.04)	-0.00 (0.04)	-0.01 (0.04)	-0.01 (0.04)	-0.01 (0.04)	0.00 (0.04)	-0.00 (0.04)
Third quintile	0.09** (0.04)	0.09** (0.04)	0.09** (0.04)	0.08** (0.04)	0.07* (0.04)	0.08** (0.04)	0.07* (0.04)
Fourth quintile	0.11*** (0.04)	0.12*** (0.04)	0.11*** (0.04)	0.10*** (0.04)	0.09*** (0.04)	0.10*** (0.04)	0.09** (0.04)
Fifth quintile	0.14*** (0.04)	0.15*** (0.04)	0.15*** (0.04)	0.12*** (0.04)	0.11*** (0.04)	0.12*** (0.03)	0.10*** (0.04)
Socio-Demographics:							
Black		-0.04 (0.04)					-0.02 (0.04)
Other		-0.04 (0.05)					-0.05 (0.05)
Male		-0.06** (0.02)					-0.03 (0.02)
Florida		0.01 (0.03)					-0.00 (0.03)
New York		-0.03 (0.03)					-0.04 (0.03)
Texas		-0.05 (0.03)					-0.04 (0.03)
56 or Older		-0.05** (0.02)					-0.02 (0.03)
Health:							
Diabetes			-0.05 (0.03)				-0.05* (0.03)
High Blood Pressure			-0.03 (0.03)				-0.01 (0.03)
Heart Disease			0.03 (0.05)				0.02 (0.05)
Asthma			0.04 (0.03)				0.04 (0.03)
Allergies			0.04 (0.02)				0.02 (0.02)
Other Condition			0.06**				0.06**

	(0.03)	(0.03)
Housing:		
Semi-urban / residential	-0.03 (0.02)	-0.03 (0.02)
Countryside	-0.08** (0.04)	-0.05 (0.04)
No Access to Open Air	-0.04 (0.03)	0.00 (0.03)
Work Arrangements-Losses:		
Stopped Working	0.10** (0.04)	0.07 (0.04)
Began Tele-Working	0.12*** (0.04)	0.09** (0.04)
Not Observed Working	0.07* (0.04)	0.04 (0.04)
Other	0.18*** (0.04)	0.13*** (0.04)
Flag extreme lost income	-0.05** (0.03)	-0.03 (0.03)
Lost HH Inc./\$1,000 Orig	-0.00 (0.00)	0.00 (0.00)
Lost HH Inc./\$1,000 Adj.	0.02*** (0.01)	0.01** (0.01)
Beliefs-Perceptions:		
Slightly effective		0.24** 0.23** (0.10) (0.10)
Moderately effective		0.25*** 0.24** (0.10) (0.10)
Very effective		0.33*** 0.32*** (0.09) (0.09)
Extremely effective		0.37*** 0.34*** (0.09) (0.09)
Infected in Area		0.04 0.01 (0.04) (0.04)
More Free Time		0.02 0.01 (0.02) (0.02)
More Family Time		0.05** 0.04* (0.02) (0.02)
Less Pollution		0.06*** 0.06**

						(0.02)	(0.02)
Less Noise						-0.03	-0.03
						(0.02)	(0.02)
Other						0.02	0.00
						(0.05)	(0.05)
Observations	955	955	955	955	955	955	955
R^2	0.033	0.049	0.046	0.040	0.064	0.099	0.138

Standard errors in parentheses

Linear probability models estimates for association between sociodemographic characteristics and behavior change.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.6: Factors Associated with Social Distancing More–Previously Not

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Income:							
Second quintile	0.10*	0.12**	0.09*	0.06	0.10*	0.08	0.08
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Third quintile	0.14***	0.16***	0.13**	0.09	0.15***	0.10**	0.09*
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Fourth quintile	0.15***	0.18***	0.14***	0.10*	0.17***	0.13***	0.12**
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Fifth quintile	0.22***	0.30***	0.22***	0.18***	0.26***	0.20***	0.20***
	(0.05)	(0.06)	(0.06)	(0.06)	(0.06)	(0.05)	(0.06)
Socio-Demographics:							
Black		-0.05					-0.01
		(0.05)					(0.05)
Other		-0.08					-0.15**
		(0.08)					(0.07)
Male		-0.19***					-0.14***
		(0.03)					(0.03)
Florida		-0.13***					-0.15***
		(0.05)					(0.04)
New York		-0.06					-0.04
		(0.04)					(0.04)
Texas		-0.12**					-0.11**
		(0.05)					(0.05)
56 or Older		0.05					0.04
		(0.03)					(0.04)
Health:							

Diabetes	-0.01 (0.05)	0.02 (0.04)
High Blood Pressure	0.01 (0.04)	-0.01 (0.04)
Heart Disease	-0.24*** (0.07)	-0.21*** (0.07)
Asthma	0.02 (0.05)	0.01 (0.05)
Allergies	0.03 (0.04)	-0.01 (0.04)
Other Condition	0.05 (0.05)	0.01 (0.05)
Housing:		
Semi-urban / residential	0.06 (0.04)	0.04 (0.04)
Countryside	0.02 (0.06)	-0.01 (0.05)
No Access to Open Air	-0.19*** (0.05)	-0.11** (0.05)
Work Arrangements-Losses:		
Stopped Working		0.14*** (0.05)
Began Tele-Working		0.14** (0.06)
Not Observed Working		0.22*** (0.05)
Other		0.19** (0.08)
Flag extreme lost income		0.03 (0.05)
Lost HH Inc./\$1,000 Orig		-0.00 (0.00)
Lost HH Inc./\$1,000 Adj.		0.04*** (0.01)
Beliefs-Perceptions:		
Slightly effective		0.16* (0.09)
Moderately effective		0.15* (0.08)
		0.11 (0.11)
		0.09 (0.10)

Very effective						0.33***	0.27***
						(0.08)	(0.10)
Extremely effective						0.39***	0.30***
						(0.08)	(0.10)
Infected in Area						-0.28***	-0.24***
						(0.07)	(0.07)
More Free Time						0.01	0.01
						(0.03)	(0.03)
More Family Time						0.08**	0.10***
						(0.03)	(0.03)
Less Pollution						0.17***	0.15***
						(0.04)	(0.04)
Less Noise						0.03	0.02
						(0.04)	(0.04)
Other						0.18**	0.15*
						(0.09)	(0.09)
Observations	918	918	918	918	918	918	918
R^2	0.019	0.071	0.030	0.038	0.057	0.135	0.206

Standard errors in parentheses

Linear probability models estimates for association between sociodemographic characteristics and behavior change.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.7: Factors Associated with Increased Washing Hands-Wearing Mask

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Income:							
Second quintile	0.07	0.07	0.06	0.05	0.07	0.06	0.05
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Third quintile	0.19***	0.21***	0.19***	0.17***	0.20***	0.17***	0.17***
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Fourth quintile	0.13***	0.15***	0.13**	0.10*	0.15***	0.13***	0.11**
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
Fifth quintile	0.20***	0.25***	0.22***	0.17***	0.24***	0.20***	0.19***
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.06)
Socio-Demographics:							
Black		0.09*					0.13***
		(0.05)					(0.05)
Other		0.08					0.06
		(0.07)					(0.07)

Male	-0.06*		-0.01
	(0.03)		(0.03)
Florida	0.02		-0.00
	(0.04)		(0.04)
New York	0.04		0.05
	(0.04)		(0.04)
Texas	-0.03		-0.03
	(0.04)		(0.04)
56 or Older	0.14***		0.13***
	(0.03)		(0.04)
Health:			
Diabetes	-0.07*		-0.06
	(0.04)		(0.04)
High Blood Pressure	0.09***		0.04
	(0.03)		(0.03)
Heart Disease	-0.05		-0.05
	(0.08)		(0.07)
Asthma	0.05		0.03
	(0.04)		(0.04)
Allergies	0.08**		0.08**
	(0.03)		(0.03)
Other Condition	0.08*		0.05
	(0.05)		(0.04)
Housing:			
Semi-urban / residential	-0.02		-0.06*
	(0.03)		(0.03)
Countryside	-0.04		-0.07
	(0.05)		(0.05)
No Access to Open Air	-0.12**		-0.07
	(0.05)		(0.05)
Work Arrangements-Losses:			
Stopped Working		0.19***	0.13**
		(0.05)	(0.05)
Began Tele-Working		0.13**	0.13**
		(0.06)	(0.05)
Not Observed Working		0.19***	0.08
		(0.05)	(0.05)
Other		0.24***	0.20***
		(0.07)	(0.06)

Flag extreme lost income						-0.02 (0.05)	-0.03 (0.05)
Lost HH Inc./\$1,000 Orig						-0.00** (0.00)	-0.00 (0.00)
Lost HH Inc./\$1,000 Adj.						0.02** (0.01)	0.02** (0.01)
Beliefs-Perceptions:							
Slightly effective						-0.03 (0.13)	-0.04 (0.13)
Moderately effective						0.10 (0.12)	0.08 (0.12)
Very effective						0.22* (0.11)	0.19 (0.12)
Extremely effective						0.30*** (0.11)	0.24** (0.12)
Infected in Area						-0.20*** (0.07)	-0.20*** (0.07)
More Free Time						-0.04 (0.03)	-0.04 (0.03)
More Family Time						-0.02 (0.03)	0.00 (0.03)
Less Pollution						0.09*** (0.03)	0.09*** (0.03)
Less Noise						0.05 (0.03)	0.04 (0.03)
Other						0.09 (0.07)	0.05 (0.07)
Observations	891	891	891	891	891	891	891
R^2	0.026	0.058	0.046	0.034	0.059	0.120	0.175

Standard errors in parentheses

Linear probability models estimates for association between sociodemographic characteristics and behavior change.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.8: Benefits of Pandemic by Income

Income Quintile	More Family Time			Less Pollution		
	Not Selected	Selected	Total	Not Selected	Selected	Total
First	0.64	0.36	168	0.64	0.36	168
Second	0.67	0.33	190	0.61	0.39	190
Third	0.59	0.41	210	0.56	0.44	210
Fourth	0.53	0.47	243	0.59	0.41	243
Fifth	0.43	0.57	196	0.61	0.39	196

Cross tabulation of US sample for two perceived benefits of the pandemic by income quintiles.

Assessing the effectiveness of alternative measures to slow the spread of COVID-19 in the United States

Michael Gapen,¹ Jonathan Millar,² Blerina Uruçi³ and Pooja Sriram⁴

Date submitted: 22 July 2020; Date accepted: 23 July 2020

To better understand the trade-offs at play as US states take measures to slow the spread of COVID-19, we investigate the effectiveness of alternative mitigation strategies using panel data estimates informed by epidemiological models. Our analysis evaluates public health outcomes using estimates of the effective reproduction number (R_t), which must drop below 1.0 to achieve sustained reductions in the infectious population. We fit outcomes for R_t on mitigation measures adopted by states, using daily data from early February through late June. Although all of the measures examined help reduce R_t , their effectiveness varies. Reductions in personal mobility on the scale achieved in April can reduce R_t by about a half and are especially effective when paired with stay-at-home orders. Alternatively, our estimates suggest that the virus could be brought under control using less-costly remediation measures. Those measures would likely involve more testing (at a rate of at least 1.75 million per day, if used in isolation), mask wearing requirements (which, in some specifications, are equivalent to testing 1.1mn persons per day), and restrictions on seated dining (which are effective when masks are not mandated). Finally, our estimates suggest that the US is nowhere near the point where “herd immunity” alone can bring infections under control.

1 Managing Director and Head of US Economics Research at Barclays BCI in New York.

2 Director and Senior US Economist at Barclays BCI in New York.

3 Director and Senior US Economist at Barclays BCI in New York.

4 Vice President and US Economist at Barclays BCI in New York.

Copyright: Michael Gapen, Jonathan Millar, Blerina Uruçi and Pooja Sriram

Introduction

With US economic activity only beginning to recover ground lost during the stay-at-home orders, experience over the past few months resoundingly confirms that measures taken by state and local governments and by the public to slow the spread of COVID-19 infections can have important implications for asset prices and economic outcomes. Several states are now tapping the brakes on plans to open up activity further while considering other measures to slow the rate of infection. Evidence regarding the effectiveness of various remediation options will be helpful to inform policymakers as they weigh likely trade-offs between virus mitigation and economic activity. Such evidence will also be helpful for economists as we consider the likely evolution of aggregate economic activity in the US.

In an effort to better understand the economic trade-offs at play as governments take measures taken to slow the spread of COVID-19, we use econometric tests informed by epidemiological models to investigate the effectiveness of alternative mitigation measures. Our analysis evaluates public health outcomes using estimates from a leading expert of the effective reproduction number (R_t). Sustained reductions in the new case count would require R_t to drop below 1.0. Using a panel data approach, we fit state-level outcomes for R_t on mitigation measures that have evolved in each state, using daily data from early February through late June. Our mitigation measures include reductions in personal mobility, formal stay-at-home orders, testing rates, facemask requirements, and reductions in seated dining at restaurants. Our panel regressions also control for other latent influences that likely confound underlying relationships, such as cross-state differences in the initial severity of the outbreak, population density, weather trends, and the exposed population (herd immunity).

Our research falls within recent threads that have looked at the effects of stay-at-home orders and other policy measures put in place by most states, for varying durations, starting in mid-March. One key area of focus has been the effects of these measures on personal mobility and economic activity more generally. Not surprisingly, these studies generally conclude that the stay-at-home orders coincided with general reductions in mobility (Goolsbee and Syverson (2020), Alexander and Karger (2020), Nguyen *et al.* (2020), Barrios *et al.* (2020), Maloney and Taskin (2020), and Chen *et al.* (2020)) and various measures of activity (Gupta *et al.* (2020, a), Gupta *et al.* (2020 b), Jiang *et al.* (2020), Coibion *et al.* (2020), and

Alexander and Karger (2020)).¹ A common finding has been that voluntary behavioral changes induced by the outbreak have been at least as important as, if not more important than, policy measures in terms of explaining reductions in mobility and activity.

Of course, evaluating effects on economic outcomes is only one part of the debate about the appropriate policy response to COVID-19, as the appropriate mix of mitigation measures would weigh the fallout on economic activity against benefits in terms of public health. This work falls more squarely under the latter consideration, which encompasses a second strand of research that examines how policy measures, and changes in mobility unrelated to policy, have affected the spread of the virus. This includes work by Dave, Friedson, Sabia and other co-authors, who, in a series of papers, have examined how stay-at-home orders have affected the measured rate of increase in COVID-19 infections.² Our approach builds upon these works by controlling for a much broader set of mitigating measures and by using benchmarks for contagion and structural controls that are informed by epidemiological theory. Another contribution of our study is that we analyse a measure of contagion that is not derived from measured case growth, which, among other things, allows us to sidestep complications that could arise from variations in testing rates across regions and through time.

To preview our results, we find that all of the mitigation measures help reduce R_t , though their effectiveness varies. In particular, reductions in personal mobility on the scale of those achieved in March and April can reduce the reproduction rate by about a half and are especially effective when combined with stay-at-home orders. However, we view this as more viable as an emergency measure to slow an out-of-control outbreak at its early stages, as stay-at-home measures have proven very costly to sustain, given the costs of foregone economic activity.

Our estimates suggest that combinations of other remediation measures, including a well-designed testing regime, could be used to bring the outbreak under control. According to our estimates, R_t could be reduced from its current level of about 1.0 (as of July 21) to 0.9 by boosting the testing rate from its average level of around

¹ See also Alfaro *et al.* (2020), who document that the anticipated severity of contagion is negatively related to stock market returns.

² See the four works by these authors cited in the appendix. Along similar lines, Deb *et al.* examined the effects of various mitigation measures on case counts and deaths, using cross-country data, while Fang *et al.* (2020) and Tian *et al.* (2020) have focused on the effects of policy measures in China.

780,000 tests per day to about 1.8 million tests per day. At this rate, the number of active infections would halve every 104 days. Boosting the testing rate even further to 4.4mn tests/day would halve the number of active infections every month. We also find evidence that rules mandating masks can reduce R_t by nearly 10%, which is roughly equivalent to testing another 1.1mn persons/day. Some evidence also suggests that restrictions on in-person dining help contain virus spread, but only in states where masks are not mandated.

Our estimates also provide some insight into the potential effectiveness and duration of a “herd immunity” strategy. According to our estimates, the US as a whole is nowhere near the point where a large exposed population, in isolation, could be relied upon to contain the spread of the virus. Given the cumulative count of roughly 3.8 million cases, our estimates indicate that the herd immunity effect has reduced R_t by only about 6%, which is well short of what would likely be needed to control infections without other remedial measures.


We begin our discussion with a brief review of epidemiology models, with an eye to identifying a basic measure of contagion and some benchmarks that can be used to assess the effectiveness of mitigation in terms of slowing the outbreak. We then discuss the data for our study in detail, including the merits of some alternative measures of the effective reproduction number and a description of the various explanatory variables that we use in our panel regressions. The remainder of the piece describe our panel regression estimates in detail, framing the discussion in terms of the effectiveness of alternative broad mitigation strategies.

Epidemic models: A basic primer

As a benchmark to reflect the speed of the spread of COVID-19, we use a basic measure of contagion that derives from epidemiology models. Although the array of models is much too complex to communicate in a brief summary, their basic building blocks are fairly straightforward. The workhorse model, called “SIR”, starts with an assumption that some population (N) can be divided between three states:

1. a susceptible population (S_t) that is vulnerable to the disease,
2. an infectious population (I_t) that can spread the disease, and
3. a “recovered” population ($N - S_t - I_t$) that, for simplicity, lumps together those who have been infected by the virus and gained immunity, others who have succumbed to the disease, and those who were never susceptible.

FIGURE 1
States of the basic SIR model



Source: Barclays Research

The model then tracks the aggregate dynamics as people flow, probabilistically, between the three states through time (Figure 1), where transitions are only possible from *susceptible* to *infectious*, and from *infectious* to *recovered*. The key equation of the model determines how the number of infectious people evolves over time:

$$1. \quad \Delta I_{t+1} = \beta \left(\frac{S_t}{N} \right) I_t - \gamma I_t$$

The two terms of this equation represent inflows and outflows from the infectious state. The number of people flowing from susceptible to infectious in each period is the product of the number already infected I_t , the number of people β with whom each of these persons comes into contact, and the share of these contacted people who are susceptible to the virus (S_t/N). The parameter $\beta > 0$ is known as the *contact rate*, which could vary through time. The second term represents the number of transitions out of the infectious state, which is the product of the infectious population I_t and the share of infected people who recover each period $\gamma > 0$, known as the *recovery removal rate* or simply the *recovery rate*. This constant rate implies that the average infected person will remain contagious for $1/\gamma$ periods.³

³ Although the equation above is deterministic, one can allow for some randomness by regarding the first and second terms as random draws from discrete probability distributions. Technically, the first term can be thought of as a draw from a Poisson distribution with expected value $\beta(S_t/N)I_t$, and the second term as a random draw from a poisson distribution with expected value γI_t .

For reasons that will become apparent, it is helpful to restate the equation as:

$$2. \quad \frac{I_{t+1}}{I_t} = 1 + \gamma(R_t - 1),$$

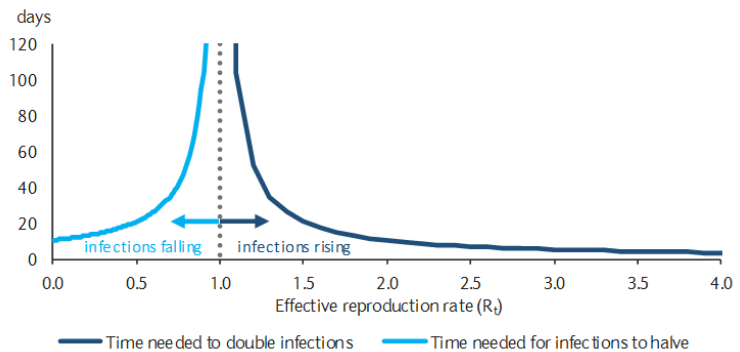
so that the growth rate in the infected group is $\gamma(R_t - 1)$. In this equation, we have made use of the definitions of the *effective reproduction rate*:

$$R_t = R_0 \left(\frac{S_t}{N} \right),$$

This definition also makes use of the *basic reproduction rate* $R_0 \equiv \beta/\gamma$, which is the number of people whom we would expect an infectious person to expose to the disease, during the entire duration of their contagion. Intuitively, this is calculated by multiplying the number of contacts per period β by the average duration of contagion $1/\gamma$. This rate is important in the model because it provides information about the inherent explosiveness of an outbreak at its initial stages, when nearly the entire population is likely to be susceptible ($S_t \cong N$). We would expect the infectious population to grow over time provided that $R_0 > 1$, so that newly infectious people are expected to generate more than enough additional infections to compensate for their eventual departure from the infected group. Conversely, if $R_0 < 1$, contagious people would not be generating enough additional infections to replace themselves, so that the number of infections would be expected to diminish over time. If $R_0 = 1$, the infectious population would be flat.

FIGURE 2

R_t determines whether the infectious population is increasing or diminishing and the magnitude of such changes



Note: Assumes average recovery period of 15 days. Source: Barclays Research

The effective reproduction rate R_t also represents the number of people a typical infectious person would be expected to infect during their contagious period, but for a more general situation where some of the population has existing immunity. It is the product of the basic reproduction rate R_0 , which is the relevant rate when there is no immunity, and the share of the exposed persons who would typically remain susceptible to the disease (S_t/N).

Returning again to equation (2), the rate of change in the infected population in any given period will depend on the share of infectious people expected to recover γ and the discrepancy $R_t - 1$. Much as with the basic reproduction rate R_0 , this discrepancy represents the number of people whom we would expect to step in to replace infected people as they recover, or “net replacement.” Hence, the sign of this rate of change will depend on whether the net replacement is positive or negative, while the magnitude of the growth rate will depend on the magnitude of net replacement and the recovery rate. For COVID-19, most estimates currently place the average duration of infectiousness at about 15 days, which would correspond to a removal/recovery rate of about $1/15$, or about 6.7% of the infected population per day. Given this number, an effective reproduction rate of $R_t = 2$ would imply a daily growth rate in infections of about 6.7%, which would imply a doubling of infections roughly every 10.4 days, while a value of $R_t = 1.1$ would imply a daily rate of about 0.67%, or a doubling roughly every 104 days.⁴ Working in the other direction, an effective reproduction rate $R_t = 0.9$ would halve the existing number of infections roughly every 104 days, and $R_t = 0.5$ would halve them every 10.4 days.

Policy options to control COVID-19

Models suggest alternative paths to slowing the virus


Given this construction, the model provides clear prescriptions about how to bring a virus outbreak under control, which amounts to pushing R_t below one. One path is to reduce the contact rate β , which, in turn, would reduce the basic reproduction rate. This option would include measures that limit close human interactions, such as through stay-at-home orders. Another approach is to limit the spread of the virus when there are interactions, such as by requiring masks. Alternatively, one could take measures to boost the recovery rate γ through better medical treatments or

⁴ Given the daily rate of change g_t , we calculate the doubling time as the quotient $\ln(2)/g_t$, and the halving time as $-\ln(2)/g_t$. Many readers will recognize this as the “rule of 72”.

other measures, which is equivalent to shortening infection duration.⁵The final path is to boost immunity, which could occur proactively through the creation of a vaccine, or, reactively, by simply allowing the virus to run its course in order to achieve herd immunity. Indeed, one of the key outcomes of the model is that even the most intense outbreak must *eventually* slow by attrition, as the diminishing share of the population that remains susceptible to the virus pushes down R_t .

These approaches seemingly vary both in effectiveness and in costliness, involving inherent trade-offs. For this reason, formulating the appropriate policy response falls squarely within the realm of economics, requiring a careful weighing of costs – such as foregone economic activity, inconvenience, and diminishing the wellbeing that comes from more limited human interaction – against benefits in terms of public health and limiting loss of life. In the remainder of this study, we mainly investigate the effectiveness part of this trade-off, exploring whether there are different mixes of policy mitigation measures that achieve similar outcomes in terms of public health.

FIGURE 3
Measures of the effective reproduction rate from different sources tell a similar story



Note: The rt.live estimate is formed by population weighting across the 50 states and the District of Columbia; Source: covid19-projections.com, rt.live.com.

R_t as a summary measure of health outcomes

For our purposes, we view R_t as an ideal measure to summarize the effect of various mitigation measures in terms of public health, since it succinctly captures the severity of the outbreak in a given area and the likely evolution of infections and deaths in the coming days. Conceptually, this measure is closely related to the growth rate in new infections, which tends to be correlated with the death count

⁵ To see this, note that our formulas imply that $\gamma(R_t - 1) = \beta \left(\frac{S_t}{N} \right) - \gamma$.

after an appropriate lag. R_t is also ideal because it helps cut through much day-to-day noise in observed variables (such as daily case counts), potentially providing a higher-quality signal about the likely evolution of the virus. That said, obtaining reliable estimates requires a great deal of expertise, in part because observed measures of new cases, recoveries, and deaths do not directly correspond to model concepts, such as susceptible and infectious populations.

Estimates using different approaches tell a consistent story

In practice, there is a variety of methods that can be used to estimate R_t , usually with an eye to obtaining a “best fit” for observed variables under the assumption that these outcomes are generated by the model. One approach estimates R_t using daily counts of new COVID-19 cases, which is the approach taken by the website *rt.live*. To be sure, these estimates are only as good as the case data that underlie them and could be misleading if testing occurs at low rates, varies in magnitude over time, or is applied only selectively to some subset of the overall population (such as persons with acute symptoms). Partly for this reason, we place most of our attention on estimates of R_t from *covid19-projections.com*, which uses a methodology that sidesteps these potential problems by applying a machine learning algorithm that generates estimates that best explain the daily evolution of COVID-19 death counts.⁶ We prefer these estimates, in part, because death counts are plausibly less susceptible to measurement bias, partly because they have outperformed alternative approaches in terms of predicting future caseloads, and partly because they are available for a longer history.

Reassuringly, estimates using these two approaches tell a very similar story. Figure 3 compares the time path of R_t estimates for the US as a whole from the two sources. As one can see, both approaches place the effective reproduction number in the vicinity of 2.3 during February and the first half of March, prior to mitigation measures implemented by many states. The rate began to plunge rapidly in mid-March, when the hardest-hit states (including California and New York) began to enact stay-at-home orders, plunging to around 0.9 by mid-April, with nearly all states adopting similar lock down measures. After lingering somewhat below 1.0 for about a month, the rate began to climb during May and much of June, when most states took measures to relax restrictions. By late June, the rate briefly settled at

⁶ The model that underlies these estimates is a bit more detailed than the basic SIR model described earlier, including a separate accounting of recoveries and deaths.

nearly 1.1, and it has gradually descended toward 1.0 over the course of during July with increased testing and a re-intensification of remedial measures, including moves by some states (including Nevada and Texas) to mandate masks and by others (including Arizona, California, Florida, Louisiana, and Texas) to restore restrictions on bars, restaurant dining, or both.

This effective reproduction rate would need to drop below 1.0 to put the overall number of infectious cases on a sustained downward trajectory. Given variations in R_t across states – with state estimates currently ranging from about 0.9 in New Hampshire to nearly 1.10 in Hawaii, according to *COVID19-projections.com* – the national figure would likely need to drop to at least 0.9 to eliminate hotspots in some parts of the country, which would be consistent with the overall number of infectious cases halving roughly every 104 days.

Estimating the effectiveness of mitigation measures

With this brief summary of measures of contagion from epidemiology in mind, we turn to estimating the effectiveness of various measures to control the spread of COVID-19. For this purpose, we use a panel analysis of state-level data, with the effective reproduction number, R_t , as the dependent variable. Our dataset is assembled from state-level daily data between February 5 and June 20, 2020, using a regression specification that is informed by the basic epidemiological model discussed above. Specifically, we estimate an equation of the form,

$$\ln(R_{i,t}) = \alpha_i + \delta_t + \beta_1 MEI_{i,t} + \beta_2 TST_{i,t} + \beta_3 SHO_{i,t} + \beta_4 MASK_{i,t} + \beta_5 DINE_{i,t} + \beta_6 X_{i,t} + \varepsilon_{i,t}$$

where $i=1,\dots,N$ is the number of US states (including the District of Columbia), and $t=1,\dots,T$ is the number of days in our sample period. The dependent variable $R_{i,t}$ is the effective reproduction number in state i for day t , which we convert to logs in order to impose the positivity constraint implied by the model.⁷ Our mitigation controls include $MEI_{i,t}$, the mobility and engagement for each state and period, which we derive from the Dallas Fed Mobility and Engagement Index; $TST_{i,t}$, the moving sum of the number of tests for COVID-19 in the state over a six-day period, which we express in percentage points of the state population; $SHO_{i,t}$, a dummy variable that

⁷ We estimate our panel regression using two alternative estimates for $R_{i,t}$ across states. The first comes from *covid19-projections.com*; the results of this specification are presented in Figure 9. Estimation results using a measure of the effective reproduction number from *rt.live* yield similar results, but are not presented here for reasons of space. We also obtain qualitatively similar results when we substitute $R_{i,t}$ with the growth rate in daily cases, but the standard errors tend to be much larger in magnitude, and the number of usable observations is reduced by about one-fifth.

takes on the value of 1 during days in which a state has a stay-at-home order in place and 0 otherwise; $MASK_{i,t}$, a dummy variable that takes on the value of 1 on days in which a state has a requirement to wear a mask in public and in enclosed spaces and 0 otherwise; and $DINE_{i,t}$, the year-on-year percentage change in seated diners in state i on day t . We also include a set of observed exogenous controls $X_{i,t}$ that vary by time and state, which include heating and cooling degree days and the cumulative number of confirmed positive cases in the state, expressed as a percentage of population.⁸

To these controls, we add a full set of state fixed effects α_i that control for other unobserved influences on the effective reproduction number that do not vary over time, which would include the initial severity of the outbreak, population density, demographics, and other influences. Finally, we include a full set of time fixed effects δ_t , which will control for any unobserved aggregate factors that affect all states in the same way and that vary over time, such as common mitigation measures taken by the federal government and government agencies.


Our regression specification was carefully calibrated for ease of interpretation. Since $R_{i,t}$ is expressed as a log, the coefficients on each of the included variables can be roughly interpreted as the percentage the effective reproduction rate changes when the control is set to one. For example, β_2 is roughly the percentage that $R_{i,t}$ will change if 1% of the population is being tested every six days. Moreover, we express each of the mitigation variables such that they equal zero in the absence of any mitigation. For example, the mobility index is normalized so that it is zero for “pre-COVID” observations, dining is zero when it is unchanged from the prior year, and so forth. Likewise, the exogenous controls can also be interpreted as being zero on a typical pre-COVID day. Given this construction, we can roughly interpret fitted values of $R_{i,t}$ for each state, with all the included mitigation and exogenous controls set to zero, as the effective reproduction number that would prevail in the state with no state-level mitigation measures. For example, fitted values for any given state using just the state fixed effect, plus the time effects from the initial day in the sample, can plausibly be interpreted as a proxy of the state’s basic reproduction rate $R_{i,0}$. Indeed, we exploit this interpretation to compare differences in basic reproduction rates across states.

⁸ In principle, the coefficient on the cumulative exposures variable should be about 1.0 if our measure is an accurate estimate of the exposed population share. We allow some flexibility for the estimate to deviate from this theoretical value, in part because the cumulative count of the infected population is likely under-measured. Later, we show how we can use this property to form an estimate of the degree of under-measurement.

Data and identification strategy


We use the Dallas Fed Mobility and Engagement Index (MEI) as a broad proxy for the degree of social distancing. Measured values summarize the information in seven different variables based on geolocation data collected from mobile devices.⁹ These are (1) the fraction of devices leaving home in a day, (2) devices away from home for three to six hours at a fixed location, (3) devices away from home longer than six hours at a fixed location, (4) devices taking trips longer than ten miles, (5) devices taking trips less than 1.2 miles, (6) the adjusted average of daytime hours spent at home and (7) the average time spent at locations far from home. These are combined using principal components analysis, with the first fitted component used to derive a county-level MEI. These MEIs are then aggregated to form measures for metro areas, states, and the nation as a whole. The Dallas Fed scales the national MEI so that January and February average zero and that the minimum in the week ended April 11 is -100. With this weighting, the MEI captures the degree of mobility during the pandemic, with values of zero indicating “normal” mobility and values of -100 the “maximal” reduction in mobility that was achieved during the height of the lockdowns that were implemented in March and April 2020. We take the 7-day moving average of the reported measures and rescale this value by dividing it by 100, so that the estimated coefficient measures the percentage change in R_t that can be achieved by reducing mobility during a lockdown, given the experience from that episode. Our rescaled Dallas Fed Mobility and Engagement Index can be found in Figure 4.

FIGURE 4 Dallas Fed Mobility and Engagement Index



Note: 7dma for the US and state-level MEI divided by 100. A value of -1 on the US index is equal to mobility observed during the week of April 11, 2020, when state-wide shutdowns were in effect. For regions with values of less than -1, mobility fell more than the national average. Source: FRB Dallas, Barclays Research

FIGURE 5 Change in cumulative COVID tests in the US



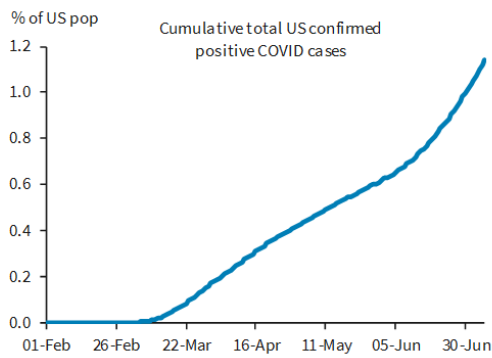
Source: Census Bureau, The COVID Tracking Project, Barclays Research

⁹ Additional information on the Dallas Fed Mobility and Engagement Index can be found at <https://www.dallasfed.org/research/mei.aspx>

Our state-level testing measures are formed using data reported by The COVID Tracking Project.¹⁰ Our preference would be to include only polymerase chain reaction (PCR) tests, which are used to detect active infections by detecting the genetic information of the virus, rather than serologic or antibody tests, which determine who has previously had the virus and is likely immune. During the early stages of the pandemic, virtually all reported tests were PCR tests, but over time this test count has been mixed by most states with the results of antibody tests, somewhat clouding the interpretation of this measure. If anything, using a combined measure of tests performed would likely reduce the likelihood of finding statistical significance between tests performed and the effective reproduction rate. Hence, if our analysis finds that testing has a statistically significant effect on R_t , then we can have high confidence in the result. To form our estimate of the testing rate, we take the cumulative number of tests at time t , deduct the cumulative total from six days earlier, divide by the state population (as reported by the US Census Bureau), and multiply by 100. Results of this procedure for the US as a whole are presented in Figure 5.

FIGURE 6


Cumulative confirmed COVID cases



Source: Census Bureau, The COVID Tracking Project, Barclays Research

FIGURE 7

Index of dining traffic



Note: 7dma in the y/y change in seated dining from *OpenTable*, divided by -100. A value of 1 is equivalent to a restriction on seated dining while values near 0 suggest a return to pre-COVID dining. Source: Haver Analytics, Barclays Research

We also use a cumulative confirmed cases variable to control for the share of the population that has been exposed to the virus. According to epidemiology models, this should have a negative effect on R_t , reflecting the degree of herd immunity. To

¹⁰ See <https://covidtracking.com/>

proxy for this share, we compute the cumulative number of confirmed COVID-19 cases each day as a share of the state's population, expressed as a percentage. An example for the US as a whole is shown in Figure 6. Data on cumulative confirmed positive cases are taken from The COVID Tracking Project and data on state populations are from the Census Bureau. Given this, the estimated coefficient on this variable will represent the percentage reduction in R_t for each percentage point of the state's population that has previously been infected by the virus.¹¹

The dummy variable for stay-at-home orders, which takes on the value of 1.0 for each day a state had a stay-at-home order for the entire state in effect, is formed using information on emergency orders from the National Governors Association and *The New York Times*.¹² Helpfully, there is a lot of variation through time and across states to identify this effect. At the peak of the initial surge in coronavirus cases, 42 US states had imposed stay-at-home orders and several others had imposed partial orders. The first was imposed by California on March 19, while the last state to formally end such orders was New Hampshire on June 15. The duration of these orders varied substantially across states, with the longest in New Hampshire and New Jersey (81 days), and the shortest in Mississippi (25 days). At the time that states moved to lift stay-at-home orders, we set the dummy equal to 0, even in cases where states engaged in a phased re-opening. In doing so, we control only for the most stringent period of lockdowns.

To form our dummy for mask requirements, we used our judgment based on an examination of state-wide orders issued by state governments and the District of Columbia. As of end-June, 24 entities had a state-wide mask order in effect. The remaining states had no formal requirement, though some issued mask recommendations. On April 24, the Center for Disease Control (CDC) issued a non-binding mask guidance for the US as whole, the effects of which should be reflected in our time fixed effects from April 24 onward. This federal guidance was within a few weeks of the earliest mask requirements issued by states (New Jersey, April 8). Hence, the coefficient on this variable can be reasonably interpreted as the percentage reduction in R_t from issuing a formal state-wide requirement to wear a face mask or appropriate covering, as opposed to a simple recommendation.

¹¹ Technically, the effect would be calculated as $1 - \exp(\beta_2 TST_{it})$

¹² Current information on the status of COVID-19 actions taken by states and territories from the National Governors Association can be found at <https://www.nga.org/coronavirus/>. Data on stay-at-home orders from The New York Times can be found [here](#) and [here](#).

We proxy for restrictions and/or the willingness of households to attend dining facilities using daily state-level data from *OpenTable*.¹³ This variable is included, in addition to mobility and stay-at-home orders, because dine-in restaurants have been subject to closures and strict operating limits even after shutdown orders were relaxed. Bars, in particular, have been the subject of increased scrutiny in many states in the South and West and have been singled out as potential high-risk settings for virus transmission. This reflects the unique features of in-person dining and drinking. While some activities such as general shopping can be done while wearing personal protective equipment like masks and gloves, such preventative behavior is more difficult to implement while dining. *OpenTable* provides information on the year-on-year percent change in seated diners in restaurants within the *OpenTable* network, which includes those derived from online reservations, phone reservations, and walk-ins. However, only states with metropolitan areas with 50 or more restaurants are included in the sample, so data availability is limited and likely reflects in-person dining in more populous regions of the country. That said, these data are derived from a large sample of approximately 20,000 restaurants across the country. To form our measure, we take the seven-day moving average in the year-on-year percentage change in each state, then re-scale by dividing by -100. Given this scaling, a value of 1 indicates that seated dining is down 100% relative to year-earlier levels, a value of 0 indicates that dining was at year-earlier levels, while negative values indicate an increase in dining. The coefficient on this variable indicates the percentage change in R_t from completely closing restaurants to seated dining in the state, a situation that should be well identified in our sample with seated diners down 100% in many states for weeks during the lockdowns. Our re-scaled measure of *OpenTable* state-level data can be found in Figure 7.

In principle, weather should affect R_t because unusually hot or cold weather tends to drive people indoors, where transmission risks are higher. To measure unusual weather, we use heating and cooling degree days from NOAA, which are estimates of energy demand over a period of time related to the sum of the daily heating and cooling degree day values, or the degree day accumulation. Each of the two series is based on the deviation of the mean temperature during the period from 65 degrees. Since data on heating and cooling degree days are available only weekly,

¹³ For information on *OpenTable* data, see <https://www.opentable.com/state-of-industry>

we divide the cumulative heating and cooling requirement readings by seven to convert to average daily values. We then re-scale each series to take on a value of 1.0 when the temperature is equal to the hottest or coldest day on record, which we define as 134 degrees Fahrenheit (hottest) and -70 degrees (coldest) based on historical US weather data. Hence, the estimated coefficient on each measure is equal to the percentage change in R_t if the state incurred the hottest or coldest degree day on record. If the temperature is a “normal” 65 degrees, the value is zero.

Effectiveness of measures to reduce community spread

Results of the panel estimations are shown in Figure 9. In total, we present nine different model specifications, each illustrating the effectiveness of various mitigation policies for reducing R_t . Our baseline specification estimates the effect of mobility and testing on reducing community spread. Specifications 3 and 4 examine the relationship between mobility and the use of stay-at-home orders, while specifications 5 and 6 examine state-wide mask requirements. Finally, specifications 7-9 include restrictions on dining and the interaction of dining restrictions and requirements to wear a mask.


For each estimation, we report a constant, which is normalized to represent the average of the (logged) fixed and time effects for all observations in our sample. As suggested earlier, we interpret these values as indicating the average effective reproduction number across states assuming no (state) mitigation effort. For example, the estimated value of the constant in our baseline specification is 0.4724, which, after converting from logs by taking the exponential, yields $R_t = 1.6$. This suggests that the average unmitigated reproduction number across all states during our sample period is consistent with a daily growth rate in the infectious population of about 4% and a doubling time of roughly 17 days. Although this seems rapid, it is probably somewhat lower than the basic reproduction number R_0 at the onset of the outbreak, prior to public awareness of community spread of the virus, which likely had some effect on behaviour and R_t over time. Across the nine models, our estimate of this “unmitigated” average R_t across states ranges from 1.2 to 2.1.

To get a sense of how much the basic reproduction number R_0 varies across states, we use our baseline specification to form estimates, using the state fixed effects and the time effect for the first day of our sample. These estimates are shown in Figure 8, with states grouped by Census region and division, for comparison purposes. In general, we find that our rough estimates of R_0 are lower for less populated, rural

states and larger for states with denser urban populations. Indeed, the five states with the lowest estimates were Vermont (1.1), Montana (1.1), Hawaii (1.1), Alaska (1.1), and Idaho (1.2), while the five states with the largest numbers were New York (2.0), Illinois (2.0), Maryland (2.0), Minnesota (2.1) and New Jersey (2.2).

With state-level reproduction numbers in mind, a natural goal for policymakers could be to choose combinations of mitigation measures that push R_t below 1.0 for an extended period, which should reduce the growth rate of new cases (and deaths) over time. Ultimately, if R_t remains below 1.0 for long enough, on average, across states, the virus could be extinguished – a process that would occur more rapidly if the value were as far below 1.0 as possible. The reality of the situation, according to our estimates, is that levels of mitigation needed to achieve this goal will differ across states, depending on the various factors that influence their basic reproduction rate.

FIGURE 8
Initial state-level R_t estimates under the baseline model specification



Source: Barclays Research

Reducing mobility versus stay-at-home orders

In our baseline specification, we test the significance of mobility as an explanatory variable, excluding our stay-at-home order dummy. Model 3 tests the reverse, including stay-at-home orders but excluding mobility. Model 4 includes both variables. We include these as separate tests because we think that each measure may be capturing effects that overlap in scope. On one hand, imposing stay-at-home orders and locking down parts of the economy most “at risk” for community spread will almost surely coincide with reductions in mobility, as people will likely become less inclined to move about in such circumstances. On the other hand, it is possible to

imagine situations where people might remain mobile but do not congregate in large numbers at high-risk locations such as stores, offices, and outdoor locations. We also think it is plausible that lifting stay-at-home orders may not lead to a full restoration of mobility if individuals lack confidence that they can move around safely. Hence, the mobility variable could capture information about individual behavior and confidence that is independent from stay-at-home orders. We also believe the converse, that stay-at-home episodes might provide independent information about the effectiveness of prohibiting perceived risky face-to-face contact within some sectors of the economy.

Our suspicion that these variables capture different aspects of mitigation seems to be borne out in the data. Indeed, mobility plummeted during the lockdown even in states that did not impose lockdowns, such as Iowa and Arkansas (Figure 10). Moreover, even though stay-at-home orders generally coincided with rapid declines in mobility in states that imposed them, mobility has remained lower than usual in states that moved aggressively to relax stay-at-home restrictions, such as Texas and Florida (Figure 11). One key lesson from these two examples appears to be that stay-at-home orders are partly a reflection of public fears of infection, which are a key source of social distancing in their own right. This suggests that lockdowns are not necessary to reduce mobility and that relaxing lockdowns will not be sufficient to restore it fully unless underlying worries about contagion are addressed.

Our estimation results confirm that both measures can play distinct roles in mitigating the spread of the virus. In the baseline specification, mobility is statistically significant at the 1% level when included by itself, as are stay-at-home orders in isolation (model 3). Including both variables, as in model 4, does not significantly alter their coefficients and statistical significance. In all estimations in which mobility is included, reductions in mobility appear to have an especially high degree of leverage on the effective reproduction rate. In the baseline specification, a reduction in mobility by the maximum extent observed during the economic lockdowns reduces a state's R_t nearly 50%. Hence, if it were possible to reduce mobility severely without stay-at-home orders, this alone would be enough to push the effective reproduction rate below 1 for any state with R_t below 2.0. According to our estimates from Figure 9, all but five states fell into this category. Ostensibly, this provides much more leverage than stay-at-home orders in isolation. According to model 3, stay-at-home orders accompanied by no reduction in mobility would reduce R_t by about 10%.

FIGURE 9

Estimation results on the effectiveness of various measures taken to slow the growth of COVID-19

Independent variables	COVID-19 Reproduction number: log(R _t estimate)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Base model	Base model, no time fixed effect	Stay at home orders	Stay at home orders & mobility	Mask requirement	Stay at home orders & mask requirement	Dining restrictions	Mask requirement & dining restrictions	Dining restriction in states with no mask requirement
Constant	0.4724*** (0.0256)	0.3901*** (0.0150)	0.1369*** (0.0108)	0.4467*** (0.0246)	0.4239*** (0.0232)	0.4062*** (0.0229)	0.7511*** (0.0441)	0.6571*** (0.0385)	0.6860*** (0.0395)
FRB Dallas ME _{it}	0.6816*** (0.0442)	0.5142*** (0.0137)		0.5940*** (0.0430)	0.5747*** (0.0391)	0.5101*** -0.0393	0.8933*** (0.0460)	0.7894*** (0.0386)	0.7877*** (0.0382)
Testing rate _{it}	-0.0631*** (0.0113)	-0.1591*** (0.0108)	-0.0843*** (0.0107)	-0.0630*** (0.0110)	-0.0518*** (0.0111)	-0.0525*** (0.0109)	-0.1055*** (0.0116)	-0.1011*** (0.0113)	-0.1002*** (0.0113)
Heating degree days _{it}	0.1652** (0.0741)	0.5843*** (0.0800)	0.1993* (0.0720)	0.1931*** (0.0732)	0.1918** (0.0719)	0.2036*** (0.0713)	-0.0037 (0.0795)	-0.0155 (0.0767)	-0.0152 (0.0767)
Cooling degree days _{it}	0.1851* (0.1074)	0.3614*** (0.0863)	0.2243** (0.1165)	0.1284 (0.1116)	0.1454 (0.1081)	0.1015 (0.1123)	-0.0637 (0.1211)	-0.2138* (0.1215)	-0.2136* (0.1217)
Cumulative exposed _{it}	-0.0533*** (0.0094)	-0.1511*** (0.0105)	-0.0887*** (0.0103)	-0.0549*** (0.0095)	-0.0591*** (0.0086)	-0.0600*** (0.0088)	-0.0248*** (0.0094)	-0.0417*** (0.0095)	-0.0401*** (0.0094)
Stay at home order _{it}			-0.1000*** (0.0087)	-0.0635*** (0.0076)		-0.0524*** (0.0076)			
Mask requirement _{it}					-0.0162*** (0.0086)	-0.0987*** (0.0085)		-0.1013*** (0.0102)	-0.1602*** (0.0362)
Dine restriction _{it}							-0.0964** (0.0455)	-0.0230 (0.0414)	
Dine restriction*(1-Mask) _{it}									-0.0631* (0.0390)
R-square	0.8243	0.7962	0.8156	0.8273	0.8309	0.8329	0.8606	0.8673	0.8675
Observations	5077	5077	5077	5077	5077	5077	3580	3580	3580
Cross-sections included	51	51	51	51	51	51	37	37	37
Sample period	Feb 5-June 20, 2020	Feb 5-June 20, 2020	Feb 5-June 20, 2020	Feb 5-June 20, 2020	Feb 5-June 20, 2020	Feb 5-June 20, 2020	Feb 26-June 20, 2020	Feb 26-June 20, 2020	Feb 26-June 20, 2021

Note: Cross-section SUR (PCSE) standard errors & covariance. Source: Barclays Research

Covid Economics 40, 30 July 2020: 46-75

Although our regression models provide evidence in support of mobility and economic lockdowns as identifying separate effects, we also suspect that the idea of achieving a reduction in mobility on the scale seen during the lockdowns without stay-at-home orders may not pass the smell test for many readers. Even though mobility declined by similar amounts in states that did and did not impose lockdowns, it seems reasonable to think that lockdowns in most of the states played a role in reducing mobility more generally. For those who prefer that interpretation, it seems reasonable to regard the overall effect of stay-at-home orders as the sum of the effects from reduced mobility and the stay-at-home dummy variable. When we add the two effects, as in model 4, the maximum reduction in R_t from combining the two measures is about 48%, similar to the effect from a maximum reduction in mobility in our baseline model. Examination of the nationwide R_t number suggests that this combined magnitude is reasonable. Prior to the drop-off in mobility and imposition of stay-at-home orders, R_t for the US as a whole was estimated at about 2.3 (Figure 3). On the eve of the first wave of measures to re-open state economies, the number had fallen to about 0.9, a reduction of slightly more than 60%.


Altogether, we draw the following conclusions about mobility and economic lockdowns. First, both lockdowns and mobility have statistically significant explanatory power in reducing community spread, as evidenced by the 1% level of significance across all model estimations. Their joint power in reducing R_t is quite strong. Second, our mobility proxy is likely capturing both voluntary and involuntary social distancing, with the latter occurring due to government-imposed lockdowns and the former through voluntary changes in behavior as people respond to perceived infection risk. While government lockdowns will likely come with reduced mobility and the removal of official restrictions will likely lead to increased mobility, the willingness of households to return to pre-COVID patterns of movement – going to work, restaurants, travel, etc. – will also be dependent on the degree that households believe the virus is under control.

Third, we can say, unequivocally, that stay-at-home orders combined with voluntary reductions in mobility can have very powerful effects on reducing the effective reproduction rate for COVID-19. Although our results also suggest that maximum reductions in mobility, taken in isolation, have a much larger effect than stay-at-home orders, we cannot state this with certainty because our analysis may not fully disentangle the scope of the two effects. That said, our estimates seem broadly

consistent with findings from studies that are able to apply careful statistical treatments to disentangle effects on mobility from these two sources more carefully (most notably, Goolsbee and Syverson (2020)).¹⁴ This gives us confidence to state a fourth conclusion, that states may be able to reduce community spread, without resorting to broad and economically costly lockdowns, by issuing social distancing recommendations alongside more targeted restrictions on particularly high-risk activities.

FIGURE 10

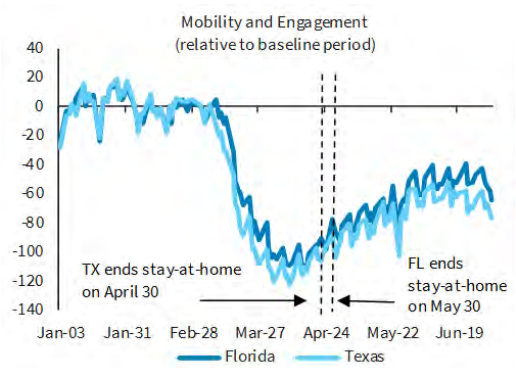
Mobility fell sharply in March and April in states that did not impose state-wide stay-at-home orders



Source: FRB Dallas, NGA, Barclays Research

FIGURE 11

Mobility has recently fallen in Florida and Texas following the sharp rise in new COVID-19 cases



Source: FRB Dallas, NGA, Barclays Research

Increased testing is effective in reducing community spread


As noted earlier, reducing R_t to a value of 0.9 means it takes about 104 days to cut the number of infectious cases in half. If the US were experiencing 10,000 new cases per day, this would mean it would take roughly 416 days for new daily cases to fall below 1000 (where daily new cases are across France, Germany, Spain, and Italy, for example). At the current seven-day average of about 66,000 cases per day, it would take about 628 days to push the daily new case count below 1000 cases per day. Although lockdowns can surely achieve this goal, the economic fallout is very costly, and we do not believe they are a viable solution on their own to address the pandemic except, in cases where governments react extremely quickly at the first

¹⁴ Using detailed data from individual businesses that draw customers from diverse commuting zones that cut across municipalities and counties with varying COVID-19 responses, Goolsbee and Syverson (2020) find that reductions in consumer traffic from voluntary distancing were seven to eight times larger than those from stay-at-home requirements. This difference in magnitudes is broadly consistent with those between the coefficients we estimate on stay-at-home orders and mobility in specifications (3) and (6).

signs of an outbreak. Our results support the conventional wisdom that lockdowns are a very effective temporary tool to buy time while more sustainable remedial measures are developed.

We turn next to COVID-19 testing to see whether the evidence supports the idea that increased testing capacity can reduce R_t by identifying infected persons and those with whom they have come into contact. If this is the case, an adequate testing regime could effectively transition policy from a macro-quarantine framework under stay-at-home orders to a micro-quarantine system whereby individuals are tested, then quarantined with contact tracing if they are positive. Our regression results support the idea that increased testing can indeed reduce community spread. The estimated coefficient on our testing variable represents the percentage change in the effective reproduction rate from testing 1% of the state's population every six days. Nationwide, this would be equivalent to about 550,000 tests per day, based on the current US population of about 330mn. In our baseline estimation, a testing regime of this magnitude would reduce R_t about 6%. As more explanatory variables are added, the coefficient rises slightly, implying a reduction in the effective reproduction rate of about 9.0%. Across all model specifications, the testing variable is statistically significant at the 1.0% level, indicating that this effect is very robust.

FIGURE 12
Fewer than half of the states appear to be testing enough to prevent the infectious COVID-19 population from growing



Note: Required testing rate is rate needed to bring reported R_t to 0.90 or below, using estimates and testing data reported on July 21, 2020 and results of our baseline regression specification. Source: COVID19-projections.com, The COVID Tracking Project, Barclays Research

The US still is not testing enough

The question then becomes: How much testing is needed? Prior to economic lockdowns, most estimates from epidemiological models showed that the national effective reproduction number was about 2.0. Relative to this benchmark, our results indicate that, absent all other remedial measures, the US would need to be testing about 12.7% of the population every six days – which amounts to nearly 7.0mn tests per day – to push R_t to 0.9.¹⁵ To reduce R_t to 0.5, testing capacity would need to be about 22% of the population every six days, or about 12.1mn tests per day. According to the data from The COVID Tracking Project through July 20, the US was testing at a seven-day average of around 780,000 per day. This, alongside other remedial steps to reduce mobility (including past stay-at-home orders) and changes in household behavior, has helped push down the *COVID19-projections.com* estimate of R_t to nearly 1.0 as of July 21, 2020. From this lower starting point, the US, as a whole, would need to boost total testing capacity to about 3.2% of the population every six days, or about 1.8mn tests per day, to push R_t to 0.9. At the moment, fewer than half of the states are testing enough to simply keep the infectious population from growing (Figure 12). Given current effective reproduction numbers for the states, six-day testing requirements to push R_t to 0.9 would range from as much as 5.2% in Alaska to as little as 0.4% in New Hampshire.

Adequate testing is 1.8-4.5mn tests per day

In addition, while 1.8mn tests per day would push the effective reproduction number to 0.9, it would still take 104 days to reduce new infections by half. In the meantime, fallout from potential adverse second-round effects on the economy – such as diminished confidence, credit tightening, or temporary unemployment turning into long-term unemployment – poses substantial downside risks. Moreover, any further easing of remedial measures already in place, such as restrictions on bars and in-person dining, will likely raise R_t . For these reasons, policymakers should consider aiming for a faster reduction in the infectious population by increasing test capacity even more.

To this end, our estimates suggest that testing at a rate of 12.9mn per day would reduce R_t to 0.25, implying a halving of new cases every 14 days. This would require testing between 20% and 26% of the population every six days, depending on the

¹⁵ Using our estimated coefficient on the testing variable from the baseline specification (-0.0630), we calculate the six-day testing requirement as $\ln(0.9/2)/(-0.063)$, which provides a number as a percentage of the population.

state. But this standard is almost surely unrealistic, in part because testing materials are in short supply and in part because the implied magnitude of increase is so large in relation to current testing rates. That said, testing at a rate of about 4.5mn per day would reduce R_t to 0.66, cutting new cases in half every month and providing room for error, potentially restoring household and business confidence by demonstrating that the pandemic is under control and on the road to extinction. This goal amounts to testing nearly 1.4% of the population per day. In sum, we interpret our results as pointing to a minimum near-term testing threshold of 1.8mn per day, with an eventual goal of testing 4.5mn per day.

Masks are powerful in combination with other measures

Our estimates provide perspective on the degree to which official masking requirements are effective in limiting infections. In each of the four model specifications that include our mask variable, we find that the addition of an official state-wide requirement to wear a mask or other face covering has statistically significant explanatory power in lowering R_t . Moreover, we find that the estimated effect of a state-wide order to wear a mask in public is quite large, particularly when used in combination with other remedial measures. For example, when combined with reduced mobility, testing, and stay-at-home orders (model 6), issuing a state-wide order to wear a mask further reduces R_t in the state by about 9.4%, which, in that specification, is nearly equivalent to testing an additional 2% of the state population every six days. This suggests that mask requirements could be particularly useful as a substitute for testing in states where capacity remains low. In states where it is at adequate levels, mask requirements could also help reduce the time needed to eradicate infections.

Restrictions on dining prove useful when states are reluctant to mandate the wearing of masks

Recently, there has been much discussion about whether the reopening of bars and restaurants has contributed to increased growth in new cases, particularly in the South and West. In light of our own results, we also question the long-term efficacy of stay-at-home orders and whether more targeted restrictions could be nearly as effective for controlling the spread of COVID-19. Both considerations led us to include a measure of dining restrictions in our regressions, to determine whether these have a statistically significant effect on the effective reproduction number. In specification 7, we replace the stay-at-home dummy variable with dining

restrictions, finding that broad restrictions on dining, combined with reduced mobility, are about as effective as a combination of stay-at-home orders and reduced mobility. The former reduces R_t about 50%, while the latter reduces it about 60%. Although there is some modest loss in statistical significance, dining restrictions are found to be significant at the 5% level. Altogether, this suggests that they might be a useful substitute for stay-at-home restrictions for reducing the growth of new cases.

We also find suggestive evidence that dining restrictions could be used as a substitute in states that are reluctant to mandate the use of masks. In model 8, when we test mask requirements and dining restrictions together, mask requirements retain their significance at the 1% level, but dining restrictions become statistically insignificant. This suggests that dining restrictions are not effective when mask requirements are in place. To examine this issue further, in specification 9, we interact dining restrictions with states that did not require masks. The requirement to wear a mask retains its sizeable and statistically significant effect in reducing the growth rate of new cases in states that have such orders in place, while we find that the re-opening of restaurants and bars likely contributed to a rise in confirmed cases in states that do not require masks, albeit with significance at about 10%. Hence, we find some evidence to support the view that re-opening these businesses has contributed to new infections in states that do not impose masking requirements. Our results could be viewed as consistent with recent steps taken by governors in Arizona, Florida, Texas, and California, among others, to close bars in all or parts of their respective states shortly after their re-opening in late May.¹⁶

What about herd immunity?

Our regression results provide some support for the theoretical result that R_t will decline as a greater proportion of the population becomes exposed to the disease over time. Indeed, our baseline specification suggests that a 1% increase in the percentage of the population that has tested positive for the virus will reduce the reproduction number about 5.3% and that this result is very robust across specifications. In fact, this baseline estimate is much larger than theory would predict: since the model implies that R_t will decline in proportion with the population

¹⁶ On July 7, California closed indoor operations for a variety of businesses, including bars, in six additional counties, bringing the total to 23 (California has 58 counties in total). Governor Ducey of Arizona ordered bars, gyms, and movie theaters closed on June 29. On June 26, Texas closed bars and limited restaurant occupancy, due to the view that congregating in close spaces was fueling an increase in new COVID-19 cases.

share of the susceptible population, a decline in the susceptible population share of about 1pp should reduce R_t about 1%.

Even if we take this estimate at face value, the notion of controlling the virus by attrition and, eventually, herd immunity seems fanciful. As of July 20, the cumulative number of confirmed positive cases in the US amounts to nearly 3.8mn, or about 1.1% of the population, with the measured case count having increased by about 0.14% of the population over the past week. The cumulative count is only enough to reduce R_t by about 5.9% at present. This is nowhere close to what would be sufficient to control the virus without other remedial measures. For instance, in the absence of testing, our baseline estimates suggest that nearly 5.0% of the population would need to be immune to the virus in order to keep R_t at 0.9. At the current infection rate, this would take a little more than half a year (27 weeks).¹⁷ The costs of this course in terms of public health would almost certainly be very high.

Moreover, given ongoing test shortages, and the likelihood that many have already had asymptomatic cases of COVID-19, there is good reason to think that the number of exposures is substantially understated in the data. Indeed, our estimate of the coefficient on the cumulative testing variable may provide a good sense of the magnitude of understatement. If we assume that the model is approximately correct, so that R_t moves in rough proportion with the susceptible population and the true coefficient on the exposed share of the population is about -0.01, then the estimated coefficient can be interpreted as the multiple that the cumulative case count is being understated. In that case, the true infection count would be about 5.3 times higher than the measured count. For what it is worth, this number is in the same ballpark as the undercount estimated by *COVID19-projections.com* using a very different methodology.¹⁸

That said, even if the number of positive cases is being undercounted by this magnitude in the official estimates, this does not improve the mathematics of herd immunity, as this would still imply that cumulative exposures to date have only reduced R_t by roughly 5.8%.

¹⁷ We calculate this number using our baseline specification by using the coefficient on the testing rate to back out a hypothetical estimate of the current effective reproduction rate if testing were reduced from its current rate (about 1.4% of the population per six days) to zero. We then use the estimated coefficient on the exposed population to determine the additional percentage of the population that would need to be exposed to reduce this hypothetical reproduction rate (1.1, according to our calculations) to 0.9.

¹⁸ As of June 21, *COVID19-projections.com* estimated that about 8.5% of the US population has been infected by COVID-19 infections, compared with the measured count of a little more than 1.1% of the population.

Taking stock

Our results suggest that daily new coronavirus cases are likely to continue to increase in the absence of policy mitigation. Thankfully, policymakers appear to have options to contain the virus that are likely to be less costly than full-scale lockdowns. A thoughtful mix of mitigation measures would likely include dramatic increases in testing, perhaps combined with mask requirements or seating dining restrictions (though not necessarily both).

Although all of the mitigation measures help reduce R_t , effectiveness varies. In particular, our estimates suggest that reductions in personal mobility on the scale of those achieved in March and April can reduce the reproduction rate by about a half and are especially effective when combined with stay-at-home orders. That said, we view such orders to be more viable as an emergency measure to slow an out-of-control outbreak at its early stages, when the effective reproduction number has not yet been mitigated by other measures. Experience suggests that stay-at-home measures are very costly to sustain for an extended period in terms of foregone economic activity and may well become less effective over time as human nature chips away at the general public's resolve.

A well-designed testing regime would likely be a more sustainable alternative. According to our estimates, R_t could be reduced to 0.9 by boosting the current testing rate from its current average level of around 780,000 tests per day to about 1.8mn tests per day. At that rate, the number of active infections would halve every 104 days. Boosting the testing rate even further to 4.5mn tests/day would halve the number of active infections every month.

But testing of this magnitude may prove challenging, in part due to the limited availability of testing agents, equipment, and other supplies. Thus, mandating mask usage seems to be a viable supplemental measure. Our estimates suggest that a mask mandate can reduce R_t nearly 10%, which is roughly equivalent to testing another 1.1mn persons/day. Alternatively, states could implement restrictions that prevent in-person dining at bars and restaurants, which appear to be a reasonable substitute if masks are not mandated.

In our view, a strategy of letting the virus run its course would be very costly in terms of public health. Given the cumulative case count, our best guess is that the herd immunity effect has reduced R_t about 0.06pts, which is equivalent to a reduction

from 1.06 to the current level of 1.0. This is not nearly enough to rely on herd immunity to control the outbreak in the absence of substantial remediation from other policy measures. For instance, we estimate that it will take about a half a year, at the current daily case rate, before the US can rely on herd immunity to keep infections under control without the need for testing. In the meantime, other mitigation measures would need to be intensified to keep infections in check.

References

- Alexander, Diane and Ezra Karger. 2020. "Do Stay-at-Home Orders Cause People to Stay at Home? Effects of Stay-at-Home Orders on Consumer Behavior." FRB of Chicago Working Paper WP 2020-12.
- Alfaro, Laura, Ester Faia, Nora Lamersdorf, and Farzad Saidi. 2020. "Social Interactions in Pandemics: Fear, Altruism, and Reciprocity." NBER Working Paper 27134.
- Barrios, John M., Efraim Benmelech, Yael V. Hochberg, Paola Sapienza, and Luigi Zingales. 2020. "Civic Capital and Social Distancing during the Covid-19 Pandemic." NBER Working Paper 27320.
- Chen, Keith, Yilin Zhuo, Malena de la Fuente, Ryne Rohla, and Elisa F. Long. 2020. "Causal Estimation of Stay-at-Home Orders on SARS-CoV-2 Transmission." UCLA Anderson School Working Paper, May 15, 2020.
- Coibon, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2020. "The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending." NBER Working Paper 27141.
- Dave, Dhaval M., Andrew I. Friedson, Kyutaro Matsuzawa, Drew McNichols, and Joseph J. Sabia, "Did the Wisconsin Supreme Court Restart a COVID-19 Epidemic? Evidence from a Natural Experiment," NBER Working Paper No. 27322, June 2020.
- Dave, Dhaval M., Andrew I. Friedson, Kyutaro Matsuzawa, and Joseph J. Sabia. 2020. "When Do Shelter-in-Place Orders Fight COVID-19 Best? Policy Heterogeneity across States and Adoption Time." NBER Working Paper 27091.
- Dave, Dhaval M., Andrew I. Friedson, Kyutaro Matsuzawa, Joseph J. Sabia, and Samuel Safford. 2020. "Were Urban Cowboys Enough to Control COVID-19? Local Shelter-in-Place Orders and Coronavirus Case Growth." NBER Working Paper 27229.
- Dave Dhaval, Andrew I. Friedson, Drew McNichols, and Joseph J. Sabia, 2020. "Did California's Shelter-in-Place Order Work? Early Coronavirus-Related Public Health Effects." NBER Working Paper 26992.
- Deb, Pragyan, Davide Furceri, Jonathan D. Ostry and Nour Tawk, (2020). "The Effect of Containment Measures on the COVID-19 Epidemic." COVID Economics, Issue 19.
- Fang, Hanming, Long Wang, and Yang Yang. 2020. "Human Mobility Restrictions and the Spread of the Novel Coronavirus (2019-nCoV) in China." NBER Working Paper 26906.
- Goolsbee, Austan, and Chad Syverson. 2020. "Fear, Lockdown, and Diversion: Comparing Divers of Pandemic Economic Decline in 2020", NBER Working Paper 27423.
- Gupta, Sumedha, Laura Montenegro, Thuy Nguyen, Felipe Lozano Rojas, Ian M. Schmutte, Kosali I. Simon, Bruce A. Weinberg, and Coady Wing. 2020. "Effects of Social Distancing Policy on Labor Market Outcomes." NBER Working Paper 27280.

- Gupta, Sumedha, Thuy D. Nguyen, Felipe Lozano Rojas, Shyam Raman, Byungkyu Lee, Ana Bento, Kosali I. Simon, and Coady Wing. 2020. "Tracking Public and Private Responses to the COVID-19 Epidemic: Evidence from State and Local Government Actions." NBER Working Paper 27027.
- Jiang, Xuan, Laura Montenegro, Felipe Lozano Rojas, Kosali I. Simon, Bruce A. Weinberg, and Coady Wing. 2020. "Is the Cure Worse than the Problem Itself? Immediate Labor Market Effects of COVID-19 Case Rates and School Closures in the U.S." NBER Working Paper 27127.
- Maloney, William and Temel Taskin. 2020. "Social Distancing and Economic Activity during COVID-19: A Global View." COVID Economics, Issue 13.
- Nguyen, Thuy D., Sumedha Gupta, Martin Andersen, Ana Bento, Kosali I. Simon, and Coady Wing. 2020. "Impacts of State Reopening Policy on Human Mobility." NBER Working Paper 27235.
- Tian, H. Yonghong Liu, Yidan Li, Chieh-Hsi Wu, Bin Chen, Moritz U. G. Kraeme, 2020. "An investigation of transmission control measures during the first 50 days of the COVID-19 epidemic in China." *Science*, 368(6491).

Who suffers from the COVID-19 shocks? Labor market heterogeneity and welfare consequences in Japan¹

Shinnosuke Kikuchi,² Sagiri Kitao³ and Minamo Mikoshiba⁴

Date submitted: 20 July 2020; Date accepted: 23 July 2020

Effects of the COVID-19 shocks in the Japanese labor market vary across people of different age groups, genders, employment types, education levels, occupations, and industries. We document heterogeneous changes in employment and earnings in response to the COVID-19 shocks, observed in various data sources during the initial months after onset of the pandemic in Japan. We then feed these shocks into a life-cycle model of heterogeneous agents to quantify welfare consequences of the COVID-19 shocks. In each dimension of the heterogeneity, the shocks are amplified for those who earned less prior to the crisis. Contingent workers are hit harder than regular workers, younger workers than older workers, females than males, and workers engaged in social and non-flexible jobs than those in ordinary and flexible jobs. The most severely hurt by the COVID-19 shocks has been a group of female, contingent, low-skilled workers, engaged in social and non-flexible jobs and without a spouse of a different group.

1 We acknowledge financial support from the Center for Advanced Research in Finance (CARF) at the University of Tokyo.

2 Massachusetts Institute of Technology.

3 University of Tokyo.

4 University of Tokyo.

Copyright: Shinnosuke Kikuchi, Sagiri Kitao and Minamo Mikoshiba

1 Introduction

The COVID-19 pandemic has brought significant shocks to the labor markets all over the world, and Japan is no exception. While Japan has not seen a sharp increase in unemployment rate, which stood at 2.9% in May 2020, compared to other countries such as 14.9% in the United States in April 2020, the shocks in the labor market are spread highly unequally across workers.^{1,2} In this paper, we first document heterogeneous responses in employment and earnings to the COVID-19 shocks observed during the initial months after onset of the crisis in Japan. We then feed these shocks in the labor market into a life-cycle model of heterogeneous agents to quantify welfare consequences of the COVID-19 shocks.

Despite a relatively small change in the overall unemployment rate, we find that negative effects of the COVID-19 shocks significantly differ across individuals workers, in various dimensions including age group, gender, employment type, education level, occupation, and industry. Moreover, in each dimension, the shock is larger for those who earned less prior to outbreak of the pandemic, amplifying inequality in the labor market across multiple dimensions.

To quantify welfare effects from the COVID-19 shocks, we build a life-cycle model and let heterogeneous individuals face unexpected changes to their earnings and employment, as observed in the data, and have them re-optimize in response to the shocks. We evaluate welfare effects on different types of individuals in terms of consumption equivalent variation that would make them as better off as before in the economy in the absence of the COVID-19 shocks.

Our findings can be summarized as follows. First, contingent workers suffer significantly, up to more than nine times as much as regular workers in terms of our welfare measure. They are more severely hurt in both employment and wages than regular workers, and we find that employment type is one of the most critical dimensions that divides the fate of individuals in the labor market after the crisis. Second, we also find that younger generations suffer more than older generations. Third, female workers fare worse than males and their negative welfare effects are three times as large as those of male workers. The difference is mainly due to the fact that the share of contingent workers is larger for females, but also because females are more concentrated in jobs that are more severely affected by the COVID-19 shocks. Forth, workers in social sectors and/or non-flexible occupations suffer more. The COVID-19 crisis differs from past recessions such as

¹The Japanese unemployment rate is from the Labor Force Survey (LFS) of the Ministry of Internal Affairs and Communications (MIC). The U.S. unemployment rate is from the Labor Force Statistics of the Current Population Survey (CPS). The U.S. unemployment rate peaked in April and declined to 13.3% and 11.1% in May and June, respectively.

²Kikuchi et al. (2020) discuss heterogeneity of potential vulnerability of workers to the COVID-19 shocks using data prior to the crisis.

the financial crisis of 2008 in that it contracts economic activities in sectors that involve more face-to-face transactions and occupations involving tasks difficult to be completed remotely from homes or in physical isolation from other people. [Kikuchi et al. \(2020\)](#) discussed heterogeneous vulnerability across occupations and industries and pointed to risks of rising inequality, which we confirm has manifested in wage and employment changes across workers in the data during the first quarter after the crisis.

We also stress caution in the interpretation of our quantitative results. As discussed above, the main focus of our paper is to assess changes in the labor market during the initial months after onset of the COVID-19 crisis, which we observed in various official data, and to quantify welfare implications from these observations. For this purpose, we build a simple life-cycle model of heterogeneous agents that enables us to focus on the analysis of these effects in the short-run. There is, however, significant uncertainty about whether various shocks we observe now will be short-lived or long-lived and whether they will be repeated multiple times over years to come. We evaluate welfare effects under some scenarios about the duration of shocks and our results may need to be re-examined once more data is available and there is less uncertainty as to the magnitude and duration of the pandemic.

Moreover, there may well be other structural changes in the economy that the COVID-19 crisis may induce over the medium and long-run. There are also many changes that the Japanese economy had been going through, including changes in the composition of employment type and gender-specific involvement in the labor market, aging demographics, fiscal challenges associated with rising expenditures on the social insurance system. The COVID-19 crisis may interact with these changes and possibly amplify challenges that Japan is faced with in some dimensions, or hopefully mitigate them in other dimensions. Although we acknowledge these topics and potential consequences of the COVID-19 crises in the medium and long-term as very important and worth exploring, they are not in the scope of the current analysis and our model intentionally abstracts from them.

Numerous studies have emerged that investigate heterogeneous consequences of the COVID-19 shocks on individuals and implications for welfare and policies, which include but are not limited to [Acemoglu et al. \(2020\)](#), [Alon et al. \(2020\)](#), [Glover et al. \(2020\)](#), [Kaplan et al. \(2020\)](#), and [Albanesi et al. \(2020\)](#), just to name a few.³ Our study complements the literature by documenting facts and analyzing welfare consequences in Japan.

This paper is also complementary to studies of various economic aspects of the COVID-19 shocks in Japan. They include [Fukui et al. \(2020\)](#) on the impact of pandemic on job vacancy postings, [Watanabe and Omori \(2020\)](#) on consumption responses across sectors, [Miyakawa et al. \(2020\)](#) on firm default, [Kawata \(2020\)](#) on occupational and spatial

³Other papers that document and study early responses to the COVID-19 shocks in the U.S. labor market include [Coibion et al. \(2020\)](#), [Gregory et al. \(2020\)](#) and [Kahn et al. \(2020\)](#).

mismatch, Kawaguchi et al. (2020) on uncertainty faced by small and medium-sized firms, and Okubo (2020) on implementation of telework across occupations.

The rest of the paper is organized as follows. Section 2 provides an overview of economic shocks triggered by the COVID-19 shocks observed in the early data and lays out facts that our model analysis in the following sections is focused on. Section 3 presents our dynamic life-cycle model and section 4 discusses parametrization of the model. Numerical results are discussed in section 5 and section 6 concludes. The appendices provide more details about the data sources and discusses our computation methods.

2 Impact of the COVID-19 Shocks on the Labor Market in Japan

This section documents changes in employment and earnings during the COVID-19 crisis. The data source of our analysis is mainly Labor Force Survey (LFS) data for monthly employment, and is supplemented by Monthly Labor Survey (MLS) data for monthly earnings and Employment Status Survey (ESS) data in 2017 for composition of workers across different categories.

2.1 Data Sources

We provide a brief explanation of the three labor market data sources: LFS, MLS, and ESS below. Detailed description of these data sets is provided in appendix A.

Labor Force Survey (LFS): The LFS is a monthly cross-sectional household survey conducted by the Ministry of Internal Affairs and Communications (MIC). It covers approximately 40 thousand households across the nation and collects detailed information about the employment status of household members. We use publicly available tabulated data to compute employment by age, gender, employment type, industry, and occupation.

Monthly Labor Survey (MLS): The MLS is a monthly cross-sectional monthly survey conducted by the Ministry of Health, Labour and Welfare (MHLW), which covers approximately 33 thousand establishments and their employees from the private and public sectors. We use publicly available tabulated data to compute earnings by employment type and industry.

Employment Status Survey (ESS): The ESS is a cross-sectional household survey conducted every five years by the MIC. For our research purpose, we use the latest data collected in October 2017. It is one of the most comprehensive surveys on employment circumstances in the nation. It covers approximately 490 thousand households and provides

detailed information about the demographic characteristics of households, employment and unemployment situations, and descriptions of current jobs held by household members. We use the “order-made” summarization system to compute joint distribution of workers and earnings prior to the crisis, across age groups, genders, education levels, employment types, industries, and occupations.⁴

Besides the three data sources for labor market statistics, we also use the Family Income and Expenditure Survey (FIES) data for changes in consumption level and allocations. More details about the data sources are provided in appendix A.

2.2 Classification of Workers

We briefly explain below how we classify workers according to three different dimensions: employment type, industry and occupation. More details about the classifications in each of the data sources are given in appendix A.

Employment-Type Categories: Employment in the Japanese labor market is characterized by a distinction in employment type: regular or contingent employment. How they are termed in the Japanese language differs depending on situations and data source. In the ESS, for example, regular employment includes executives of companies and staff members who are termed regular (*seiki*) employees. Contingent (*hiseiki*) employment includes part-time workers, albeit (temporary workers), dispatched workers, contract employees and others. Contingent workers are sometimes termed irregular or non-regular workers as well.⁵

The distinction is different from that between full-time and part-time workers in other countries. Contingent workers may well work for the same number of hours as regular workers but they tend to receive lower wages, fewer fringe benefits, and much less job security than regular workers. As documented in papers such as [Kitao and Mikoshiba \(2020\)](#) and [İmrohoroğlu et al. \(2016\)](#), earnings of contingent workers are much lower among both males and females. Females have a higher fraction of contingent workers than males and so do less educated workers than those with higher education. Moreover and most importantly, contingent workers are subject to more frequent employment adjustment and job instability, as shown in empirical studies including [Yokoyama et al. \(2019\)](#) and [Esteban-Prete et al. \(2011\)](#). In the analysis below, we include employment status as one of the key dimensions of heterogeneity across workers in evaluating effects of the COVID-19 crisis.

⁴The ESS data is based on statistical products provided by the Statistics Center, an independent administrative agency based on the Statistics Act, as a tailor-made tabulation of the 2017 ESS compiled by the MIC.

⁵How workers are divided into the two employment types in each database we used is explained in appendix A.

Sectoral Categories: Following Kaplan et al. (2020), we classify industries into two sectoral categories: ordinary and social.⁶ Based on the distribution of workers across sectors in the ESS, 48% of total employment is classified into the ordinary sector, and the remaining 52% is classified into the social sector, prior to the COVID-19 shocks.

- **Ordinary Sector:** agriculture and forestry, fisheries, mining and quarrying of stone and gravel, construction, manufacturing, electricity, gas, heat supply and water, information and communications, road freight transport, water transport, warehousing, service incidental to transport, postal service including mail delivery, wholesale, finance and insurance, real estate and goods rental and leasing, postal service.
- **Social Sector:** railway transport, road passenger transport, air transport, retail trade, scientific research, professional and technical services, accommodations, eating and drinking services, living-related and personal services and amusement services, education, learning support, medical, health care and welfare, cooperate associations, services, N.E.C., government except elsewhere classified.

Note that not all data sources provide sector information of the same accuracy, and we use a broader classification for the MLS. Also, we use a slightly different categorization for the expenditure data from the FIES. For more details, see sections A.2 and A.4, respectively.

Occupational Categories: We classify occupations into two occupational categories, flexible and non-flexible occupations, based on the fraction of workers in each occupation who are likely to work remotely and less affected by difficulty in commuting to and working in their regular workplace.⁷ Following Mongey et al. (2020), we construct measures of the fraction of flexible-type workers in each occupation. Figure 1 shows the result. We then classify occupations as flexible if the measure is larger than 0.75. As a result, 60% of total employment is classified into flexible occupation, and the remaining 40% is classified into non-flexible occupation.

⁶We use industrial categories defined in the Japan Standard Industrial Classification (JSIC), as revised in 2013.

⁷We use occupational categories defined in the Japan Standard Occupational Classification (JSOC), as revised in December 2009.




Figure 1: Work-from-home Measures: JSOC

Note: This figure shows the fraction of workers who are able to work from home in each occupation. To compute the measure, we follow ? and convert the Standard Occupational Classification (SOC) to the Japan Standard Occupational Classification (JSOC).

- Flexible Occupation: administrative and management workers, professional and engineering workers, clerical workers, sales workers.
- Non-flexible Occupation: service workers, security workers, agriculture, forestry and fishery workers, manufacturing process workers, transport and machine operation workers, construction and mining workers, carrying, clearing and packaging, and related workers.

2.3 Changes in Employment

This section documents changes in employment in Japan during the COVID-19 crisis. The data source is LFS data for most of the analysis, and ESS data for compositional analysis.

By Employment Type, Sector and Occupation: Figure 2a shows the number of employed by employment type (regular and contingent). We normalize to 100 the level of employment for each type in January 2020. While regular workers' employment declined by around 1% in April and May compared to January, contingent workers' employment declined more sharply by around 4% to 5%. This is consistent with previous episodes in

Japan where contingent workers have been more vulnerable to business cycle shocks, as documented by [Yokoyama et al. \(2019\)](#).

Figure 2b shows the number of employed according to the sectoral and occupational categories defined above. The number of workers in the social sector and non-flexible occupations declined the most, by more than 5% from January to April 2020. The difference across sectors and occupations highlights the importance of the feasibility of completing work from home, as emphasized by [Dingel and Neiman \(2020\)](#) in the case of the US labor market and [Fukui et al. \(2020\)](#) based on changes in the pattern of job vacancy postings in Japan after the COVID-19 shocks.

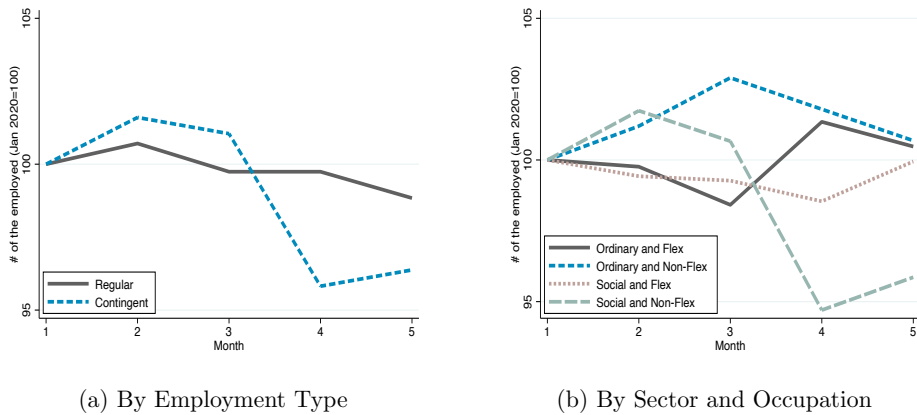


Figure 2: Changes in Employment (Jan. 2020 = 100)

Note: Figure 2a shows the number of employed by employment type in each month between January and May 2020. We restrict samples to workers aged 25 to 64. Figure 2b shows the number of employed by sector and occupation categories by monthly frequency. The samples are all workers aged 15 to 64, including not only regular and contingent workers but also other types of workers such as the self-employed, since the more granular age and employment-type categories cannot be obtained from publicly available aggregated data. In both figures, the values in January 2020 are normalized to 100, and series are not seasonally adjusted. The data is from Labor Force Survey (LFS) by the Ministry of Internal Affairs and Communications (MIC).

By gender: Figure 3 shows changes in the number of employed by gender, where the level in January 2020 is normalized to 100. While both males' and females' employment declined since February 2020, the decline is larger for females. This is similar to what occurred in the U.S. where female workers were hit harder by the COVID-19 shocks than male workers, as emphasized by [Alon et al. \(2020\)](#).




Figure 3: Changes in Employment by Gender (Jan. 2020 = 100)

Note: Figure 3 shows the number of employed by gender in each month between January and May 2020. We restrict samples to workers aged 25 to 64. The values in January 2020 are normalized to 100, and series are not seasonally adjusted. The data is from Labor Force Survey (LFS) by the Ministry of Internal Affairs and Communications (MIC).

Why have female workers suffered more from the COVID-19 shocks? Figure 4 shows the characterization of workers by gender based on the ESS data prior to the COVID-19 crisis. Figure 4a displays the share of contingent workers out of total employment by gender. While the share of contingent workers is less than 10% for males, more than 50% of female workers work a contingent job. This difference partially contributes to larger decline for female employment, since contingent workers are subject to more employment adjustment during economic downturns as discussed above, and in fact, there was a larger decline in employment among contingent workers as we show below.

Figure 4b shows the share of workers in the social sector out of total employment by gender. Again, female workers are more concentrated in the social sector (69%) than male workers (39%). Figure 4c shows the share of workers in non-flexible occupations out of total employment by gender. In contrast to employment type and sector, male workers appear to be more vulnerable in terms of the non-flexibility of the work arrangement, though the difference is relatively small.⁸ Figure 4d, however, which shows the joint distribution of employment across sectors and occupations, reveals that the share of the most vulnerable workers engaged in social *and* non-flexible jobs is higher for females than males. The share of the least vulnerable workers in ordinary and flexible jobs is larger for

⁸The share of non-flexible occupations is 46% for males and 34% for females.

males than females as well.

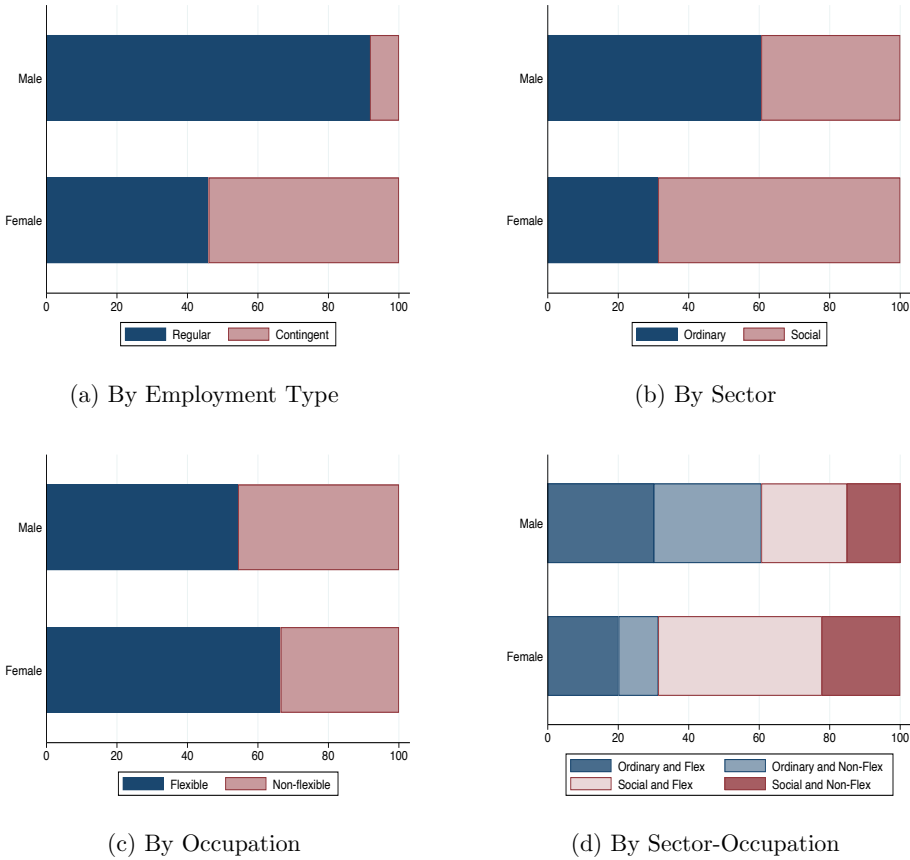


Figure 4: Share of Each Characteristics by Gender

Note: Figure 4 shows the employment share for each characteristic by gender. We restrict samples to workers aged between 30 and 59 because the data is available only for 10-age bin. The data is from Employment Status Survey (ESS) conducted in 2017 by the Ministry of Internal Affairs and Communications (MIC).

By Age Group: Figure 5a and Figure 5b show the number of employed by age for regular workers and contingent workers, separately. We normalize the level in January 2020 to 100. For regular workers, changes during the first five months of the year are modest. For contingent workers, the decline by April 2020 is much larger in the range of 4 to 5% relative to the level in January 2020. Across age groups, changes from January 2020 to April 2020 are similar, but the decline from the first quarter to April and May of 2020 is larger for younger cohorts. We discuss this heterogeneity in employment across

age groups and employment types in more details in section 5.2.

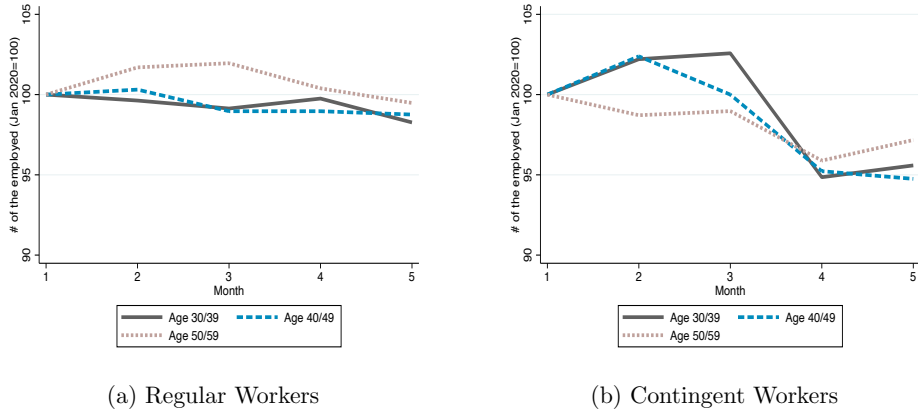


Figure 5: Changes in Employment by Age Group (Jan. 2020 = 100)


Note: Figure 5a shows the number employed by age for regular worker in each month between January and May 2020. Figure 5b shows the number of employed by age for contingent workers during the same period. The values in January 2020 are normalized to 100. Samples are restricted to workers aged 25 to 64. Series are not seasonally adjusted. The data is from Labor Force Survey (LFS) by the Ministry of Internal Affairs and Communications (MIC).

2.4 Changes in Earnings

This section documents changes in earnings in Japan during the COVID-19 crisis, based on the MLS data. Figure 6 shows year-on-year changes in earnings in the ordinary and social sectors for regular workers and contingent workers, separately.

As shown in Figure 6a, earnings of regular workers barely changed during the first quarter of 2020 compared to the same months of the previous year. The average earnings in both sectors declined in April 2020, but only by approximately 1% compared to the April 2019 and the magnitude of the change is similar in both ordinary and social sectors.

6b shows a very different picture for year-on-year changes in earnings for contingent workers. In April 2020, the average earnings for all sectors declined by 4% compared to the April 2019. There are also significant differences in the changes across sectors. For workers in the social sector, earnings declined by 7% while those in ordinary sectors did not experience a decline.



(a) Regular Workers

(b) Contingent Workers

Figure 6: Changes in Earnings by Sector (in YOY change, %)

Note: Figure 6a and 6b show changes in fixed earnings (excluding seasonal bonus) by sector for regular and contingent workers, respectively. The values are in year-on-year change by monthly frequency, that is, they compare changes in earnings from the same month in the previous year. The data is from the Monthly Labor Survey (MLS) by the Ministry of Health, Labour and Welfare (MHLW).

3 Model

Demographics: At age $j = 1$, individuals enter the economy with initial assets denoted as a_1 . Individuals face probability s_j of surviving from age $j - 1$ to j . S_j denotes unconditional survival probability that an individual lives up to age j . We assume that they retire at the age of $j = J^R$ and live up to the maximum age of $j = J$. The deceased will be replaced by the newborn. Population is assumed to be constant and age distribution is stationary.

Endowment and Earnings: Individuals are born with gender $g = \{M, F\}$, male or female, and a skill type $s = \{H, L\}$, high or low. Upon entering the labor market, they are also assigned to an employment type $e = \{R, C\}$, regular or contingent, an occupation $o = \{o_1, o_2\}$, and sector $d = \{d_1, d_2\}$.

The two occupation types, o_1 and o_2 , are associated with different levels of work flexibility, i.e. whether the job can be done remotely from home or not. The two sectors, $d = \{d_1, d_2\}$, produce different types of goods and services. Sector d_1 produces ordinary goods while sector d_2 produces social goods, which are more immune to infection risk in terms of consumption.

We let $x = \{j, g, s, e, o, d\}$ denote a state vector of each individual. We denote by μ_x the population share of individuals in state x , that is, age j , gender g , skill s , employment

type e , occupation o , and sector d . Each individual's efficiency units of labor depend on the state vector x and are denoted as η_x , which varies over a life-cycle and approximates human capital that grows in age for each type of workers.

Earnings of an individual in state x at time t are given by

$$y_{x,t} = \lambda_{x,t} \eta_x w_t.$$

$\lambda_{x,t}$ summarizes shocks that affect earnings of type- x individuals at time t , which will be discussed in detail in section 5.2. w_t denotes the market wage per efficiency unit of labor.

Preferences: Individuals derive utility from consumption of two types of goods, c_1 and c_2 , representing ordinary and social goods, respectively. We assume a period utility function:

$$U(c_1, c_2) = \xi_t \frac{[c_1^{\gamma_t} c_2^{1-\gamma_t}]^{1-\sigma}}{1-\sigma}, \tag{1}$$

where ξ_t represents an intertemporal preference shifter that affects marginal utility from consumption in each period. It is a weight on utility from consumption at time t relative to other times and may change with the arrival of the COVID-19 shocks, but it is assumed to be constant in normal times.

γ_t is a preference weight on ordinary goods, which, similarly to ξ_t , is constant in normal times, but may vary upon the arrival of the COVID-19. σ represents risk aversion. Individuals discount future utility at constant rate β .

There are no bequest motives and assets a_{t+1} left by the deceased are collected and transferred to all surviving individuals as accidental bequests, denoted as b_t , which satisfies the following equation.

$$b_t = \frac{\sum_x a_{t+1}(x)(1 - s_{j+1})\mu_x}{\sum_x \mu_x} \tag{2}$$

Government: The government operates a social security program, which provides a pension benefit p_t to each retiree. Individuals are taxed on their consumption, labor income and capital income at proportional rates, $\tau_{c,t}$, $\tau_{l,t}$, and $\tau_{a,t}$, respectively. We assume that the government budget is balanced each period and let a lump-sum transfer $\tau_{ls,t}$ absorb an imbalance from the period budget constraint (3).

$$\sum_x [\tau_{c,t}(c_{1,t}(x) + c_{2,t}(x)) + \tau_{a,t}r_t(a_t(x) + b_t) + \tau_{l,t}\lambda_{x,t}\eta_x w_t] \mu_x = \sum_{x|j \geq j^R} p_t \mu_x + \sum_x \tau_{ls,t} \mu_x \tag{3}$$

Life-cycle Problem: The intertemporal preference ordering of an individual of type x born at time t is given by:

$$U(\{c_{1,t+j-1}, c_{2,t+j-1}\}_{j=1}^J) = \sum_{j=1}^J \beta^{j-1} S_j \xi_{t+j-1} \frac{[c_{1,t+j-1}^{\gamma_{t+j-1}} c_{2,t+j-1}^{1-\gamma_{t+j-1}}]^{1-\sigma}}{1-\sigma}$$

subject to:

$$\begin{aligned} (1 + \tau_{c,t})(c_{1,t} + c_{2,t}) + a_{t+1} &= (1 - \tau_{l,t})\lambda_{x,t}\eta_x w_t + R_t(a_t + b_t) + \tau_{ls,t} \text{ for } j < j^R \\ (1 + \tau_{c,t})(c_{1,t} + c_{2,t}) + a_{t+1} &= p_t + R_t(a_t + b_t) + \tau_{ls,t} \text{ for } j \geq j^R \end{aligned}$$

where $R_t = 1 + (1 - \tau_{a,t})r_t$ denotes net-of-tax gross interest rate at time t .

Initial Economy and Transition Dynamics The initial economy is stationary and characterized by demographics, $\{s_j\}_{j=1}^J$ and μ_x , type-specific labor productivity, η_x , a set of fiscal variables, $\{\tau_c, \tau_l, \tau_a, p\}$, factor prices, $\{r, w\}$, where individuals choose the optimal path of consumption and assets $\{c_1, c_2, a'\}$ at each age j . In equilibrium a lump-sum tax, τ_{ls} , balances the government budget (3) and the accidental bequest, b , satisfies the condition (2).

At time 1, we assume that individuals are hit by employment and wage shocks summarized in $\lambda_{x,t}$, which we will fully characterize in section 5.2, as well as by preference shocks, ξ_t and γ_t . Given the new paths of earnings and preferences, individuals re-optimize and choose a new path of consumption and assets. We let $\tau_{ls,t}$ adjust to balance the government budget to satisfy (3) in each period as well bequests b_t to meet the condition (2).

4 Calibration

This section describes parametrization of the economy presented above. The model frequency is quarterly. The initial economy approximates the Japanese economy prior to onset of the COVID-19 shocks. We compute the transition dynamics starting in the first quarter of 2020, which corresponds to our initial economy. Parametrization of the initial economy is explained in this section and summarized in Table 1. The shocks that characterize the COVID-19 crisis are discussed in section 5.2.

4.1 Demographics

Individuals of the model enter the economy and start working at the age of 25, and they may live up to the maximum age of 100 years subject to age-specific survival probabilities s_j . The retirement age j^R is set at 65 years old. We calibrate the probabilities based on the estimates of the National Institute of Population and Social Security Research (IPSS) for the year 2020. We abstract from population growth and age distribution is stationary.

4.2 Preferences

The risk aversion parameter, σ , in the utility function (1) is set to 2.0. The parameter γ in the initial economy represents a weight on ordinary goods relative to social goods and it is set at 0.789 so the model matches the ratio of consumption expenditures of the two types of goods, based on the Family Income and Expenditure Share (FIES) from the Ministry of Internal Affairs and Communications (MIC). The parameter ξ that represents an intertemporal weight on consumption is set at 1 in the initial economy. In section 5.4, we simulate time-varying preference weights to approximate consumption data observed during the initial months of the COVID-19 crisis.

The subjective discount factor β is set at 1.00532 (or 1.0215 on an annual basis) to match the average growth of consumption between ages 25 and 50 as observed in the FIES data estimated in İmrohoroğlu et al. (2019).

4.3 Endowment and Human Capital

Each individual is endowed with a unit of time and supplies labor inelastically until they reach the retirement age j^R . The labor productivity $\eta_{j,g,s,e,o,d}$, which represents human capital of an individual worker and evolves over a life-cycle, is calibrated with the ESS data. Details about the categorization of individual workers into employment type, education level, industry and occupation are provided in appendix A.

We assume that the type of individual worker is determined upon entry to the labor market and fixed throughout their life-cycle. The share of each type is based on the distribution from the ESS data, and we take the average share of types among individuals aged between 30 and 59.

4.4 Government and Other Parameters

The pay-as-you-go social security program provides pension benefits p to each retiree. We assume that benefits are set to 30% of average earnings in the initial economy, based on the estimated replacement rate of social security benefits by the OECD.⁹

The consumption tax rate, τ_c , is set to 10%. Labor and capital income tax rates, τ_l and τ_a , are set to 13% and 20%, respectively, following İmrohoroğlu et al. (2019). The lump-sum transfer τ_{ls} is determined in equilibrium to absorb an imbalance from the government budget and is set to 4.84% of average earnings in the initial economy.

We set the interest rate at 0% and wage rate is normalized so that the average earnings in the initial economy is 1.

⁹OECD Pension at a Glance, 2020.

Table 1: Parameters of the Model: Initial Economy

Parameter	Description	Value
<i>Demographics</i>		
J^R	Retirement age	65 years
J	Maximum age	100 years
$\mu_{j,g,s,e,o,d}$	Population share	ESS data
<i>Preference</i>		
β	Subjective discount factor	1.0215 (annual)
σ	Risk aversion parameter	2.0
γ	Expenditure share on regular goods	0.789 (FIES)
ξ	Intertemporal weight	1 (before shock)
<i>Human Capital</i>		
$\eta_{j,g,s,e,o,d}$	Life-cycle human capital	ESS data
λ	Shocks to earnings	1 (before shock)
<i>Government</i>		
τ_c	Consumption tax rate	10%
τ_l	Labor income tax rate	13%
τ_a	Capital income tax rate	20%
τ_{ls}	Lump-sum tax/transfer	4.8% of avg. earn
p	Social security benefit	30% of avg. earn
<i>Other Parameters</i>		
r	Interest rate	0%
w	Wage rate	Normalization

5 Numerical Results

5.1 Baseline Model: Initial Economy

Figure 7 shows the earnings profile based on ESS data as discussed in section 4, for selected types of workers. The left panel shows average earnings of all workers at each age, normalized to the average earnings of all workers. It exhibits a hump-shaped profile, where earnings rise monotonically after the entry and peak at around age 55, when they start to decline. The right panel shows profiles for each gender and employment type and highlights a stark difference in earnings by individual characteristics.

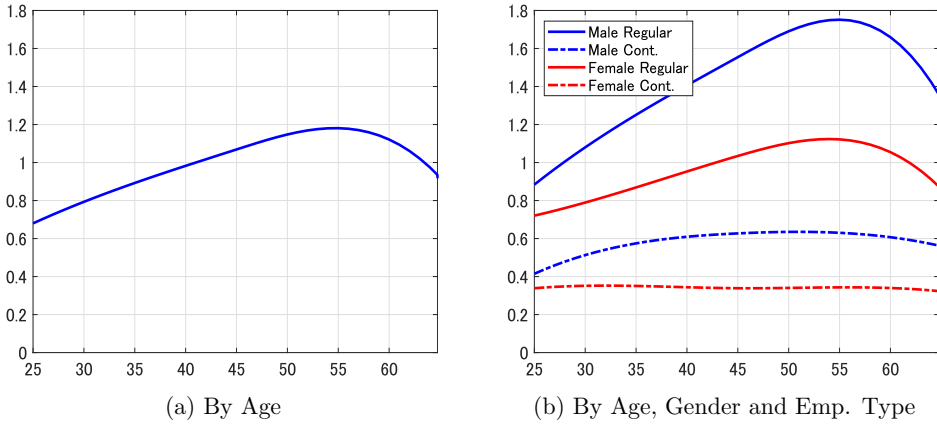


Figure 7: Earnings in the Initial Economy (in model units; average earnings=1)

Solving the model described above, we obtain consumption and asset profiles of individuals averaged for each age, as shown in Figure 8.¹⁰




Figure 8: Consumption and Assets in the Initial Economy (in model units; average earnings=1)

5.2 The COVID-19 Shocks

We will next discuss the COVID-19 shocks that are introduced in the initial economy described above, before we study how they affect welfare of heterogeneous individuals in the model economy in section 5.3. This section revisits the data description presented

¹⁰Note that assets are expressed in terms of average annual earnings, with an adjustment for quarterly frequency of the model.

in section 2 and explains how we process them as shocks that we feed into our model. We will decompose shocks into five, three associated with wage and employment shocks and two associated with preferences. Our main focus will be the first three. Table 2 summarizes five different types of shocks that we consider in the simulations.

Wage and Employment Shocks: Earnings of an individual in state x are hit by wage and employment shocks, summarized in $\lambda_{x,t} \equiv \omega_{e,d,t} \phi_{o,d,t} \nu_{j,e,t}$. This decomposition captures shocks to wages, $\omega_{e,d,t}$, and to employment, $\phi_{o,d,t}$ and $\nu_{j,e,t}$.

Wage shocks, $w_{e,d,t}$, are specific to the industry and vary by employment type, and they are measured as a change in earnings between the first and the second quarters of 2020, using the MLS data.¹¹ The shocks vary across the combination of employment type and industry, $(e, d) = (1,1), (1,2), (2,1), (2,2)$, independently of other states of an individual, and are set to $\{w_{1,1}, w_{1,2}, w_{2,1}, w_{2,2}\} = \{1.000, 0.999, 0.990, 0.946\}$ based on the quarterly change in the data. Workers with contingent employment type in the social sector experience a wage decline of 5.4% and are the most severely hurt, while the change is relatively small for those in the ordinary sector and the social sector but with the regular employment type.

Employment shocks consist of two parts, employment type shock, $\nu_{j,e,t}$, and occupation-sector specific shock, $\phi_{o,d,t}$. We calculate the employment type shock, $\nu_{j,e,t}$, from a change in the number of employees between the first and the second quarters of 2020, using the LFS data. Changes in employment by employment type vary by age, and we assume that the shock is age dependent. Figure 9 displays the decline in employment of contingent workers relative to regular workers and shows that employment type shocks hit younger workers harder than older workers.

¹¹We use monthly data since January 2013. Before calculating the shocks, we seasonally adjust raw data by converting data from monthly to quarterly frequency. We use the data in April and May, and assume that the level in June remains unchanged from that of May in computing the quarterly change in the labor market. Please see appendices A and B for detailed data structures and definitions.




Figure 9: Employment-type Shocks: Change in Employment of Contingent Workers Relative to Regular Workers (Regular=1, 2020Q1 vs 2020Q2)

Note: This graph shows changes in the number of contingent workers relative to regular workers from age 25 to 65. The data is from the Labor Force Survey (LFS) by the Ministry of Internal Affairs and Communications (MIC).

The occupation-sector specific employment shocks, $\phi_{o,d,t}$, are computed for each combination of $(o, d) = (1,1), (1,2), (2,1), (2,2)$ and are set at $\{\phi_{1,1}, \phi_{1,2}, \phi_{2,1}, \phi_{2,2}\} = \{1.003, 0.996, 0.990, 0.956\}$. Employment of workers engaged in non-flexible occupations in the social sector is the most severely hurt, falling by 4.4%, while the change is relatively small for those in flexible occupations, or non-flexible but in the ordinary sector.¹²

Preference Shocks: Preference shocks are captured by share parameter shock, γ_t , and intertemporal preference shock, ξ_t .¹³ The preference parameters are summarized in Table 2.

Figure 10 shows the expenditure share for social goods from the FIES data. Until the

¹²In computing the decline of employment by occupation and sector, we also use the LFS and ESS data of MIC. Since the LFS data only observe employment change of all type- (o, d) workers, shocks using only LFS may be biased by age-composition. Therefore, we use computed employment shocks $\nu_{j,e,t}$ and the ESS data to isolate shocks associated with industry and occupation in a way that is consistent with the aggregate changes in employment for each occupation and sector. More details of the computation are given in appendix B.

¹³Similarly to wage and employment shocks, we use monthly consumption data from January 2013 by converting to quarterly data and seasonally adjusting them. We use consumption data in April and May and assume that the level in June remains unchanged from that of May in computing the quarterly change in the consumption shares and levels. Please see appendices A and B for detailed data structure and definitions.

first quarter of 2020, the expenditure share of social goods remained stable at 21.1% on average, and it plummeted by 6.2 percentage points, to 14.9% in the second quarter of 2020. We take this decline in the expenditure share as reflected in the share parameter shock γ_t .

We calibrate intertemporal preference shock, ξ_t , to match the change in total expenditures from the fourth quarter of 2019 to the second quarter of 2020 by using the FIES, which stands at minus 8.5%. The value of ξ_t in the first quarter of the shock that generates a decline in consumption in the observed magnitude is 0.839.




Figure 10: Expenditure Share of Social Goods

Note: This graph shows the expenditure share of social goods. The data is from the Family Income and Expenditure Survey (FIES) by the Ministry of Internal Affairs and Communications (MIC).

Table 2 summarizes the shocks observed during the first quarter of the COVID-19 crisis. As we stand, we do not know how long the shocks will remain after the second quarter of 2020. In the next section, we simulate the transition under some scenarios about the duration of the shocks.

Table 2: The COVID-19 Shocks in 2020Q2

Parameter	Description	Values, source
<i>Wage Shocks</i>		
$\omega_{e,d,t}$	Wage shock	{1.000, 0.999, 0.990, 0.946}, MLS
<i>Employment Shocks</i>		
$\nu_{j,e,t}$	Employment-age specific shock	Figure 9, LFS
$\phi_{o,d,t}$	Industry-occupation specific shock	{1.003, 0.996, 0.990, 0.956}, LFS and ESS
<i>Preference Shocks</i>		
γ_t	Share parameter shock	6.2ppt, FIES
ξ_t	Intertemporal preference shock	0.839, FIES

5.3 Transition Dynamics and Welfare Analysis

As discussed in section 5.2, COVID-19 brought sizable shocks to the labor market but the effects are far from uniform across heterogeneous groups of individuals. We now simulate the transition of our model economy assuming that individuals in the initial economy are hit by the shocks at time 1 and make a transition back to normal times over time.

In this section, we first focus on effects of labor market shocks through employment and wage shocks, explained in section 5.2. In the next section, we will also add shocks to preferences to account for changes in consumption shares and levels observed in the data. Our main focus, however, is on effects of heterogeneous labor market shocks on individuals' welfare.

As discussed above, it is very difficult, if not entirely impossible, to conjecture how long the shocks will persist. We assume that the shocks are temporary and disappear eventually, but will last for multiple periods. In the computation, we let the shocks diminish at rate ρ each period, with expected duration of $1/\rho$.

In the baseline scenario, we assume that shocks last for one year (four quarters) in expectation and set $\rho = 0.25$. In section 5.4, we also consider more and less optimistic scenarios, in which shocks diminish more quickly with expected duration of two quarters, and more slowly over six quarters, respectively.

Given the size of initial shocks as summarized in Table 2, the average earnings exhibit a decline of 1.5% in the first quarter of the crisis, which gradually diminishes over the following quarters, as shown in Figure 11. Note that the decline takes into account changes in both employment and earnings of individuals.




Figure 11: Changes in Average Earnings Relative to the Initial Economy (%)

The shocks, however, do not hit individuals equally. Figure 12 shows heterogeneity in the magnitude of shocks by gender, education level, and employment type under the baseline scenario where expected duration of shocks is four quarters. They are expressed as a percentage change in earnings of each type of worker relative to the levels in the initial economy.

As shown in Figure 12a, females on average experience a 2.8% drop in earnings while the decline is less than 1% for males. Figures 12b and 12c show an even starker difference in the decline of earnings across employment types and education levels of workers. Contingent workers experience a drop of 6.5% for males and almost 8% for females, while that of regular workers is less than 1% for both genders. Individuals with less than a college degree experience a much sharper decline than those with a college degree. Note that we do not have any education-specific shock in the model and the difference comes from different compositions of workers within each group that are hit by the COVID-19 shocks.

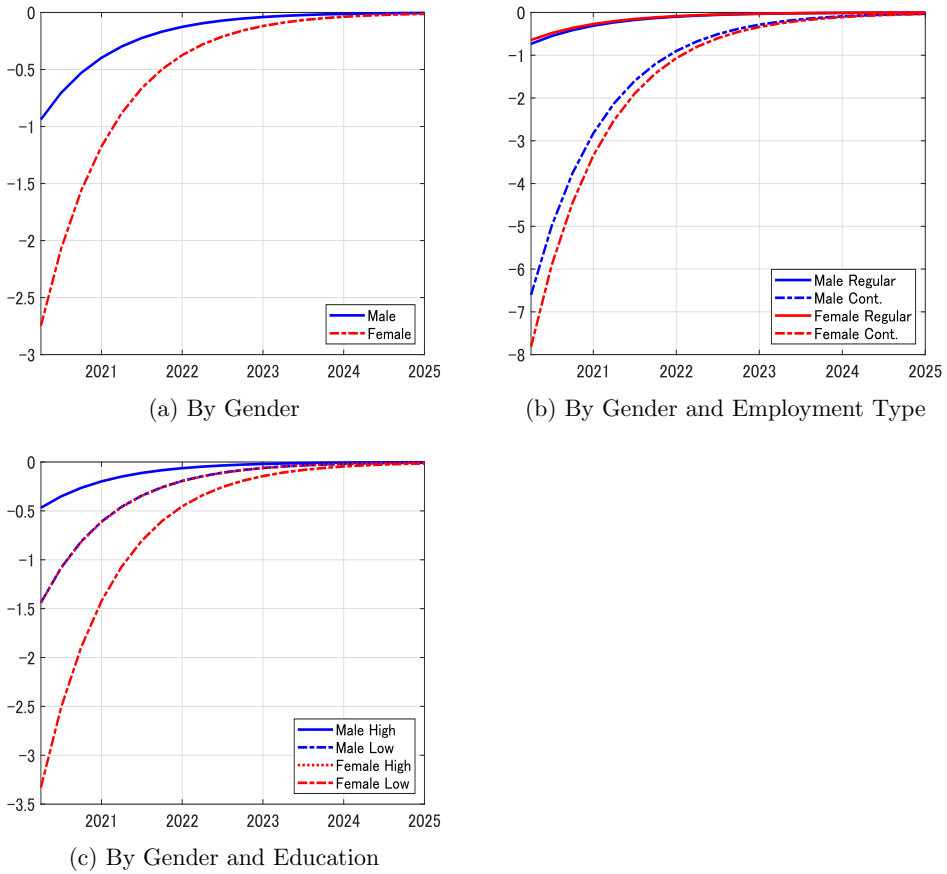


Figure 12: Changes in Average Earnings Relative to the Initial Economy (%)

We feed these shocks into our model in transition and compute welfare effects on different types of individuals. We use the initial economy as a basis of comparison and consider how individuals’ welfare changes once the COVID-19 shocks hit the economy and they live through the new paths of earnings.

More precisely, we compute welfare of individuals under the initial economy as well as welfare of all types of individuals in an economy that experiences the COVID-19 shocks at time 1, which corresponds to the second quarter of 2020. We then compute consumption equivalent variation, “*CEV*,” which equals a percentage change in consumption in the initial economy that would make an individual indifferent between living in the initial economy versus the economy facing COVID-19 shocks.

In order to account for difference in the expected duration of remaining life, which varies by individuals of different ages, we compute the present discounted value of consumption adjustment for the rest of an individual’s life, which we call “*PV-CEV*,” that

Covid Economics 40, 30 July 2020: 76-114

will be needed to make the individual indifferent.

Tables 3 and 4 show the *PV-CEV* of different groups of workers relative to average earnings of each group. Table 3 shows average welfare effects by gender, employment type and education level. Females on average face a welfare loss equivalent to 3.4% of their earnings, while the loss is more moderate at 1.1% for males. The table also shows a significant welfare loss for contingent workers, in a magnitude that corresponds to 7.5% and 9.5% of earnings for males and females, respectively.

Table 3: Welfare Effects by Gender, Employment Type and Education (aged 25-64, in PV-CEV)

	All	Emp. type		Education	
		Regular	Cont.	High	Low
All	-1.87	-0.92	-9.05	-0.90	-2.66
Male	-1.14	-0.91	-7.47	-0.64	-1.68
Female	-3.44	-0.94	-9.47	-1.68	-4.15

Table 4 shows welfare effects that differ across occupations and industries of individual workers. Workers in the social sector suffer significantly more from the COVID-19 crisis than those in the ordinary sector. The negative effect is much larger among those in non-flexible occupations, conditional on industry. Workers in the ordinary and flexible jobs experience a small loss of 0.16%, while those in the social and non-flexible jobs suffer from a large welfare loss of 6.8% relative to their earnings. Within each occupation and industry, females face a more significant welfare loss than males.

Table 4: Welfare Effects by Gender, Industry and Occupation (aged 25-64, in PV-CEV)

	Ordinary		Social	
	Flexible	Non-flex.	Flexible	Non-flex.
All	-0.16	-1.75	-1.82	-6.83
Male	+0.05	-1.44	-0.83	-5.16
Female	-0.85	-3.90	-2.90	-9.75

We now turn our attention to heterogeneity in welfare effects across age groups when the COVID-19 shocks hit the economy. Figure 13 plots the welfare effects by gender and age in 2020, in terms of *PV-CEV* in units of average earnings across all workers in the initial economy. On average, younger individuals suffer more from the COVID-19 shocks in the labor market than those approaching a retirement age, because the young must endure full length of shocks. Retirees are not affected directly by the wage shocks but their welfare declines slightly as we assume lump-sum transfers are adjusted to make up for a decline in tax revenues so the government can pay its social security expenditures.

In addition to the longer duration of the shocks that young individuals must suffer than the old, as we saw in Figure 9, employment of contingent workers is more severely hurt among the young, which adds to a larger welfare cost for them. The effects more clearly manifest among young female workers, whose share of contingent workers is much larger than males.

Besides the shape, the magnitude of the welfare costs is significantly larger for females, who are concentrated in the types of jobs that are more severely hit by the COVID-19 shocks.




Figure 13: Welfare Effects by Age and Gender (in PV-CEV)

Figure 14 shows welfare effects by other dimensions of heterogeneity across workers. As shown in Figure 14a, contingent workers suffer more from the shocks than regular workers and the difference is larger among younger workers who are hit harder by the employment type shocks, as discussed in section 5.2. Figure 14b demonstrates that the low-skilled workers suffer by more than the high-skilled workers.




Figure 14: Welfare Effects by Age, Employment Type and Education (in PV-CEV)

The analysis reveals the fact that negative effects of the COVID-19 crisis in the labor market have very different implications for people of different age, gender, employment type, education and job type in terms of industry and occupation. In each dimension, the shock is larger for those who earn less initially.

Our model captures heterogeneity across workers in many dimensions that turn out to be critical in evaluating welfare effects the COVID-19 crisis in Japan. There are, however, other dimensions that are not captured in our model. For example, our model assumes full insurance within each group and does not account for within-type heterogeneity in other dimensions such as wealth, health status, family structure, etc, which presumably may be important dimensions to analyze once a model is properly extended and calibrated to data that will eventually become available.

In the following section, we run a few additional experiments to consider alternative scenarios about duration of the COVID-19 shocks, and to introduce preference shocks to account for changes in consumption level and relative allocation across different types of goods. We will also consider welfare of some hypothetical households that consist of different types of individuals.

5.4 Sensitivity Analysis and Alternative Scenarios

5.4.1 Preference Shocks

We now consider shocks to preferences upon outbreak of the COVID-19 crisis. As summarized in section 5.2, there was a sizeable shift in the shares of consumption goods allocated to ordinary and social goods. The share of the latter was very stable at around 21% before the crisis and plummeted to less than 15% in the second quarter of 2020. At the same time, when we compare the level between the fourth quarter of 2019 and the

second quarter of 2020, we found the average consumption level also fell by 8.5%.¹⁴ We adjust preference parameters ξ_t and γ_t so that the model approximates these changes in consumption shares and average levels observed in the data. Similarly to the shocks to the labor market considered in section 5.3, we assume that the shocks will last for one year on average and diminish at rate $\rho = 0.25$.

Table 5 shows welfare effects from the transition incorporating preference shocks. With preference shocks, quantifying welfare effects of the COVID-19 becomes challenging since a new set of preference parameters directly affects welfare. Therefore, we compute welfare effects from different paths of consumption before and after the COVID-19 shocks, evaluated in terms of utility function in the initial economy. Although the level of welfare effects requires caution in interpretation, we confirm the same pattern of heterogeneous impact across different types of individuals, as shown in Table 5.¹⁵ Negative welfare effects are larger for females than males, contingent workers are hit harder than regular workers and so are the low-educated than the high-skilled.

Table 5: Welfare Effects with Preference Shocks (aged 25-64, in PV-CEV)

	All	Emp. type		Education	
		Regular	Cont.	High	Low
All	-0.43	+0.45	-7.09	+0.48	-1.17
Male	+0.23	+0.45	-5.90	+0.73	-0.31
Female	-1.85	+0.45	-7.40	-0.42	-2.49

¹⁴We approximate the effect of the COVID-19 shocks on the consumption level by a change between the fourth quarter of 2019 and the second quarter of 2020, rather than between the first and second quarters of 2020. We note some caution in quantifying the impact of COVID-19 on consumption from the time series data over this short time horizon before and after the crisis. Some decline in consumption had already begun in the latter half of the first quarter of 2020, in March in particular, and we avoid using this quarter as a basis of comparison. Also, there was a hike in the consumption tax rate from 8% to 10% in October 2019. The government implemented tax credits under some conditions for purchases until June 2020, in order to alleviate negative effects on consumption caused by the tax increase and to encourage more “cashless” transactions. Isolating pure effects of the COVID-19 crisis on consumption from these and other factors would be a non-trivial task. For these reasons, we use a quarterly change in consumption from 2019Q4 to 2020Q2 as approximating the COVID-19 shocks. Although the estimated change may vary under alternative assumptions, we think the main message from the welfare comparison across heterogeneous individuals presented in this section would remain intact.

¹⁵Although the focus of the analysis is a relative difference of welfare effects across different types of individuals, the levels of welfare effects also differ from those in the baseline without preference shocks since we are imposing the same pre-crisis preference in the computation. For example, shocks to the share parameter induce more consumption of ordinary goods, which carry more weight in the pre-crisis preference and raise the level of welfare effects, compared to the welfare effects evaluated without preference shocks. Other equilibrium effects also affect the magnitude of the welfare evaluated under the pre-crisis preference. We note, however, that since preferences are not type-specific, these effects do not affect our relative comparison of welfare across different types of individuals.

5.4.2 Duration of Shocks

In the baseline simulations, we assume that the COVID-19 shocks will diminish at rate $\rho = 0.25$ on a quarterly basis and last for 4 quarters in expectation. We consider two alternative scenarios in which shocks last for 2 and 6 quarters on average. Table 6 shows how welfare effects vary by duration of the shocks in the labor market. Not surprisingly, welfare loss is magnified when shocks last longer and exacerbate welfare loss of the vulnerable more. The table shows the difference across genders, but the pattern of heterogeneous welfare effects across other dimensions remains the same as in the baseline simulations presented above.¹⁶

Table 6: Welfare Effects and Shock Durations (aged 25-64, in PV-CEV)

Duration	Baseline		
	6 months	12 months	18 months
All	-0.94	-1.87	-2.78
Male	-0.58	-1.14	-1.71
Female	-1.74	-3.44	-5.11

5.4.3 Welfare Effects across Household Types

The unit of our analysis is an individual, and we do not explicitly consider a family structure in the baseline simulations. We observed a significant difference in the labor market experience across individuals by their characteristics. An especially large difference was observed between regular and contingent workers.

In this section, we simulate a model to infer how a household that consists of two earners of particular types may fare against other types of married households. We hypothetically construct earnings of a typical male and female individual engaged in a regular or contingent job. Four types of households that differ by gender and employment type of spouses are constructed. We then quantify welfare effects of the COVID-19 shocks on these four types of households and compare them.

Figure 15 shows the welfare effects married individuals in terms of *PV-CEV*, present discounted value of consumption equivalent variation, for each individual in a two-earner household of different combinations of spouses' employment type. Not surprisingly, members of two-earner households that consist of two contingent workers suffer the most. The negative effect of the COVID-19 is the smallest for married households with two regular workers.

¹⁶We do not show all the results under alternative duration assumptions due to a space constraint, but they are available from the authors upon request.




Figure 15: Welfare Effects of Married Individuals by Family Type (in PV-CEV)

6 Conclusion

In this paper, we document heterogeneous responses in employment and earnings to the COVID-19 shocks during the initial months after onset of the crisis in Japan. We then feed these changes in the labor market into a life-cycle model and evaluate welfare consequences of the COVID-19 shocks across heterogeneous individuals.

We find that negative effects of the COVID-19 shocks in the labor market significantly vary across people of different age group, gender, employment type, education level, industry and occupation. In each dimension, the shock is amplified for those who earn less prior to the crisis. Contingent workers are hit harder than regular workers, younger workers than older workers, females than males, workers engaged in social and non-flexible jobs than those in ordinary and flexible jobs. Our study identifies groups of individuals that are more severely hurt than others from the COVID-19 crisis, and suggests how the policy could be structured, which aims to reach the most vulnerable and the most severely affected.

Although the scope of the paper is to evaluate short-run impacts of COVID-19 in the labor market during the initial months of the crisis, there may well be other effects triggered by the crisis, such as structural changes in the labor market or in other dimensions of the economy over the medium and long-run. Such changes may also depend on how long various shocks we observe now will persist and whether they will be repeated multiple times. These topics which cover a longer time horizon are left for future research.

References

- Acemoglu, D., V. Chernozhukov, I. Werning, and M. D. Whinston (2020). A multi-risk SIR model with optimally targeted lockdown. NBER Working Paper No. 27102.

- Albanesi, S., R. Gihleb, J. Kim, and J. Huo (2020). Household insurance and the macroeconomic impact of the novel Corona virus. Working Paper.
- Alon, T. M., M. Doepke, J. Olmstead-Rumsey, and M. Tertilt (2020). The impact of COVID-19 on gender equality. NBER Working Paper No. 26947.
- Coibion, O., Y. Gorodnichenko, and M. Weber (2020). Labor markets during the COVID-19 crisis: A preliminary view. NBER Working Paper No. 27017.
- Dingel, J. I. and B. Neiman (2020). How many jobs can be done at home? NBER Working Paper No. 26948.
- Esteban-Pretel, J., R. Nakajima, and R. Tanaka (2011). Are contingent jobs dead ends or stepping stones to regular jobs? evidence from a structural estimation. *Labour Economics* 18, 513–526.
- Fukui, M., S. Kikuchi, and C. Goalist (2020). Job creation during the COVID-19 pandemic in Japan. Working Paper.
- Glover, A., J. Heathcote, D. Krueger, and J.-V. Rios-Rull (2020). Health versus wealth: On the distributional effects of controlling a pandemic. *Covid Economics* 6, 22–64.
- Gregory, V., G. Menzio, and D. Wiczer (2020). Pandemic recession: L- or V-shaped? *Quarterly Review* 4011.
- İmrohoroğlu, S., S. Kitao, and T. Yamada (2016). Achieving fiscal balance in Japan. *International Economic Review* 57(1), 117–154.
- İmrohoroğlu, S., S. Kitao, and T. Yamada (2019). Fiscal sustainability in Japan: what to tackle? *Journal of the Economics of Ageing* 14. Article 100205.
- Kahn, L., F. Lange, and D. Wiczer (2020). Labor supply in the time of COVID-19. Working Paper.
- Kaplan, G., B. Moll, and G. Violante (2020). Pandemics according to HANK. Working Paper, Princeton University.
- Kawaguchi, K., N. Kodama, and M. Tanaka (2020). Small business under the COVID-19 crisis: Expected short- and medium-run effects of anti-contagion and economic policies. Working Paper, Hong Kong University of Science and Technology.
- Kawata, K. (2020). Job seeker-vacancy mismatch with COVID-19. Working Paper.
- Kikuchi, S., S. Kitao, and M. Mikoshiba (2020). Heterogeneous vulnerability to the covid-19 crisis and implications for inequality in Japan. RIETI Discussion Paper, 20-E-039.
- Kitao, S. and M. Mikoshiba (2020). Females, the elderly and also males: Demographic aging and macroeconomy in Japan. *Journal of the Japanese and International Economics* 56. Article 101064.

- Miyakawa, D., K. Oikawa, and K. Ueda (2020). Firm exit during the covid-19 pandemic: Evidence from japan.
- Mongey, S., L. Pilossoph, and A. Weinberg (2020). Which workers bear the burden of social distancing policies? NBER Working Paper No. 27085.
- Okubo, T. (2020). Spread of COVID-19 and telework: Evidence from Japan. *Covid Economics* 32, 1–25.
- Watanabe, T. and Y. Omori (2020). Online consumption during the COVID-19 crisis: Evidence from Japan. *Covid Economics* 32, 208–241.
- Yokoyama, I., K. Higa, and D. Kawaguchi (2019). Adjustments of regular and non-regular workers to exogenous shocks: Evidence from exchange-rate fluctuation. *Industrial and Labor Relations Review*. forthcoming.

A Data Appendix

A.1 Labor Force Survey (LFS)

Sample: The Labor Force Survey (LFS) is a cross-sectional household survey conducted by the Ministry of Internal Affairs and Communications (MIC). The LFS is established to elucidate the current state of employment and unemployment in Japan. The survey was first conducted in July 1947. For our research propose, we use the monthly data, known as the “Basic Tabulation,” for the period from January 2013 to May 2020. The survey unit is a household residing in Japan, excluding foreign diplomatic and consular corps, their family members, and foreign military personal and their family members. For the “Basic Tabulation,” approximately 40 thousand households are selected. The questions on employment status are asked to only members aged 15 years or over. The LFS is conducted as of the last day of each month (except for December), and the employment status is surveyed for the week ending the last day of month.¹⁷

Definition of Variables: Employment status of the population aged 15 years and above is classified according to activity during the reference week. Our interest is the number of employed persons among the population aged 15 years and above. Employed persons consist of the employed at work and the employed not at work. Employed persons at work are defined as all persons who worked for (1) pay or profit, or (2) worked as unpaid family workers for at least one hour. Thus, we do not include people with jobs but not at work as employed at work. For example, those who did not work but received or were expected to receive wages or salary are classified as an employed person not at work.

¹⁷More detailed information can be found here: <https://www.stat.go.jp/english/data/roudou/pdf/1.pdf>

Employed people also consist of employees, self-employed worker, and family workers according to their main job. We use employees (those who work for wages or salaries) and classify them as regular or contingent (non-regular) based on what they are termed by their employers.

Industry classification follows the basis of the Japan Standard Industrial Classification (JSIC) according to the main types of business and industries of establishments, as revised in October 2013. We allocate industries into two sectors, which we call ordinary and social sectors.

Occupations are classified based on the Japan Standard Occupational Classification (JSOC), as revised in December 2009. We allocate them into two occupations, which we call flexible and non-flexible occupations.

Note that the samples of both industry and occupation are all workers aged 15 to 64, including not only employees (regular and contingent workers) but also other types of workers (self-employed worker and family workers), since more granular age and employment type categories cannot be obtained from publicly available aggregate data.

A.2 Monthly Labor Survey (MLS)

Sample: The Monthly Labor Survey (MLS) is a cross-sectional monthly survey conducted by the Ministry of Health, Labour and Welfare (MHLW). The MLS is established to measure changes in employment, earnings, and hours worked on both national and prefectural levels. The survey was first conducted in July 1923. For our research propose, we use the monthly national data for the period from January 2013 to May 2020. The MLS was conducted on approximately 33 thousands establishments, selected from all private and public sector establishments normally employing five or more regular employees and belonging to 16 categorized sectors. Surveys are conducted monthly and use values as of the end of each month.¹⁸

Definition of Variables: In this paper, we use the monthly data for contractual cash earnings of regular employees. The regular employees are defined as workers who satisfy condition (1) those who are employed for an indefinite period of time, or (2) those employed for a fixed term of one month or more. Then, the regular employees are classified as “full-time employees” and “part-time workers.” In section 5, we follow this definition as employment type. The part-time workers are those who satisfy condition (1) whose scheduled working hours per day are shorter than ordinary workers, or (2) whose scheduled working hours per day are the same as ordinary workers, but whose number of scheduled working days per week is fewer than ordinary workers.

¹⁸More detailed information can be found here: <https://www.mhlw.go.jp/english/database/db-slms/dl/slms-01.pdf>

The 16 industry categories follow the basis of the JSIC according to the main types of business and industry of establishments, as revised in October 2013. The 16 industry categories are a more granular categorization than that of the LFS. Then we similarly allocate industry into two categories, which we call ordinary and social by following the strategy taken in [Kaplan et al. \(2020\)](#).

- Ordinary Sector: mining and quarrying of stone and gravel, construction, manufacturing, electricity, gas, heat supply and water, information and communications, transport and postal service, wholesale, finance and insurance, real estate and goods rental and leasing.
- Social Sector: retail trade, scientific research, professional and technical services, accommodations, eating and drinking places, living related and personal services and amusement service, education, learning support, medical, health care and welfare, compound services, services, N.E.C.

We use contractual cash earnings as earnings in this paper. Cash earnings are the amount before deducting taxes, social insurance premiums, trade union dues or purchase price, etc. Contractual cash earnings are defined as earnings paid according to a method and conditions previously determined by labor contract, collective agreement, or wage regulations of establishments. The contractual cash earnings consist of scheduled cash earnings and non-scheduled cash earnings, which are overtime pay. Overtime pay is the wages paid for work performed outside scheduled working hours, such as at night and in the early morning. Note that contractual cash earnings include a salary paid without actual labor, such as leave pay.

A.3 Employment Status Survey (ESS)

Sample: The Employment Status Survey (ESS) is a cross-sectional household survey conducted by the Ministry of Internal Affairs and Communications (MIC). The ESS aims to obtain basic data on actual conditions of the employment structure at both national and regional levels by surveying the usual labor force status in Japan. The ESS was conducted every three years between 1956 and 1982, and has been conducted every five years since 1982. For our research propose, we use the latest data collected in October 2017. The survey unit is a household of members aged 15 years and above residing in Japan except for (1) foreign diplomatic corps or consular staff (including their suite and their family members), (2) foreign military personnel or civilians (including their family members), (3) persons dwelling in camps or ships of the Self-Defense Forces, (4) persons serving sentences in prisons or detention houses, and (5) inmates of reformatory

institutions or women's guidance homes. Approximately 490 thousand households living in sampled units are selected.¹⁹

Definition of Variables: To obtain the distribution of employees with various characteristics, we use the “order-made” data and focus on employees aged 20 and over. For characteristics of employees, we follow the information about age, gender, education, employment type, sector, occupation, and income.

Age is counted as of September 30, 2017. In this paper, we use data for the 10-year age groups: 30s, 40s and 50s. Education status is defined according to the information on the survey date. In this paper, we allocate education status into two types, which we call high and low. We define employees as high-skilled if they have a college or higher degree, and low skilled otherwise.

In this paper, we focus on employees and classify them into two types of employment: regular and contingent. The regular employment type includes executives of companies or corporations and regular staff who are termed “regular employees.” The contingent employment type includes part-time workers, albeit (temporary workers), dispatched workers from a temporary labor agency, contract employees, entrusted employees, and others.

Industry classification follows the basis of the JSIC for the main types of business and industries of establishments, as revised in October 2013. We allocate industries into two sectors, which we call ordinary and social sectors.

Occupations are classified based on the JSOC, as revised in December 2009. We allocate them into two groups, which we call flexible and non-flexible occupations.

Income is defined as the sum of annual income from October 2016 to September 2017 that workers earn from their main jobs excluding non-monetary income. Note that the income of those who changed their jobs or took up a new job during the past year is calculated based on income from the day when they start a new job up to the reference day assuming that they keep working for a year. The income of employees is gross earnings inclusive of tax gained during the past year from wages, salaries, charges for labor, various allowances, bonuses, and the like. Incomes are grouped into 17 categories: less than 50, 50-99, 100-149, 150-199, 200-249, 250-299, 300-399, 400-499, 500-599, 600-699, 700-799, 800-899, 900-999, 1000-1249, 1250-1499, over 1500 (in 10 thousand yen). When we calculate average income, we use the middle value of income categories for all categories but the smallest and largest groups. For the group with less than 50, we use 25, and for the group with over 1500, we use 1500.

¹⁹More detailed information can be found here: <https://www.stat.go.jp/english/data/shugyou/2017/outline.html>

A.4 Family Income and Expenditure Survey (FIES)

Sample: The Family Income and Expenditure Survey (FIES) is a cross-sectional household survey conducted by the Ministry of Internal Affairs and Communications (MIC). The survey was first conducted in September 1950. For our research propose, we use the “Monthly Report on the Family Income and Expenditure Survey” of two-or-more-person households (multiple-person households) for the period from January 2013 to May 2020. The survey unit is a household residing in Japan, except for (1) one-person student households, (2) inpatients in hospitals, inmates of reformatory institutions, etc., (3) households which manage restaurants, hotels, boarding houses, or dormitories, sharing their dwellings, (4) households which serve meals to boarders even though not managing boarding houses as an occupation, (5) households with 4 or more live-in employees, (6) households whose heads are absent for a long time (three months or more), (7) foreigner households. The entire land of Japan is stratified into 168 strata. Approximately 8,000 multiple-person households and 750 one-person households are surveyed every month from the strata. Multiple-person households are surveyed for six consecutive months, while one-person households are surveyed for three consecutive months, but only after 2002.²⁰

Definition of Variables: In this paper, we use monthly multiple-person household’s income and expenditure data. We allocate commodities into two types from two different sectors, which we call ordinary and social sectors, and closely follow the strategy taken in Kaplan et al. (2020).

- **Ordinary Sector:** cereals, fish and shellfish, meat, daily products and eggs, vegetables and seaweeds, fruits, oils, fats and seasonings, cakes and candies, cooked food, beverages, rents for dwelling and land, tools and materials for repairs and maintenance, fuel, light and water changes, durable goods assisting housework, heating and cooling appliances, interior furnishings and decorations, interior furnishings and decorations, bedding, domestic utensils, domestic non-durable goods, Japanese clothing, clothing, shirts and sweaters, underwear, cloth and thread, other clothing, footwear, medicines, health fortification, medical supplies and appliances, purchase of vehicles, purchase of bicycles, maintenance of vehicles, communications, school textbooks and reference books for study, culture and recreation, personal care goods, personal effects, tobacco, other miscellaneous, pocket money, money gifts, remittance, other social expenses.
- **Social Sector:** meals out side the home, service charges for repairs and maintenance, domestic service, services related to clothing, medical service, public transportation, school fees, tutorial fees, recreational service, personal care services.

²⁰More detailed information can be found here: <https://www.stat.go.jp/english/data/kakei/1560.html>

B Calibration of Shocks

Seasonal Adjustment and Conversion of Frequency: As discussed in appendix A, we use the monthly labor and consumption data to calculate the shocks, which we feed into the model. The frequency of our model, however, is quarterly, and we use changes between the first quarter and the second quarter of 2020 as the COVID-19 shocks. For the purpose of the calibration in section 5.2, we convert monthly data into quarterly data and seasonally adjust it by using X12 ARIMA.²¹

Occupation-sector specific shocks: The occupation-sector specific shock $\phi_{o,d,t}$ is one of the two employment shocks and this shock hits workers of each combination of occupation and sector $(o, d) = (1,1), (1,2), (2,1), (2,2)$, independently of the other individual characteristics.

We first compute changes in employment between the first and the second quarters of 2020 for each combination. Note that the LFS's aggregate data only provide changes in employment of "all" type- (o, d) workers and do not represent pure (o, d) shocks associated with occupation and industry.²² If, for example, social and non-flexible workers are disproportionately contingent, their employment may decline sharply, not because of the (o, d) shock, but because of the employment-type shock. Thus, we use the employment type shock $\nu_{j,e}$ by the LFS and, the distribution $\mu_{j,e|o,d}$ over employment type and age, conditionally on (o, d) . Note $\sum_{j,e} \mu_{j,e|o,d} = 1$. Denoting the employment changes of all type- (o, d) workers as $x_{f,d}$, we calculate the occupation-sector specific shocks $\phi_{o,d}$ so that they satisfy

$$x_{o,d} = \sum_{j,e} \mu_{j,e|o,d} (1 - \nu_{j,e}) \phi_{o,d}$$

for each combination of (o, d) .

C Computation Algorithm

This appendix describes computation of equilibrium of our model. First, we compute an equilibrium of the initial economy and second, the transition from the initial economy to the final economy. The final economy is assumed to be the same as the initial economy and effects of the shocks disappear in the long-run. The transition dynamics are computed in the following three steps. We assume that the transition takes T periods, which is long enough so that the economy converges to the final economy smoothly.

²¹We use the R package "x12". <https://cran.r-project.org/web/packages/x12/x12.pdf>

²²Note that the samples of both industry and occupation are all workers aged 15 to 64, including not only employees (regular and contingent workers) but also other types of workers such as the self-employed, since more granular age and employment type categories cannot be obtained from publicly available aggregate data.

1. Guess the paths of two equilibrium objects, $\{\tau_{ls,t}, b_t\}_{t=1}^T$; lump-sum taxes and bequests.
2. Solve individuals' problems. See below for details.
3. Check if the government budget constraint is satisfied. If not, adjust $\tau_{ls,t}$. Check if assets of the deceased equal accidental bequests. If not, adjust b_t . Continue until the conditions are satisfied for all $t = 1, \dots, T$.

The equilibrium of the initial economy is computed in similar steps, with only one time period and by setting $T = 1$.

Individuals' Life-cycle Problem: We now describe individuals' life-cycle problem and details of step 2 above. Recall the utility function

$$U(c_{1,t}, c_{2,t}) = \xi_t \frac{[c_{1,t}^{\gamma_t} c_{2,t}^{1-\gamma_t}]^{1-\sigma}}{1-\sigma} \tag{4}$$

where $c_{1,t}$ and $c_{2,t}$ denotes an individual's consumption of ordinary and social goods by individual at time t . Recall also the budget constraint

$$(1 + \tau_{c,t})(c_{1,t} + c_{2,t}) + a_{t+1} = y_{x,t} + R_t(a_t + b_t) + \tau_{ls,t} \tag{5}$$

where $y_{x,t}$ denotes after-tax earnings of an individual of a working age in state x or pension benefits in case of a retiree.

From an intratemporal condition

$$c_{2,t} = \frac{1 - \gamma_t}{\gamma_t} c_{1,t} \equiv \Lambda_t c_{1,t} \tag{6}$$

where

$$\Lambda_t \equiv \frac{1 - \gamma_t}{\gamma_t}.$$

Plug (6) in (4),

$$U(c_{1,t}, c_{2,t}) = \xi_t \frac{[c_{1,t}^{\gamma_t} (\Lambda_t c_{1,t})^{1-\gamma_t}]^{1-\sigma}}{1-\sigma} = \Omega_t \frac{c_{1,t}^{1-\sigma}}{1-\sigma} \equiv u(c_{1,t}) \tag{7}$$

where

$$\Omega_t \equiv \xi_t \Lambda_t^{(1-\gamma_t)(1-\sigma)}$$

Now consider an intertemporal decision of individuals. Plug (6) in (5),

$$(1 + \tau_c) \frac{1}{\gamma_t} c_{1,t} + a_{t+1} = y_{x,t} + R_t(a_t + b_t) + \tau_{ls,t} \tag{8}$$

Rewrite an individual's life-cycle problem in terms of $c_{1,t}$ as

$$\max \sum_{j=1}^J \beta^{j-1} \left(\prod_{k=1}^j s_k \right) u(c_{1,j,t})$$

where $u(c_{1,t})$ is defined as in (7) subject to (8).

From the Euler equation

$$\frac{c_{1,t+1}}{c_{1,t}} = \left(\beta s_{j+1} R_{t+1} \frac{\Omega_{t+1} \gamma_{t+1}}{\Omega_t \gamma_t} \right)^{\frac{1}{\sigma}} \equiv g_{1,t+1}^c$$

where $g_{1,t+1}^c$ denotes gross growth rate of consumption of goods 1 between time t and $t+1$.

Consumption of goods 2 is given as (6), and we have

$$\frac{c_{2,t+1}}{c_{2,t}} = \frac{\Lambda_{t+1} c_{1,t+1}}{\Lambda_t c_{1,t}} = \frac{\Lambda_{t+1}}{\Lambda_t} g_{1,t+1}^c \equiv g_{2,t+1}^c$$

Consumption of goods 1 and goods 2 of an individual aged j born in time t is

$$c_{1,t+j-1} = c_{1,t} \prod_{k=1}^j g_{1,t+k-1}^c \tag{9}$$

$$c_{2,t+j-1} = c_{2,t} \prod_{k=1}^j g_{2,t+k-1}^c \tag{10}$$

where $g_{1,t}^c = g_{2,t}^c = 1$.

Present discounted values of expenditures for consumption goods 1 and 2, $C_{1,t}$ and $C_{2,t}$, for an individual born at time t , are given as

$$\begin{aligned} C_{1,t} &= c_{1,t} + \sum_{j=2}^J \left(\prod_{k=2}^j \frac{s_k}{R_{t+k-1}} \right) c_{1,t+j-1} \\ &= c_{1,t} \left[1 + \sum_{j=2}^J \left(\prod_{k=2}^j \frac{s_k}{R_{t+k-1}} \right) \left(\prod_{k=1}^j g_{1,t+k-1}^c \right) \right] \\ C_{2,t} &= c_{2,t} + \sum_{j=2}^J \left(\prod_{k=2}^j \frac{s_k}{R_{t+k-1}} \right) c_{2,t+j-1} \\ &= c_{2,t} \left[1 + \sum_{j=2}^J \left(\prod_{k=2}^j \frac{s_k}{R_{t+k-1}} \right) \left(\prod_{k=1}^j g_{2,t+k-1}^c \right) \right] \\ &= c_{1,t} \frac{1-\gamma_t}{\gamma_t} \left[1 + \sum_{j=2}^J \left(\prod_{k=2}^j \frac{s_k}{R_{t+k-1}} \right) \left(\prod_{k=1}^j g_{2,t+k-1}^c \right) \right] \end{aligned}$$

Define $\tilde{y}_{x,t}$ as total income given as

$$\tilde{y}_{x,t} = y_{x,t} + R_t b_t + \tau_{l,t}$$

Present discounted value of income is given as

$$Y_t = \tilde{y}_{1,t} + \sum_{j=2}^J \left(\prod_{k=2}^j \frac{s_k}{R_{t+k-1}} \right) \tilde{y}_{j,t+j-1}$$

Since

$$(1 + \tau_c) (C_{1,t} + C_{2,t}) = Y_t,$$

$c_{1,t}$ is computed as

$$c_{1,t} = \frac{Y_t / (1 + \tau_c)}{\left[1 + \sum_{j=2}^J \left(\prod_{k=2}^j \frac{s_k}{R_{t+k-1}} \right) \left(\prod_{k=1}^j g_{1,t+k-1}^c \right) \right] + \frac{1-\gamma_t}{\gamma_t} \left[1 + \sum_{j=2}^J \left(\prod_{k=2}^j \frac{s_k}{R_{t+k-1}} \right) \left(\prod_{k=1}^j g_{2,t+k-1}^c \right) \right]}$$

Then compute $c_{1,t}$ and $c_{2,t}$ using (6), (9) and (10). Finally, compute assets from (5) recursively.

Optimal social distancing in SIR-based macroeconomic models¹

Yoseph Y. Getachew²

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

The paper introduces voluntary social distancing to the canonical epidemiology model, integrated into a conventional macroeconomic model. The model is extended to include treatment, vaccination, and government-enforced lockdown. Infection-averse individuals face a trade-off between a costly social distancing and the risk of getting infected and losing next-period labor income. We find an individual's social distancing is proportional to the welfare loss she incurs when moving to the infected compartment. It increases in the individual's psychological discount factor but decreases in the probability of receiving a vaccination. Quantitatively, a laissez-faire social distancing flattens the infection curve that minimizes the economic damage of the epidemic. A government-enforced social distancing is more effective in flattening the infection curve but has a detrimental effect on the economy.

1 We are grateful for the comments from the participants of the Mini Online Economics Conference (26 May 2020), University of Hull, particularly to Parantap Basu and Keshab Bhattacharai. The usual disclaimer applies.

2 Associate professor, University of Pretoria.

Copyright: Yoseph Y. Getachew

1. Introduction

At the time this paper is written, more than 15.4 million individuals are infected by COVID-19 worldwide and more than 631 thousand died while the spread of the epidemic shows no sign of slowing down.¹ However, many countries have already relaxed the strict lockdown measures that they implemented at the early stage of the epidemic to ease the pressure on the economy.² Controlling the spread of the infection has thus been mainly left for choices made at individual levels. Unfortunately, the canonical epidemiology models that are often adopted to track the spread of the COVID-19 epidemic do not have the necessary tool to account for individual behaviors despite some of the variables in these models largely depend on how individuals behave in the presence of the epidemic.

The present paper aims to contribute to fill this gap. It complements the recent macroeconomic literature that gives microfoundations to the canonical epidemiology models used to track the spreads of COVID-19 and assesses the spread of the outbreak and its macroeconomic impact. The paper in particular develops a SIR (Susceptible-Infected-Recovered) macroeconomic model where individuals face a trade-off between practicing a costly social distancing and increasing the risk of getting infected and losing next-period labor income. Their optimal decisions eventually determine the dynamics of the epidemic, aggregate income, and welfare. The model is further ex-

¹COVID-19 was first reported in Wuhan, China, in December 2019.

²When the South Africa government had relaxed a five-week-long strict lockdown measure (including a ban on jogging, cycling, and dog-walking) in the 1st of May 2020, the number of confirmed COVID-19 cases in the country was less than six thousand, after ten weeks the country has breached the 300 thousand mark of COVID-19 cases.

tended to include treatment, vaccination, and government-enforced social distancing.

The basic model considers an economy that faces an epidemic where individuals are categorized into three compartments – Susceptible, Infected and Recovered; hence the name SIR. Initially, there are two types of individuals – susceptible and infected – while in the following periods some of the infected persons get recovered. Similar to conventional economic dynamic models, agents derive utility from consumption and leisure. In contrast, they derive utility from social closeness such as hugging, kissing, and shaking hands of their loved ones although this could expose susceptible individuals to the virus. Infected individuals work less time, due to sickness, and hence lose labor income. They could also die from the infection. However, if they get recovered they resume a normal life. Infected and recovered individuals practice the minimum social distancing, which is zero. While the latter develop immunity, the former have nothing to lose.

We thus model social distancing as costly to individuals but it does not involve consumption goods or time, which is in sharp contrast to the recent literature on macroeconomics and epidemics (e.g., Eichenbaum et al., 2020, Krueger et al., 2020). In particular, we incorporate social distancing into an otherwise standard utility function where individuals derive utility from consumption and leisure. In this context, individuals could optimally decide on social distancing along with the labor-leisure trade-off. Our rationale for providing microfoundations to SIR models does not rest upon individuals' consumption, work, or leisure activities. Instead, we assume susceptible individuals are infection-averse when it comes to social distancing. Accordingly, two different individuals may choose a similar bundle of consumption

goods or leisure time but may experience different social distancing.

We modify the basic SIR model with laissez-faire social distancing to include treatment and vaccination (could be any other similar controlling mechanisms). In the SIVTR model (V and T stand for Vaccination and Treatment), individuals are categorized into five compartments every period: Susceptible, Infected, Vaccinated, Treated, and Recovered. Treatment is believed to decrease the infectivity of the epidemic as it often involves the identification and quarantining of infected individuals and increases the recovery rate of infected individuals. Vaccination or other controlling practices such as wearing masks, education, or washing hands, could significantly reduce susceptibility to infection as it reduces the number of susceptible individuals. We also extend the original SIR model to accommodate a government-enforced social distancing.

In the SLIR (Susceptible-Lockdown-Infected-Recovered) model, a fraction of susceptible individuals will leave the susceptible compartment starting from the initial period of the epidemic, which leads to a staggering job loss in the economy. We think of the latter as those individuals who work in industries such as hotels and tourism whose employment status is severely affected by the lockdown measure. The government may subsidize the resulting unemployment through lump-sum taxes, levied on the general population. But the government's revenue could quickly dwindle, as the infection soars and consequently people die, call in sick and their ability to pay taxes decreases. To cope up with the pressure of a declining economy, the government is allowed to relax the lockdown measure through time, which has been an empirical regularity. In such a scenario, with lump-sum taxes, we establish a second-best

condition by equating the lifetime utility of individuals in lockdown to that of the recovered persons.

Among the findings, we show that a susceptible individual's current optimal social distancing is the difference between her value function (welfare) of having remained susceptible and being infected in the next period. It increases in her psychological discount factor but decreases at the probability of receiving a vaccination or her likelihood of developing immunity. Quantitatively, aggregate income, consumption, and welfare increase in laissez-faire social distancing but they decrease in government-enforced social distancing. The latter is more effective in flattening the curve but leads to a higher unemployment rate. If available, treatment and vaccination could affect aggregate welfare positively but in different ways. The availability of treatment has a positive influence on all individuals' welfare, including that of the susceptible individuals as it increases their likelihood of getting treatment (if they get infected) and hence getting recovered quickly. Whereas, the availability of vaccination pulls individuals out of the susceptible compartment from the outset and enables them to avoid costly social distancing.

Calibrating the model to the U.S. economy, a laissez-faire social distancing is found to have a strong impact on delaying and flattening the infection curve. It delays the peak period by about 20 more periods, flattens the curve at the peak by about 10 percentage pts, and decreases the death rate by more than 3 percentage pts from the baseline case of no social distancing. The decrease in the infection and fatality rates translates to a positive impact on the economy due to the boost in the labor supply. Aggregate income and consumption increase, which in turn leads to an

increase in the aggregate welfare. During the early periods, when the infection and fatality rates are small, there is no much difference in the macroeconomic variables between practicing and not practicing laissez-faire social distancing. This would change quickly once the epidemic gains momentum, more people get infected and hence lose their labor income due to sickness and death. At the peak of the epidemic, there is a 20 percentage pts difference in aggregate income between practicing and not practicing social distancing and there is a permanent 5 percentage pts difference after herd immunity is achieved.

If treatment is available, it will have a relatively modest impact on the evolution of the outbreak. The infection curve flattens by about 1.84 percentage pts and the death rate decreases by 1.12 percentage pts, from the baseline case of no treatment is available. Although the impacts on aggregate income and consumption are relatively small, the impact on aggregate welfare could be quite important due to its positive impact on the lifetime welfare of all individuals, including susceptible individuals. However, if vaccination is available, it would have a much stronger influence on the dynamics of the epidemic and consequently on the economy. A 0.03% vaccination rate per period (similar to vaccinating 0.1 million susceptible individuals per week) cuts down the death rate by a 3 percentage pts and flattens the infection curve by about 4.5 percentage pts from the baseline case of no vaccination is available. Increasing the vaccination rate to 0.06 per period (similar to vaccinating 0.2 million susceptible individuals per week) could lead to herd immunity after only 2% of the population gets infected. Without vaccination, 73% of the population has to be infected to achieve herd immunity. With the 0.03% vaccination rate per period,

aggregate income increases by 16 percentage pts, at the peak of the infection, from the baseline case of no vaccination is available.

The numerical simulation shows that, compared to a laissez-faire policy of do nothing, a government-enforced social distancing has a much stronger impact on the spread of the epidemic. With the latter, at the peak of the epidemic, only 0.74% of the susceptible individuals get infected and 2.43% of them die; however, with the laissez-faire social distancing, 11.43% of the susceptible individuals get infected and 17% of them die. The lockdown has a strong negative impact on the macroeconomy, due to losses in aggregate labor income, however. With our parametrization of the lockdown for the U.S. economy, aggregate income and welfare reduce by 46 and 75 percentage pts, respectively. When comparing between a more and a less strict lockdown, a more strict lockdown leads to a relatively higher welfare loss at the early stage of the epidemic; however, during and after the peak of the infection, it leads to a relatively lower welfare loss as some of the welfare loss are offset by the live-savings effects of the lockdown.

The work contributes to the very recent debate in the macroeconomic impact of the epidemic. In the last couple of months, many works have appeared that combine macroeconomic models with SIR models in response to the COVID-19 crisis. An incomplete list of these works includes Acemoglu et al. (2020), Alvarez et al. (2020), Atkeson (2020), Bodenstein et al. (2020), Chang and Velasco (2020), Eichenbaum et al. (2020), Farboodi et al. (2020), Garibaldi et al. (2020), Glover et al. (2020), Greenstone and Nigam (2020), Jones et al. (2020), Krueger et al. (2020), Toxvaerd (2020) and Fernandez-Vallaverde and Jones (2020).

Atkeson (2020) provides an early summary of SIR models from the perspective of macroeconomics. Jones et al. (2020) compare a social planner's mitigating incentives with that of private agents. They argue the planner's mitigation policy (that encourages working from home) could be much more effective in reducing the death rate despite that results in a significant drop in consumption. Acemoglu et al. (2020) focus on optimal targeted lockdown policy in a multi-group SIR model. While infection reduces in a strict and long lockdown of the most vulnerable group (the oldest group), this also enables to impose a lesser lockdown in the lower-risk group (the young). Bodenstein et al. (2020) look into the impact of public health measures such as social distancing or lockdown on the death rate in a model that combines a multi-sectoral model with the SIR model. Glover et al. (2020) examine the distributional impact of optimal mitigation policy across different groups, categorized in terms of age, sector, and health status.³

The current work is more closely related to the work of Eichenbaum et al. (2020), Farboodi et al. (2020), Krueger et al. (2020), and Toxvaerd (2020). Eichenbaum et al. (2020) and Krueger et al. (2020) attach individuals' consumption and labor activities to the contact rate that increases their likelihood of getting infected. Thus, a consumption tax could be considered as a containment policy. We share with them in our modeling approach to the extent that we introduce the SIR model to conventional macroeconomic models through the contact rate. We share with

³There are also other many recent macroeconomic works in the COVID-19 that abstract from the SIR model (e.g., Baker et al., 2020 and Barrot et al., 2020). Baker et al. (2020), for instance, examine how household spending responds to the COVID-19. Barrot et al. (2020) look at the impact of the weeks' long lockdown in France's and other European countries' output while focusing on the sectoral effects.

Farboodi et al. (2020), Basu et al. (2020), and Toxvaerd (2020) that agents in these models derive utility from social activity. Similar to them we focus on individual behaviors towards optimal social distancing, in contrast, we approach the problem from a macroeconomic point of view.

We organize the next sections as follows. Section 2 introduces the basic SIR model to the household problem. Section 3 models treatment and vaccination. In section 4, we introduce and examine a government-enforced social distancing. We calibrate the models in Section 5. Section 6 provides the numerical results and Section 7 concludes.

2. The SIR Model

We suppose an economy that faces an epidemic. There are in general three types of individuals, namely susceptible, infected, and recovered individuals, as in the standard epidemiological SIR models. For the susceptible individuals, the probability to remain susceptible in the next period is $1 - p_t$, where p_t is the probability of getting infected. For the infected individuals, the probability to remain infected in the next period is $1 - \gamma - \iota$, where γ and ι are the probability of recovering and dying, respectively.

As in the standard representative household models, agents derive utility from consumption and leisure. In contrast to these models, susceptible individuals dislike social distancing. They derive utility from social closeness (such as kissing, hugging friends & relatives and shaking hands) although it increases the likelihood of getting infected, and thence, being off work and losing some of their earnings, and risk of

dying in the next period. Infected and recovered individuals practice the minimum amount of social distancing, which is zero. While the latter develop immunity, the former has nothing to lose.

2.1. Household

At time $t = 0$ there are two types of individuals – susceptible and infected individuals. In the following periods, some of the infected individuals get recovered. Recovered individuals develop immunity to the virus and resume normal life. Denote infected individuals as 0, susceptible individuals as 1 and recovered individuals as 2. The problem for the infected and susceptible individuals can be represented as two state process $(\pi_{i,t})$, where i takes 0 or 1. At time t , an infected person is represented, by $\pi_{1,t} = 1$, and a non-infected person, by $\pi_{0,t} = 0$. The utility of the i th person is then given by:

$$U(C_{i,t}, L_{i,t}, \chi_{i,t}) = \frac{C_{i,t}^{1-\sigma} - 1}{1 - \sigma} - \frac{L_{i,t}^{1+\theta}}{1 + \theta} - \pi_{i,t} \frac{\chi_{i,t}^2}{2} \tag{1}$$

where $U'_c > 0$, $U''_c < 0$, $U'_l < 0$, $U''_l < 0$, $U'_\chi < 0$ and $U''_\chi < 0$.

The budget constraint is given by:

$$C_{i,t} = w_t ((1 - \pi_{i,t})(L_{0,t} - L_s) + \pi_{i,t}L_{1,t}) - M_t \tag{2}$$

where $C_{i,t}$ and $L_{i,t}$ are the i th individual consumption and leisure; w_t and M_t are the wage rate and lump-sum tax respectively. $\chi_{i,t}$ represents social distancing by a susceptible individual and L_s is work time lost due to sickness absence by an infected person. From (1), infected individuals ($\pi_{0,t} = 0$) do not practice social distancing,

and from (2), they do not work full time, $L_s \neq 0$.

The utility function for recovered individuals is given by,

$$U(C_{2,t}, L_{2,t}) = \frac{C_{2,t}^{1-\sigma} - 1}{1-\sigma} - \frac{L_{2,t}^{1+\theta}}{1+\theta} \quad (3)$$

subject to the budget constraint:

$$C_{2,t} = w_t L_{2,t} - M_t \quad (4)$$

The individual's consumption is simply her wage income minus lump-sum tax.

2.2. SIR

The *transmission risk* p_t is the probability of a susceptible individual encountering an infected individual and thence getting infected. We suppose it decreases in the individual's level of social distancing $\chi_{1,t}$ and has the following simple form:

$$p_t = 1 - a\chi_{1,t} \quad (5)$$

where $p_t \in [0, 1]$ and $a \in (0, 1]$ is a parameter.

In the typical SIR model, the total number of individuals infected at time t is the number of encounters between infected ($N_{0,t}$) and susceptible ($N_{1,t}$) individuals times the contact rate (β).

$$\beta N_{0,t} N_{1,t}$$

Our strategy of providing microfoundations to the SIR model is modifying the contact

rate to account for a voluntary social distancing as follows:

$$\beta_t N_{0,t} N_{1,t} \tag{6}$$

where $\beta_t \equiv \beta p_t$ is the *effective* contact rate. The probability of recovering of an infected individual is γ and the total number of recovered individuals at time t is a fraction of infected people, $\gamma N_{0,t}$.⁴

In the SIR model, we have the following relations:

$$N_{1,t+1} = N_{1,t} - \beta_t N_{0,t} N_{1,t} \tag{7}$$

$$N_{0,t+1} = N_{0,t} + \beta_t N_{0,t} N_{1,t} - (\gamma + \iota) N_{0,t} \tag{8}$$

$$N_{2,t+1} = N_{2,t} + \gamma N_{0,t} \tag{9}$$

where $N_{1,t}$, $N_{0,t}$ and $N_{2,t}$ represent the number of susceptible, infected and recovered individuals in the economy, respectively. Eqs. (7)-(9) show the dynamics of susceptible, infected and recovered individuals. From (7)-(8), we see every period $\beta_t N_{0,t} N_{1,t}$ number of individuals leaves the susceptible compartment and joins the infected compartment in the next period. Similarly, from (8)-(9), $(\gamma + \iota) N_{0,t}$ number of individuals leaves the infected compartment, out of which $\gamma N_{0,t}$ number of individuals joins the recovered compartment and $\iota N_{0,t}$ number of individuals dies every period.

$\gamma + \iota$ is the removal rate and $1/(\gamma + \iota)$ is the mean periods that an infected

⁴Note that if we do not account for voluntary social distancing, $\chi_{1,t} = 0$, then $\beta_t = \beta$. We make the assumption, β and ι are constant, for simplicity.

individual remains in the infected compartment, leading to the basic reproduction number in the SIR model:

$$\mathbf{R}_0 = \frac{\beta_0}{\gamma + \iota}$$

An infected individual should at least transmit to more than one individual ($\mathbf{R}_0 > 1$) for the infection to have a first phase of an upward dynamics.

The size of the population at time t (N_t) is the total number of susceptible, infected and recovered individuals. At time $t + 1$, this is equal to the population size at t net of infected individuals who died from the infection.

$$N_t = N_{1,t} + N_{0,t} + N_{2,t} \quad (10)$$

$$N_{t+1} = N_t - \iota N_{0,t} \quad (11)$$

$$D_{t+1} = D_t + \iota D_{0,t} \quad (12)$$

The last equation captures the dynamics for the death rate where D_t is the number of dead people at time t . We set initial population to be one ($N_0 = 1$) and assume zero population growth rate and zero *natural* death rate, with no loss of generality.

2.3. The Households' Problem

The lifetime problem of the agent who is susceptible at time t , recursively, is

$$V_1 = \max_{\{C_{1,t}, L_{1,t}, \chi_{1,t}\}} U(C_{1,t}, L_{1,t}, \chi_{1,t}) + (1 - p_t) \rho V'_1 + p_t \rho V'_0 \quad (13)$$

and that of the person who is infected at time t , is

$$V_0 = \max_{C_{0,t}, L_{0,t}} U(C_{0,t}, L_{0,t}) + (1 - \gamma - \iota) \rho V'_0 + \gamma \rho V'_2 \quad (14)$$

subject to (1) and (2). ρ is the discount rate; and, " V " indicates the next period value function.

Similarly, the problem of a recovered individual is to maximize

$$V_2 = \max_{C_{2,t}, L_{2,t}} U(C_{2,t}, L_{2,t}) + \rho V'_2 \quad (15)$$

subject to her utility function and budget constraint. Implicit in (15), recovered individuals develop immunity to the infection and resume normal life.

2.4. Solution to the Household Problem

From the first order conditions of the susceptible individual and the budget constraint,

$$w_t C_{1,t}^{-\sigma} = L_{1,t}^\theta \quad (16)$$

$$C_{1,t} = w_t L_{1,t} - M_t \quad (17)$$

$$\chi_{1,t} = \rho a (V'_1 - V'_0) \quad (18)$$

The first is the trade-off between the individual's consumption and leisure and the second is her budget constraint. The last equation captures the individual's optimal social distancing, which depends on the individual's discount rate ρ , and the welfare

loss she incurs when moving from the susceptible to the infected compartment.

Proposition 1. (i) *A susceptible individual's optimal social distancing is proportional to the welfare loss she incurs if she moves from the susceptible to the infected compartment. (ii) It increases in the individual's discount factor.*

The solution for the infected individual is,

$$w_t C_{0,t}^{-\sigma} = L_{0,t}^\theta \quad (19)$$

$$C_{0,t} = w_t (L_{0,t} - L_s) - M_t \quad (20)$$

and the solution for the recovered individual is

$$w_t C_{2,t}^{-\sigma} = L_{2,t}^\theta \quad (21)$$

$$C_{2,t} = w_t L_{2,t} - M_t \quad (22)$$

(19) and (21) show the labor-leisure trade-off for the infected and recovered individuals while (20) and (22) show their respective budget constraints. Infected individuals have the lowest individual consumption (20) due to time lost from sickness absentees.

2.5. Aggregate output and labour

Aggregate output is produced using aggregate labor:

$$Y_t = AL_t \quad (23)$$

where A is total factor productivity (TFP). Aggregate labor at time t is given by,

$$L_t = N_{0,t}(L_{0,t} - L_s) + N_{1,t}L_{1,t} + N_{2,t}L_{2,t}$$

which is the sum of labor supply by the infected, susceptible and recovered individuals in the economy at time t .

3. Treatment and Vaccination

We modify the SIR model to include treatment and vaccination (or any other controlling mechanisms that help to reduce the spread of the epidemic by removing some individuals from the susceptible compartment). When treatment is available, the infectivity of the epidemic is believed to decrease and the recovery rate to rise. The reduction in infectivity could happen as treatment often requires certain identification and quarantining of infected individuals. Vaccination or any other controlling practices such as wearing masks, education, or washing hands, could significantly reduce susceptibility to infection as it reduces the number of susceptible individuals.

3.1. SIVTR

In the SIVTR model, individuals can be categorized into five compartments: Susceptible, Infected, Vaccinated, Treated, and Recovered. We model treatment by letting ω number of infected individuals to receive treatment every period that decreases the infectivity by ε rate. We suppose ψ number of treated individuals leave the treatment room (or recover) and $1/\psi > 1/\gamma$, that is, the recovery period of individuals receiving treatment is shorter than that of individuals who do not

receive treatment. We model vaccination letting v fraction of susceptible individuals to be vaccinated every period. For simplicity, we assume that the vaccination or the control measure implemented will completely eliminate susceptibility to infection.

Then, following the approach of Feng et al. (2011), the SIVTR model could have the following form:

$$N_{1,t+1} = N_{1,t} - \beta_t (N_{0,t} + \varepsilon N_{T,t}) N_{1,t} - v N_{V,t} \quad (24)$$

$$N_{0,t+1} = N_{0,t} + (1 - \omega) \beta_t (N_{0,t} + \varepsilon N_{T,t}) N_{1,t} - (\gamma + \iota) N_{0,t} \quad (25)$$

$$N_{V,t+1} = N_{V,t} + v N_{V,t} \quad (26)$$

$$N_{T,t+1} = N_{T,t} + \omega \beta_t (N_{0,t} + \varepsilon N_{T,t}) N_{1,t} - \psi N_{T,t} \quad (27)$$

$$N_{2,t+1} = N_{2,t} + \gamma N_{0,t} + \psi N_{T,t} \quad (28)$$

where $N_{T,t}$ and $N_{V,t}$ denote the number of individuals treated and vaccinated at period t respectively.

Eqs. (24)-(28) show the dynamics for susceptible, infected, vaccinated, treated and recovered individuals. From (24), every period, $\beta_t (N_{0,t} + \varepsilon N_{T,t}) N_{1,t}$ individuals leave the susceptible compartment and $1 - \omega$ of these individuals join the infected compartment (25) while the rest join the treatment compartment (27). The term $\varepsilon \beta_t N_{T,t} N_{1,t}$ captures the encounter of susceptible and treated individuals, which decreases infectivity by $\varepsilon \in (0, 1)$ rate. From the treatment compartment, $\psi N_{T,t}$ individuals leave the treatment room every period and join the recovered compartment (28).

As shown in (24), $v N_{V,t}$ number of susceptible individuals leave the susceptible

compartment every period and enter the vaccinated compartment (26). Eq. (28) presents the dynamics for the recovered individuals, those who leave the infection and treatment compartments.

The basic reproduction number for the SIVTR model is

$$R_0 = (1 - v) \beta_0 \left(\frac{1 - \omega}{\gamma + \iota} + \frac{\omega \varepsilon}{\psi} \right) \tag{29}$$

The first term in the big bracket is the average number of periods that an infected individual spends in the infected compartment; the second is the fraction of infected individuals who receive treatment. $1/\psi$ is the average time an infected individual stays in the treatment compartment and it decreases by ε rate.

With the availability of vaccination, the lifetime utility of susceptible individuals would change. A susceptible individual receives vaccination with probability v and with the assumption that vaccination will eliminate susceptibility to infection, the lifetime utility of the person changes as follows:

$$V_1 = \max_{\{C_{1,t}, L_{1,t}, \chi_{1,t}\}} U(C_{1,t}, L_{1,t}, \chi_{1,t}) + (1 - v) [(1 - p_t) \rho V'_1 + p_t \rho V'_0] + v \rho V'_2 \tag{30}$$

Her optimal social distancing considers her likelihood of receiving vaccination and is summarized in the following Proposition:

Proposition 2. *A susceptible individual optimal social distancing,*

$$\chi_{1,t} = (1 - v) \rho a (V'_1 - V'_0) \tag{31}$$

will reduce at the rate of the availability of a vaccine, v .

The value function for a vaccinated individual is similar to that of a recovered individual as both develop immunity and hence practice the minimum social distancing, which is zero. There is no change to the lifetime utility of infected and recovered individuals. With the availability of treatment, the lifetime utility of infected individuals would change though:

$$V_0 = \max_{C_{0,t}, L_{0,t}} U(C_{0,t}, L_{0,t}) + \rho [(1 - \gamma - \omega - \iota) V'_0 + \gamma V'_2 + \omega V'_T] \quad (32)$$

Infected individuals get a treatment with a probability of ω , get recovered with a probability of γ , remain sick and do not receive treatment, or die from the infection with a probability of ι , in the next period.

The lifetime utility of treated individual is

$$V_T = \max_{C_{T,t}, L_{T,t}} U(C_{T,t}, L_{T,t}) + \rho [(1 - \psi - \iota) V'_T + \psi V'_2] \quad (33)$$

For simplicity, we suppose there is no difference between an infected and treated individual in terms of labor supply and death rate. The only difference between the two is that individuals in the treatment compartment have a relatively higher recovery rate.

Aggregate labor supply at time t changes from the SIR model as it includes now treated and vaccinated individuals:

$$L_t = N_{T,t}(L_{T,t} - L_s) + N_{I,t}(L_{0,t} - L_s) + N_{1,t}L_{1,t} + N_{2,t}L_{2,t} + N_{V,t}L_{V,t} \quad (34)$$

where $L_{T,t}$ and $L_{V,t}$ are labor supply by treated and vaccinated individuals. One may note that the number of working time is similar for an infected and treated individual. Also, there is no difference in terms of labor supply between a vaccinated and a recovered person.

4. Lockdown

In this section, we suppose a government lockdown during the epidemic period that makes α_t number of susceptible individuals ($\alpha_t N_{1,t}$) unemployed. One may think of these individuals as those who work in industries (such as hotels and tourism) that are severely affected by the lockdown. Given that the main purpose of a lockdown is to cut down the number of susceptible individuals, we consider those individuals who are out of work also to be out of the susceptible compartment.

The government subsidizes the resulting unemployment through lump-sum taxes,

$$N_t M = \alpha_t N_{1,t} C_L \quad (35)$$

where M and C_L denote the lump-sum tax and the consumption of the individual who is affected by the lockdown respectively. $N_t M$ is the aggregate tax revenue, which will be used to subsidize the consumption of $\alpha_t N_{1,t}$ unemployed individuals. Note that initially, at $t = 0$, almost everyone is susceptible thus $N_{1,t} \approx N_t$. But, later on, because more and more people die from the infection, the size of the total population N_t will decline, resulting in declining government revenue. To hold a balanced budget, the government needs to relax the lockdown at the rate that keeps individual consumption constant. Considering that, we have from (35):

$$\alpha_t = \alpha \frac{N_t}{N_{1,t}} \quad (36)$$

α is the initial lockdown rate when $N_0 \approx N_{1,0} = 1$. Substituting (36) into the above, we get the consumption of an individual who loses her labor income due to government-enforced social distancing: $C_L = \frac{M}{\alpha}$. Because the individual is neither susceptible nor employed, $L_u = \chi_u = 0$ and her utility function is given by

$$U(C_L) = \frac{C_L^{1-\sigma} - 1}{1-\sigma} \quad (37)$$

Second best condition can be obtained by equating the lifetime utility of this individual to that of a recovered person:

$$V_2 = V_L = U(C_L) + \rho V'_L \quad (38)$$

where V_L is the lifetime utility of the individual who loses her job due to the lockdown. Combining (15) and (38),

$$U(C_L) + \rho V'_L = U(C_2, L_2) + \rho V'_2 \quad (39)$$

which equates the lifetime utility of a recovered person to that of an unemployed person. The sufficient condition for (39) to be satisfied is

$$U(C_L) = U(C_2, L_2) \quad (40)$$

One easily solves the level of lump-sum tax M associated to a given α , after substituting (3), (4) and (37) into (40):

$$M^* = \alpha \left((1 - \sigma) \left(\frac{C_2^{1-\sigma} - 1}{1 - \sigma} - \frac{L_2^{1+\theta}}{1 + \theta} \right) + 1 \right)^{\frac{1}{1-\sigma}} \quad (41)$$

where M^* is the optimal lump-sum tax that each working individual pays, and C_2 and L_2 are given by (21) and (22) respectively.

5. Calibration

We calibrate the baseline model for the COVID-19 and the U.S. economy. A period is a week as in Eichenbaum et al. (2020). We let $\theta = 1$; estimates for the Frisch elasticity of labor supply θ often range between 0.5 and 2. We set $\sigma = 2$ for the curvature of the utility function and $\rho = 0.96^{1/52}$ for the weekly discount rate. We compute $A = 24$, using a \$50,000 per year income target and 40 weekly work hours. $L_s = 0.1$ that implies a 10% less consumption for individuals who do not work full time, a regularity in the incomplete market literature. For the SIR model, we assume $M = 0$.

Values for the COVID-19 parameters are largely varied between estimates and quickly change. We mainly rely on data from the Center for Disease Control and Prevention (CDC).⁵ The national U.S. infection fatality rate among people infected with the COVID-19 is about 1.3%. We suppose a 18 days recovery time for infected individuals that implies $\frac{1}{\gamma+\iota} = 18/7$ removal weeks in our model. This gives $\gamma =$

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

0.32. An average initial reproduction number of $\mathbf{R}_0 = 2.2$ implies a contact rate of $\beta = 0.86$. Initial population size is $P = 330 \times 10^6$ which is standardized to one ($N_0 = 1$). We start with 50 infected individuals, $N_{0,0} = 50/P$ and zero recovered and death rate, $D_0 = N_{2,0} = 0$.

For the SIVTR model, we set ω and ε at 5.5%, which is the total percentage of all COVID-19 cases that are hospitalized.⁶ The average length of hospital stay ranges from 8 days (with no ICU), 10 days (with ICU and without ventilators) to 16 days (with ICU and ventilators), which is about 11 days or 1.57 weeks on average. This implies $\psi = 1 \div 1.57 = 63\%$. Apparently, there is no value for vaccination ν thus we start with some small number such as $\nu = 1/3300$, which is equivalent to vaccinating 100 thousand people weekly, and then experiment on the level of a vaccination rate that is required to achieve herd immunity at a very small infection rate.

We calibrate α , the fraction of initial susceptible individuals that leave the susceptible compartment due to government-enforced social distancing, based on the resulting unemployment. We then calibrate the lump-sum taxes M corresponding to these values from (41). The U.S. went on lockdown in March 2020 to prevent the further spread of the epidemic. According to the U.S. Bureau of Labor Statistics, the number of Americans drawing unemployment benefit at the end of May 2020 was 20.9 million people. This is equivalent to 0.5% of susceptible individuals leaving the susceptible compartment per period, in our model.⁷ We experiment between $\alpha = 0.001$ and 0.005. The table below lists the full calibrated values.

⁶<https://www.cdc.gov/coronavirus/2019-ncov/hcp/COVIDSurge.html>

⁷Dividing 21 million by the total number of susceptible individuals, 330 million, gives 0.06 and dividing that by 12 gives about 0.005.

Table 1: Baseline values

Preference	$\sigma = 2; \theta = 1; \rho = 0.96^{1/52}; a = 1$
Technology and policy	$A = 24; L_s = 0.1; M = 0$
SIR	$\beta = .86; \gamma = 0.32; \iota = 0.013$
SIVTR	$\psi = 0.63; \varepsilon = \omega = 5.5\%; \nu = 1/3300$
Lockdown	$\alpha = 0.001, \alpha = 0.005$
Baseline population	$P = 330 \times 10^6; N_0 = 1; N_{V,0} = N_{T,0} = N_{2,0} = D_0 = 0$
	$N_{0,0} = 50/P; N_{1,0} = 1 - N_{0,0}$

6. Results and Discussion

6.1. Baseline SIR

We start by examining the epidemiological SIR model (Figure 1). Figure 1a depicts the dynamics of susceptible, infected, and recovered individuals. Figure 1b is similar to Figure 1a except that it includes the population and death dynamics. Initially, almost all individuals are susceptible.⁸ It takes a while for the epidemic to build momentum as shown in the curve for the susceptible individuals, which is

⁸Only 50 individuals out of 330 million are infected at $t = 0$.

almost flat for the first thirty periods. During these periods, the number of infected and recovered individuals is close to zero. But once the number of infected individuals starts to rise, the number of susceptible individuals will decline sharply. And, the number of recovered individuals will rise quickly because as more and more people get infected, more and more people get recovered. At the peak of the epidemic, more than 21% of susceptible individuals get infected. Herd immunity could be achieved after 88% of susceptible individuals are infected. And, the death toll from the infection could pass more than 20% of the population.⁹

Figures 1c and 1d show the dynamics for aggregate consumption and labor supply that are largely determined by the dynamics of the outbreak. During the early stages of the epidemic, labor is mainly supplied by susceptible individuals, as there are only a few infected and recovered persons. As the number of susceptible individuals decreases, following the increase in the infection rate, labor supply also decreases, which in turn leads to a decline in consumption. The macroeconomic variables start to stabilize once herd immunity is achieved or the epidemic dies out.

6.2. *Laissez Faire Social Distancing*

Figure 2 compares the economic and epidemiological impacts of a laissez-faire social distancing to the baseline case of no social distancing. The former has a significant impact on the epidemics, through delaying and flattening the infection curve (Figures 2a and 2b). While it takes about 20 more periods to reach the peak

⁹The dynamics of the population and the fatality of the infection behave similarly but conversely. During the early stages of the outbreak, the latter is almost zero. The associated curve starts to incline later on, following the increase in the infection rate, eventually, it stabilizes as the epidemic dies out.

with voluntary social distancing, the infection rate decreases by about 10 percentage pts at its peak. Herd immunity could be achieved with a 15% lesser infection rate. And, the death rate declines by more than 3 percentage pts from the baseline.

The effect on the economy is positive. Aggregate income, consumption, and welfare increase compared to the baseline (Figures 2c-2f). The decrease in the infection and death rates increase aggregate labor supply, which in turn increases aggregate consumption and welfare. As shown in Figure 2e, the difference between the macroeconomic variables with and without social distancing follows the path of the infection curve. At the early stage of the epidemics, when the infection rate is too low, it is close to zero. However, at the peak of the epidemic, aggregate income is higher by more than 24 percentage pts compared to the baseline. The gap then decreases as the infection rate slows down while it remains constant once herd immunity has been achieved. The latter represents the long term macroeconomic effect of voluntary social distancing, which is the result of the decline in the fatality rate.

6.3. Treatment and Vaccination

As shown in Figure 3, treatment has a relatively smaller effect on the spread of the outbreak, particularly when compared to other controlling measures.¹⁰ It flattens the infection curve by only 1.84 percentage pts while it decreases the death rate by 1.12 percentage pts (Figures 3a and 3b). The high recovery rate also implies that herd immunity could be achieved relatively quickly.

How does that translate to the economy? The impact on aggregate income and

¹⁰In this and the next sections, laissez-faire social distancing is assumed in all of the numerical simulations.

consumption is modest as shown in Figures 3c-3e, which entirely depends on the impacts of treatment on the infection and death rates. At the early stage of the epidemic, there is no difference between aggregate income and consumption, with or without treatment. But at the peak of the infection, aggregate income is higher by 6 percentage pts from the baseline case of no treatment (Figure 3e).

The impact of treatment on welfare is more important for two reasons (Figure 3f). First, treated individuals have relatively higher welfare because of their high recovery rate. Second, the lifetime welfare of susceptible individuals is higher with the availability of treatment. Because, if they get infected, they could receive treatment and quickly recovered. The same works for infected individuals, they are better off with the prospect of receiving treatment in the future.

Even a small vaccination rate greatly influences the dynamics of the outbreak. Figure 4 demonstrates the effects of a 0.03% vaccination rate per period on the epidemics, vis-à-vis the baseline case of no vaccination. It decreases the death rate by 3 percentage pts and flattens the infection curve by about 4.5 percentage pts (Figures 4a and 4b). Herd immunity could be achieved at a much lesser infection rate (by 13 percentage pts) than the baseline. Increasing the vaccination rate to 0.06 per period, which is equivalent to vaccinating 200 thousand individuals per period, will have a tremendous impact on the outbreak. Herd immunity will be achieved after only 2% of the population gets infected.

Figures 4c-4f show the macroeconomic effects of vaccination. As shown in Figure 4e, the impact on aggregate income and consumption mainly follows that of the impact on the infection curve. Aggregate income increases by 16 percentage pts at

the peak of the infection, and by about 5 percentage pts permanently after herd immunity is achieved, compared to the case of no vaccination. However, its influence in aggregate welfare rather starts from the outset (Figure 4f). The intuition is that vaccinated individuals develop herd immunity, leave the susceptible compartment, and hence do not incur any more disutility from social distancing.

6.4. Government-Enforced Social Distancing

The numerical simulation shows that government-enforced social distancing could be among the most effective controlling mechanisms of the spread of the outbreak. Figure 5 compares a *laissez-faire* (do nothing) policy with two different lockdown levels – when $\alpha = 0.001$ and $\alpha = 0.005$. Figure 5a captures a quite interesting dynamics of susceptible individuals under lockdown. During the early stages, more individuals leave the susceptible compartment being in lockdown than being infected. During the latter stages, however, more people leave the susceptible compartment being infected than being in lockdown. The latter roughly matches the U.S. unemployment data. The rates of infection at the peak of the epidemic are 0.74%, 8.95%, and 11.43% for the lockdown levels of $\alpha = 0.005$, $\alpha = 0.001$, and do nothing (Figure 5b). The respective death rates are 2.43%, 17%, and 14.75% (Figure 5c).

The lockdown has a strong negative impact on the macroeconomy, however (Figures 5d-5f). Despite saving lives, it leads to job loss and thence a loss in aggregate labor income. From Figures 5d and 5e, aggregate income and consumption decrease. A 0.005 lockdown rate reduces aggregate income by 46 percentage pts while a 0.001 lockdown rate reduces it by about 10 percentage pts, at the end of the simulation periods (Figure 5f).

The dynamics of aggregate welfare is different from that of aggregate income and consumption as it includes the consumption of individuals under lockdown whose income comes from government transfer (Figure 5g). Figure 5h shows the percentage loss in aggregate welfare due to lockdown could go up to more than 75% pts, depending on the level of the lockdown. However, a more strict lockdown doesn't always mean a bigger welfare loss. Particularly, at later stage, a more strict lockdown could mean a lower welfare loss, as some of the negative job-loss effects are offset by the positive life-saving effects.

7. Final Remark

The paper provided an alternative framework of the SIR (Susceptible-Infected-Recovered) epidemiology model integrated into the standard economic dynamic model through a voluntary social distancing. The rationale for providing microfoundations to the SIR models does not rest upon individuals' consumption, work, or leisure activities but on infection-averse individuals who have a taste for non-pecuniary social closeness. In addition to leisure and consumption, individuals care for social closeness (e.g., hugging, kissing, socializing) although these could cost their life or income. Accordingly, two different individuals may choose a similar bundle of consumption goods or leisure time but may experience different social distancing. In their leisure choice, one person may go out to a beach for three hours and the other may stay at home watching *The Wolf of Wall Street* for the same amount of time, for instance.

Susceptible individuals face a trade-off between practicing social distancing and increasing their likelihood of being infected and thence losing labor income in the

following periods. Infected individuals work less time, due to sickness, and hence lose some labor income while recovered individuals resume a normal life. Infected and recovered individuals do not practice social distancing. While the latter develop immunity, the former have nothing to lose. Optimal individual-level social distancing determines the dynamics of the epidemic, which in turn determine the dynamics of the macroeconomic variables.

An individual's optimal social distancing is the difference between her value function of remaining in the susceptible compartment and moving to the infected compartment in the next period. It increases in her psychological discount factor but decreases at the probability of receiving a vaccination or her likelihood of developing immunity. From the numerical simulation, a laissez-faire social distancing is important in terms of delaying and flattening the infection curve that minimizes the economic damage from the outbreak. A government-enforced social distancing or a lockdown is highly effective in flattening the infection curve. But it would have a detrimental effect on the economy, through a negative job-loss effect. Treatment and vaccination positively influence aggregate welfare but through different mechanisms. The former increases aggregate welfare by increasing individuals' likelihood of getting recovered quickly. The latter pulls out individuals of the susceptible compartment from the outset and enables them to avoid a costly social distancing.

The paper is part of the primary efforts to provide microfoundations to the canonical epidemiology model, used to track the spread of the recent outbreak, and to integrate it into conventional macroeconomic models. A simple approach was adopted to deal with the problem in a tractable manner, without loss of generality. The qual-

itative results, however, should be read as illustrative and caution should be taken while interpreting the results from the numerical simulations. A strong quantitative prediction of the course of the epidemic could be obtained through adopting a more elaborated version of the model, which considers the different stages of the outbreak such as asymptomatic and symptomatic cases, different severity of illness (non-life-threatening cases and cases that require ICU admission). The work can also be extended to include more detailed non-pharmaceutical interventions (e.g., a ban on gathering, stay at home orders and the closures of industries and school), and the financing of the health sector.

References

- ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. D. WHINSTON (2020): “Optimal Targeted Lockdowns in a Multi-Group SIR Model,” *NBER Working Paper*, 27102.
- ALVAREZ, F. E., D. ARGENTE, AND F. LIPPI (2020): “A simple planning problem for covid-19 lockdown,” *Covid Economics*, 14, 6–28.
- ATKESON, A. (2020): “What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios,” Discussion paper, National Bureau of Economic Research.
- BAKER, S. R., R. A. FARROKHNIYA, S. MEYER, M. PAGEL, AND C. YANNELIS (2020): “How does household spending respond to an epidemic? consumption during the 2020 covid-19 pandemic,” *Covid Economics*, 18, 78–108.
- BARROT, J.-N., B. GRASSI, AND J. SAUVAGNAT (2020): “Sectoral effects of social distancing,” *Available at SSRN*.
- BASU, P., C. BELL, AND H. EDWARDS (2020): “Modelling social distancing and infection rate in a Covid-19 Economy,” *Powerpoint presentation, LSE*, 31.
- BODENSTEIN, M., G. CORSETTI, AND L. GUERRIERI (2020): “Social distancing and supply disruptions in a pandemic,” *Covid Economics*, 19, 1–53.
- CHANG, R., AND A. VELASCO (2020): “Economic policy incentives to preserve lives and livelihoods,” Discussion paper, National Bureau of Economic Research.

- EICHENBAUM, M. S., S. REBELO, AND M. TRABANDT (2020): “The macroeconomics of epidemics,” Discussion paper, National Bureau of Economic Research.
- FARBOODI, M., G. JAROSCH, AND R. SHIMER (2020): “Internal and external effects of social distancing in a pandemic,” *Covid Economics*, 9, 22–58.
- FENG, Z., S. TOWERS, AND Y. YANG (2011): “Modeling the effects of vaccination and treatment on pandemic influenza,” *The AAPS journal*, 13(3), 427–437.
- FERNÁNDEZ-VILLAVERDE, J., AND C. I. JONES (2020): “Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities,” Discussion paper, National Bureau of Economic Research.
- GARIBALDI, P., E. R. MOEN, AND C. A. PISSARIDES (2020): “Modelling contacts and transitions in the SIR epidemics model,” *Covid Economics Vetted and Real-Time Papers, CEPR*.
- GLOVER, A., J. HEATHCOTE, D. KRUEGER, AND J.-V. RÍOS-RULL (2020): “Health versus wealth: On the distributional effects of controlling a pandemic,” Discussion paper, National Bureau of Economic Research.
- GREENSTONE, M., AND V. NIGAM (2020): “Does social distancing matter?,” *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2020-26).
- JONES, C. J., T. PHILIPPON, AND V. VENKATESWARAN (2020): “Optimal mitigation policies in a pandemic: Social distancing and working from home,” Discussion paper, National Bureau of Economic Research.

KRUEGER, D., H. UHLIG, AND T. XIE (2020): “Macroeconomic dynamics and reallocation in an epidemic,” *Covid Economics*, 5, 22–55.

TOXVAERD, F. (2020): “Equilibrium social distancing,” *Covid Economics*, 15, 110–133.

Figures

Figure 1a: Baseline SIR for susceptible, infected and recovered persons




Figure 1b: Baseline SIR including death and population dynamic




Figure 1c: Baseline SIR for the dynamics of aggregate consumption




Figure 1d: Baseline SIR for the dynamics of aggregate labour supply




Figure 2a: Infection dynamics with laissez-faire social distancing




Figure 2b: Death dynamics with laissez-faire social distancing




Figure 2c: Aggregate consumption dynamics and laissez-faire social distancing




Figure 2d: Aggregate Labour dynamics and laissez-faire social distancing




Figure 2e: Percentage differences in aggregate consumption and income, with and without optimal social distancing





Figure 2f: Aggregate welfare dynamics and laissez-faire social distancing



Covid Economics 40, 30 July 2020: 115-163

Figure 3a: Infection dynamics with and without treatment




Figure 3b: Death dynamics with and without treatment




Figure 3c: Aggregate consumption with and without treatment




Figure 3d: Aggregate labour with and without treatment




Figure 3e: Percentage differences in consumption and income with and without treatment





Figure 3f: Aggregate welfare dynamics with and without treatment



Covid Economics 40, 30 July 2020: 115-163

Figure 4a: Infection dynamics with and without vaccination





Figure 4b: Death dynamics with and without treatment



Covid Economics 40, 30 July 2020: 115-163

Figure 4c: Aggregate consumption with and without vaccination




Figure 4d: Aggregate labour with and without vaccination




Figure 4e: Percentage differences in consumption and income, with and without vaccination





Figure 4f: Aggregate welfare dynamics with and without vaccination



Covid Economics 40, 30 July 2020: 115-163

Figure 5a: Dynamic of susceptible individuals with government-enforced social distancing




Figure 5b: Infection dynamics with government-enforced social distancing




Figure 5c: Death dynamics with government-enforced social distancing




Figure 5d: Aggregate consumption dynamics with government-enforced social distancing




Figure 5e: Aggregate labour dynamic with government-enforced social distancing




Figure 5f: Percentage differences in consumption and income, with and without government-enforced social distancing




Figure 5g: Aggregate welfare with government-enforced social distancing





Figure 5h: Percentage differences in aggregate welfare with and without government-enforced social distancing



V -, U -, L - or W-shaped economic recovery after COVID-19? Insights from an agent-based model¹

Dhruv Sharma,² Jean-Philippe Bouchaud,³ Stanislao Gualdi,⁴ Marco Tarzia,⁵ Francesco Zamponi⁶

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

We discuss the impact of a Covid-like shock on a simple toy economy, described by the Mark-o Agent-Based Model that we developed and discussed in a series of previous papers. We consider a mixed supply and demand shock, and show that depending on the shock parameters (amplitude and duration), our toy economy can display V-shaped, U-shaped or W-shaped recoveries, and even an L-shaped output curve with permanent output loss. This is due to the existence of a self-sustained "bad" state of the economy. We then discuss two policies that attempt to moderate the impact of the shock: giving easy credit to firms, and the so-called helicopter money, i.e. injecting new money into the households savings. We find that both policies are effective if strong enough, and we highlight the potential danger of terminating these policies too early. Interestingly, when policy is successful, inflation post-crisis is significantly increased. While we only discuss a limited number of scenarios, our model is flexible and versatile enough to allow for a much wider exploration, thus serving as a useful tool for the qualitative understanding of post-Covid recovery.

¹ We provide an on-line version of the code at <https://gitlab.com/sharma.dhruv/markovid>.

² Laboratoire de Physique de l'Ecole Normale Supérieure, ENS, Université PSL, CNRS; Chair of Econophysics & Complex Systems, Ecole polytechnique.

³ Capital Fund Management; Chair of Econophysics & Complex Systems, Ecole polytechnique.

⁴ Chair of Econophysics & Complex Systems, Ecole polytechnique.

⁵ LPTMC, CNRS-UMR 7600, Sorbonne Université; Institut Universitaire de France.

⁶ Laboratoire de Physique de l'Ecole Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université.

Copyright: Dhruv Sharma, Jean-Philippe Bouchaud, Stanislao Gualdi, Marco Tarzia, Francesco Zamponi

1 Introduction

The coronavirus pandemic has buffeted the world economy and induced one of the most abrupt drops in output ever recorded. What comes next? Will the economy recover quickly as lock-down measures are lifted, or will the damage inflicted by the massive waves of layoffs be more permanent? In pictorial terms, will the economic crisis be V-shaped, as commentators were initially hoping for, or U-shaped (prolonged drop followed by a quick recovery), or perhaps W-shaped, with a relapse due either to a second outburst of the illness, or to a premature lifting of the economic support to households and firms? The possibility of an L-shaped crisis, with a permanent loss of output, is also discussed. Or else, maybe, a “swoosh”, with a rapid drop followed by an excruciatingly slow recovery? [1]

There has been a flurry of activity to understand the consequences of the economic shock due to widespread lock-downs and loss of economic activity. While some have coupled classical economic models with SIR-like epidemic models, with the underlying assumption that the economy is somehow slaved to the dynamics of COVID [2], others have reasoned in terms of traditional economic models. There has been analytical support for both quick (V or U shaped) recoveries [3] and prolonged (L-shaped) crisis due to a stagnation trap (poor economic forecasts leading to lower consumption leading to lower investment) [4]. Given how different sectors of the economy are effected disproportionately (some completely shut, some are not), there are fears of deep recession due to a Keynesian supply shock - a deep demand shock greater in magnitude to the supply shock that cause them [5].

In this short note, we want to explore how the economic system by itself can recover from such a rapid drop of both supply and demand, even assuming quick return to normal in terms of sanitary measures. We perform numerical experiments using a prototype Agent Based Model that we have studied in depth in the past, in the context of monetary policy and inflation targeting.

Within our (highly simplified) model, we find that the length and severity of the crisis (its “typographical shape”) can be strongly affected by policy measures. We argue that, as was done in most European states, generous policies that avoid (as much as possible) bankruptcies and redundancies, allow the economy to recover rapidly, although endogenous relapses are possible (i.e. a W-shape without a second lock-down period).

In our model, U-shaped or L-shaped recoveries occur when the economy falls into what we called a “bad phase” in [6], characterized by a self-consistently sustained state of economic depression and deflation. The time needed for the “good phase” of the economy to re-establish itself when the shock is over can be extremely long (so long that it might exceed the simulation time). As a function of the parameters describing the crisis (amplitude and duration of the shock), we find that there is a discontinuous transition between V-shape recoveries and L-shape recessions. The main message of our numerical experiments is that policy should “do whatever it takes” [7] to prevent the economy tipping into such a “bad phase”, taking all measures that are seen to help the economy recover and shorten the recession period, such as “helicopter money” and easy access to credit for firms.

Although our model is not realistic on several counts and should no doubt be enriched, we believe that it offers interesting scenarios for recovery that helps sharpening one’s intuition and anticipating consequences that are often outside of the grasp of traditional approaches, where non-linear feedback effects and collective phenomena are absent. Because of heterogeneities and non-linearities, these emerging surprises are hard to anticipate and we need to develop qualitative numerical simulations, aka *telescopes for the mind* [8]. Although ABMs are spurned because they are hard (perhaps impossible) to calibrate, we have long argued [9] that one should abandon the “pretense of knowledge” and false sense of control provided by mainstream models and opt for a more qualitative, scenario driven approach to macroeconomic phenomena, with emphasis on mechanisms, feedback loops, etc. rather than on precise, but misleading, numbers. As Keynes famously said: *It is better to be roughly right than precisely wrong*. This is all the more so for policy makers in the face of a major crisis, such as the Covid shock.

Although highly stylized, our ABM generates a surprisingly rich variety of behaviour, in fact all the recovery letters listed above. Many parameters can be changed, such as those setting the “equilibrium” output and inflation levels, but also the length and severity of the shock, the amplitude of the policy response, etc. In the present note, we have only explored a small swath of possibilities. In order to allow our readers to experiment more and explore the variety of possible outcomes, we have put a version of our code on-line [here](#).

2 A short recap on Mark-0

The Mark-0 model with a Central Bank (CB) and interest rates has been described in full details in [6, 9, 10], where pseudo-codes are also provided. It was originally devised as a simplification of the Mark family of ABMs, developed in [11, 12]. We will not repeat here the full logic of the model, but only focus on the elements that are relevant for our crisis/recovery experiments.

First, we need some basic notions. The model is defined in discrete time, where the unit time between t and $t + 1$ will be chosen to be ~ 1 month. Each firm i at time t produces a quantity $Y_i(t)$ of perishable goods that it attempts to sell at price $p_i(t)$, and pays a wage $W_i(t)$ to its employees. The demand $D_i(t)$ for good i depends on the global consumption budget of households $C_B(t)$, itself determined as an inflation rate-dependent fraction of the household savings. D_i is a decreasing function of the firm price p_i , with a price sensitivity parameter that can be tuned. To update their production, price and wage policy, firms use reasonable “rules of thumb” [9] that also depend on the inflation rate through their level of debt (see below). For example, production is decreased and employees are made redundant whenever $Y_i > D_i$, and vice-versa.¹ The model is fully “stock-flow consistent” (i.e. all the stocks and flows within the toy economy are properly accounted for). In particular, there is no uncontrolled money creation or destruction in the dynamics.² We will actually allow some money creation below, when “helicopter money” policies will be investigated.

In Mark-0 we assume a linear production function with a constant productivity, which means that output Y_i and labour N_i coincide, up to a multiplicative factor ζ : $Y_i = \zeta N_i$. The unemployment rate u is defined as:

$$u(t) = 1 - \frac{\sum_i N_i(t)}{N}, \tag{1}$$

where N is the number of agents. Note that firms cannot hire more workers than available, so that $u(t) \geq 0$ at all times – see Eq. (5) below.

We assume that the banking sector – described at the aggregate level by a single “representative bank” – sets the interest rates on deposits and loans ($\rho^d(t)$ and $\rho^l(t)$ respectively) uniformly for all lenders and borrowers. Therefore, the rate ρ^l increases and ρ^d decreases when the firm default rate increases, in such a way that the banking sector (i.e., the representative bank) – which fully absorbs these defaults – makes zero profit at each time step (see [6, 10] for more details).

Although we have explored at length the effect of monetary policy and inflation anticipations in [10], we disregard these aspects of the problem in the present study: the baseline interest rate fixed by the central bank is set to zero, and inflation expectations of both firms and households are also zero. There is no Taylor rule coupling between inflation and interest rates either. The rationale for this choice is that we expect classical monetary policy tools to be quite ineffective as emergency measures, although they might be important to determine the long term fate of our toy economies. We leave this issue for further investigations.

2.1 Households

We assume that the total consumption budget of households $C_B(t)$ is given by:

$$C_B(t) = c [S(t) + W(t) + \rho^d(t)S(t)], \tag{2}$$

where $S(t)$ is the savings, $W(t) = \sum_i W_i(t)N_i(t)$ the total wages, and $\rho^d(t)$ is the interest rate on deposits, and c is the “consumption propensity” of households. If c is chosen to increase with increasing inflation [6, 10], then Eq. (2) describes a feedback of inflation on consumption similar to the standard Euler equation of DSGE models (see e.g. [13]). However, we neglect this effect in the present note. The total household savings evolve according to:

$$S(t + 1) = S(t) + W(t) + \rho^d(t)S(t) - C(t), \tag{3}$$

¹As a consequence of these adaptive adjustments, the economy is on average always ‘close’ to the global market clearing condition one would posit in a fully representative agent framework. However, small fluctuations persists in the limit of large system sizes giving rise to a rich phenomenology [9], including business cycles.

²In our baseline simulation, the total amount money in circulation is set to 0 at $t = 0$. This choice is actually irrelevant in the long run, but may have important short term effects.

where $C(t) \leq C_B(t)$ is the actual consumption of households, determined by the matching of production and demand, see [9].

2.2 Firms

2.2.1 Financial fragility

The model contains N_F firms (we chose $N_F = N$ for simplicity [9]), each firm being characterized by its workforce N_i and production $Y_i = \zeta N_i$, demand for its goods D_i , price p_i , wage W_i and its cash balance \mathcal{E}_i which, when negative, is the debt of the firm. We characterize the *financial fragility* of the firm through the debt-to-payroll ratio

$$\Phi_i = -\frac{\mathcal{E}_i}{W_i N_i}. \tag{4}$$

Negative Φ 's describe healthy firms with positive cash balance, while indebted firms have a positive Φ . If $\Phi_i < \Theta$, i.e. when the flux of credit needed from the bank is not too high compared to the size of the company (measured as the total payroll), the firm i is allowed to continue its activity. If on the other hand $\Phi_i \geq \Theta$, the firm i defaults and the corresponding default cost is absorbed by the banking sector, which adjusts the loan and deposit rates ρ^ℓ and ρ^d accordingly. The defaulted firm is replaced by a new one at rate φ , initialised at random (using the average parameters of other firms). The parameter Θ controls the maximum leverage in the economy, and models the risk-control policy of the banking sector.

2.2.2 Production update

If the firm is allowed to continue its business, it adapts its price, wages and production according to reasonable (but of course debatable) “rules of thumb” – see [6, 9]. In particular, the production update is chosen as follows:

$$\begin{aligned} \text{If } Y_i(t) < D_i(t) &\Rightarrow Y_i(t+1) = Y_i(t) + \min\{\eta_i^+(D_i(t) - Y_i(t)), \zeta u_i^*(t)\} \\ \text{If } Y_i(t) > D_i(t) &\Rightarrow Y_i(t+1) = Y_i(t) - \eta_i^-[Y_i(t) - D_i(t)] \end{aligned} \tag{5}$$

where $u_i^*(t)$ is the maximum number of unemployed workers available to the firm i at time t , which depends on its wage (see [10, Appendix A]). The coefficients $\eta^\pm \in [0, 1]$ express the sensitivity of the firm’s target production to excess demand/supply. We postulate that the production adjustment depends on the financial fragility Φ_i of the firm: firms that are close to bankruptcy are arguably faster to fire and slower to hire, and vice-versa for healthy firms. In order to model this tendency, we posit that the coefficients η_i^\pm for firm i (belonging to $[0, 1]$) are given by:

$$\begin{aligned} \eta_i^- &= \llbracket \eta_0^- (1 + \Gamma \Phi_i(t)) \rrbracket \\ \eta_i^+ &= \llbracket \eta_0^+ (1 - \Gamma \Phi_i(t)) \rrbracket, \end{aligned} \tag{6}$$

where η_0^\pm are fixed coefficients, identical for all firms, and $\llbracket x \rrbracket = 1$ when $x \geq 1$ and $\llbracket x \rrbracket = 0$ when $x \leq 0$. The factor $\Gamma > 0$ measures how the financial fragility of firms influences their hiring/firing policy, since a larger value of Φ_i then leads to a faster downward adjustment of the workforce when the firm is over-producing, and a slower (more cautious) upward adjustment when the firm is under-producing. Since the “dangerous” level of fragility is $\Phi = \Theta$, we assume that $\Gamma = \Gamma_0/\Theta$, where Γ_0 is an adjustable parameter. However, we neglected this effect in the present work and set $\Gamma_0 = 0$.

2.2.3 Price update

Following the initial specification of the Mark series of models [11], prices are updated through a random multiplicative process, which takes into account the production-demand gap experienced in the previous time step and if the price offered is competitive (with respect to the average price). The

update rule for prices reads:

$$\begin{aligned}
 \text{If } Y_i(t) < D_i(t) &\Rightarrow \begin{cases} \text{If } p_i(t) < \bar{p}(t) &\Rightarrow p_i(t+1) = p_i(t)(1 + \gamma\xi_i(t)) \\ \text{If } p_i(t) \geq \bar{p}(t) &\Rightarrow p_i(t+1) = p_i(t) \end{cases} \\
 \text{If } Y_i(t) > D_i(t) &\Rightarrow \begin{cases} \text{If } p_i(t) > \bar{p}(t) &\Rightarrow p_i(t+1) = p_i(t)(1 - \gamma\xi_i(t)) \\ \text{If } p_i(t) \leq \bar{p}(t) &\Rightarrow p_i(t+1) = p_i(t) \end{cases}
 \end{aligned} \tag{7}$$

where $\xi_i(t)$ are independent uniform $U[0, 1]$ random variables and γ is a parameter setting the relative magnitude of the price adjustment, chosen to be 1% (per month) throughout this work.³

2.2.4 Wage update

The wage update rule follows the choices made for price and production. Similarly to workforce adjustments, we posit that at each time step firm i updates the wage paid to its employees as:

$$\begin{aligned}
 W_i^T(t+1) = W_i(t)[1 + \gamma(1 - \Gamma\Phi_i)(1 - u(t))\xi'_i(t)] &\text{ if } \begin{cases} Y_i(t) < D_i(t) \\ \mathcal{P}_i(t) > 0 \end{cases} \\
 W_i(t+1) = W_i(t)[1 - \gamma(1 + \Gamma\Phi_i)u(t)\xi'_i(t)] &\text{ if } \begin{cases} Y_i(t) > D_i(t) \\ \mathcal{P}_i(t) < 0 \end{cases}
 \end{aligned} \tag{8}$$

where $\mathcal{P}_i = p_i \min(Y_i, D_i) - W_i N_i$ is the profit of the firm at time t and $\xi'_i(t)$ an independent $U[0, 1]$ random variable. If $W_i^T(t+1)$ is such that the profit of firm i at time t with this amount of wages would have been negative, $W_i(t+1)$ is chosen to be exactly at the equilibrium point where $\mathcal{P}_i(t) = 0$; otherwise $W_i(t+1) = W_i^T(t+1)$. Here, Γ is the same parameter introduced in Eq. (6).

Note that within the current model the productivity of workers is not related to their wages. The only channel through which wages impact production is that the quantity $u_i^*(t)$ that appears in Eq. (5), which represents the share of unemployed workers accessible to firm i , is an increasing function of W_i . Hence, firms that want to produce more (hence hire more) do so by increasing W_i , as to attract more applicants (see [6, Appendix A] for details).

The above rules are meant to capture the fact that deeply indebted firms seek to reduce wages more aggressively, whereas flourishing firms tend to increase wages more rapidly:

- If a firm makes a profit and it has a large demand for its good, it will increase the pay of its workers. The pay rise is expected to be large if the firm is financially healthy and/or if unemployment is low because pressure on salaries is high.
- Conversely, if the firm makes a loss and has a low demand for its good, it will attempt to reduce the wages. This reduction is more drastic if the company is close to bankruptcy, and/or if unemployment is high, because pressure on salaries is then low.
- In all other cases, wages are not updated.

The model, as presented above, has several free parameters. Some values are fixed throughout this work, using values that have been found in previous work to yield reasonable results [6, 9, 10]: their list is given in Table 1.

3 A Covid-like shock to the economy: phenomenology

The baseline values of the parameters, summarized in Table 1, allow our economy to settle in a rather prosperous state, with a low level of unemployment and, therefore, a near maximum output given the level of productivity $\zeta = 1$. The inflation level is $\approx 1.3\%$ /year and the average financial fragility (Φ) of

³In [10], we introduced a factor $(1 + \hat{\pi}(t))$ in the price and wage update rules, to model the fact that firms also factor in an anticipated inflation $\hat{\pi}(t)$ when they set their prices and wages. This effect is neglected here, as it plays a minor role in the present discussion.

Number of firms	N_F	10000
Consumption propensity	c	0.5
Price adjustment parameter	γ	0.01
Firing propensity	η_-^0	0.2
Hiring propensity	η_+^0	$R\eta_-^0$
Hiring/firing ratio	R	2
Bankruptcy threshold	Θ	3
Rate of firm revival	φ	0.1
Productivity factor	ζ	1
Financial fragility sensitivity	Γ_0	0

Table 1: Parameters of the Mark-0 model that are relevant for this work, together with their symbol and baseline values. For a comprehensive list of parameters, see [10].




Figure 1: Some recovery patterns following the coronavirus shock, which starts at time $t = 0$. We show, as a function of time, the fall in output relative to the no-shock scenario. For mild shocks ($\Delta c/c = 0.3$, $\Delta\zeta/\zeta = 0.1$, lasting six months), the economy contracts but quickly recovers (V-Shaped recovery). For more severe shocks over the same time period ($\Delta c/c = 0.3$, $\Delta\zeta/\zeta = 0.2$), the economy contracts permanently and never recovers (at least on the time scale of the simulation). This is the dreaded L-shaped scenario, in the absence of any government policy. An increase in consumption propensity to $c = 0.7$ (0.2 above pre-crisis levels) and a helicopter money drop at the end of the shock leads to a faster recovery (U-shaped recovery). Finally, a W-shaped scenario can also be found for stronger shocks ($\Delta c/c = 0.3$, $\Delta\zeta/\zeta = 0.5$, lasting nine months) with strong government policy and a helicopter drop (see next section).

the firms is ≈ 1 (i.e. a debt equal to one month of wages), far from the baseline bankruptcy threshold $\Theta = 3$.

The specificity of the Covid crisis is that it induced both a supply and a demand shock [14]. We can model this effect by a sudden drop of the productivity of firms, i.e. $\zeta \rightarrow \zeta - \Delta\zeta$, and of the consumption propensity of households, i.e. $c \rightarrow c - \Delta c$. These drops are meant to mimic the effect of a lock-down on the economy, that leads both to a drop of supply (employees must stay home and either not work at all or work remotely with lower productivity, while keeping their salaries) and a drop of demand (customers cannot go shopping, or are afraid to spend). It is uncertain how long the effects of the crisis would last. Hence, an important parameter describing the shock is its duration T . We will choose henceforth three benchmark values: 3 months, 6 months and 9 months. These values are meant to represent an effective length of the shock, accounting for the fact that lock-down measures can be partially lifted, which leads to an increased value of both ζ and c during the shock period and hence a shorter effective shock duration.

In Fig. 1, we show several typical crisis and recovery shapes, depending on the strength of the shock and the policy used to alleviate the severity of the crisis. For small enough Δc and/or $\Delta\zeta$, there is no drop of output at all. For larger shock amplitude or duration, one observes a V-shape recovery, as expected when the shock is mild enough not to dent the financial health of the firms. Stronger shocks can however lead to a permanent dysfunctional state (L-shape), with high unemployment, falling wages and savings, and a high level of financial fragility and bankruptcies. An L-shaped scenario can however be prevented if after the shock, consumer demand picks. To facilitate and boost consumption, a one-time policy of helicopter money can move the economy towards a path of recovery over the scale of a

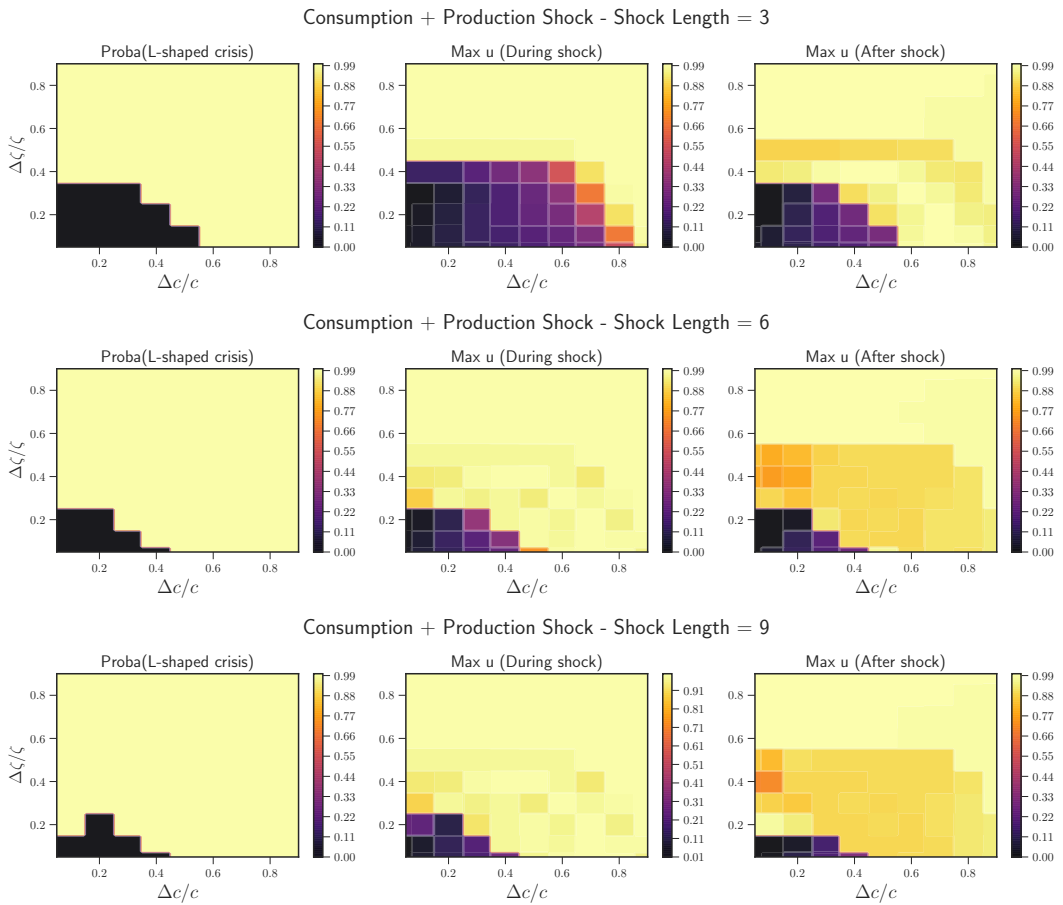


Figure 2: Phase diagrams in the $\Delta c/c - \Delta \zeta / \zeta$ plane for different shock lengths. As the length of the shock is increased, the probability of having a long-crisis increases even for mild shocks. *Top Row:* Shock lasting for 3 months. The region of parameter space with no L-shaped crisis is quite large allowing for strong consumption shocks ($\Delta c/c \lesssim 0.5$) and mild productivity shocks ($\Delta \zeta / \zeta \lesssim 0.3$). Note that for such a short shock, the effects on unemployment are seen after the shock has passed. *Middle Row:* Shock lasts for 6 months. A decrease in the region of no-crisis is observed. Mild shocks ($\Delta c/c \sim 0.4$) can also lead to prolonged crises. During the shock itself, extremely high rates of unemployment can be seen. *Bottom Row:* Shock lasts for 9 months. Only for extremely mild shocks does the economy not undergo a prolonged crisis.

few years (U-shape).

For a more complete picture of the influence of these shocks, we plot the “phase diagrams” of the crises in absence of any policy, in the plane $\Delta c/c, \Delta \zeta / \zeta$, for $T = 3, 6$ and 9 months in Fig. 2. We show (a) the probability of a “dire” (L-shaped) crisis, (b) the peak value of unemployment during the shock and (c) the peak value of unemployment after the shock. Black regions indicate that the economy survives well (i.e. no dire crisis, or short crises with little unemployment). This occurs, as expected, in the lower left corner of the graphs (small $\Delta c/c, \Delta \zeta / \zeta$). These black regions shrink as T increases. We also note that mild shocks lasting only a short time ($T = 3$ months) can cause lasting damage. Indeed, we observe that low rates of unemployment during the shock are not representative of the future evolution. Interestingly, there is an abrupt, first order transition line (a “tipping point” in the language of [9]) beyond which crises have a very large probability to be permanent, with high levels of unemployment (yellow regions). Because the location of such a tipping point in the real-world




Figure 3: Phase diagram for a pure consumption shock ($\Delta\zeta/\zeta = 0$). We observe an abrupt transition to an L-shaped crisis for consumption shocks beyond $\Delta c/c = 0.5$ for $T = 3$ months, and beyond $\Delta c/c = 0.4$ for $T = 6$ or 9 months. Below these shock amplitudes, there is no prolonged crises but short-lived crises are observed. This can be seen by observing the maximum unemployment rate during the shocks: for $\Delta c/c = 0.4$, we reach about 15% unemployment.

economy is extremely hard to estimate⁴, our results suggest that governments should be very cautious and do as much as possible to prevent a possible collapse of the economy.

It is useful to focus on the line $\Delta\zeta = 0$ of these two-dimensional plots, corresponding to a consumption shock without productivity shock. We show in Fig. 3 the same three quantities as in Fig. 2. An abrupt transition between no dire crises and dire crises can be seen for $\Delta c/c \sim 0.4$ when $T = 9$ months.

We now implement, within our model, some emergency governmental policy inspired from those that are actually currently in place in different countries. From now on, we choose $\Delta c/c = 0.3$ and $\Delta\zeta/\zeta = 0.5$ as reasonable values to represent the severity of the Covid shock [15, 16], and let T again take the values 3, 6 and 9 months. In the absence of any active policy, the economy collapses into a deep recession, with an output reduced by 2/3 compared to pre-shock levels.

4 Securing a quick recovery?

The toolbox developed in the aftermath of the Global Financial Crisis (GFC) of 2008 puts monetary policy at the center of economic crisis management. This takes the form of either direct interest rate cuts or, as was seen recently for the GFC interventions, even stronger measures such as quantitative easing. Our Agent Based Model provides several channels through which the economy can be propped up, including interest rate cuts [10]. However, given that the interest rates are already very low, the interest-rate channel itself might not be effective, and might lead to a stagnation trap and a L-shaped recovery [4]. Hence, in this work, we disregard the interest-rate channel, since it cannot be used as an emergency measure in the face of a collapsing supply sector. We focus on two possible channels: easy credit for firms, and “helicopter money” for households.

A way to loosen the stranglehold on struggling firms is to give them easy access to credit lines, independently of their financial situation. In our model, this amounts to a significant increase of the bankruptcy threshold Θ . So the policy we investigate is the following: during the whole duration of the shock, we set $\Theta = \infty$, i.e. all firms are allowed to continue their business and accumulate debt. When the shock is over, the value of Θ is taken back down. This can be done in several ways. One extreme possibility (that we call naive below) is to set Θ to its pre-shock value as soon as the shock is over. Intuitively, when the shock is short enough, allowing endangered firms to survive might be enough. For long shocks however, such a naive policy is not going to be very helpful as firms that have muddled through the shock have become much more fragile at the end of the shock. So in this case, many will fail when credit is tightened, and the economy plunges into recession as if no policy was applied. This is precisely what is shown in Figs. 4, 5, second column (“Naive Policy”), where we plot (as in Fig. 6 below) a dashboard of the state of the economy: output and unemployment, financial

⁴In physics, it is well known that there is no way to know that one is approaching a first order transition between two phases, just looking from within a single phase. Hence, no simple indicator would predict the tipping point of the economy from within its “good” phase.

Consumption + Production Shock - Shocklength = 3, $\Delta c/c = 0.3$, $\Delta \zeta/\zeta = 0.5$

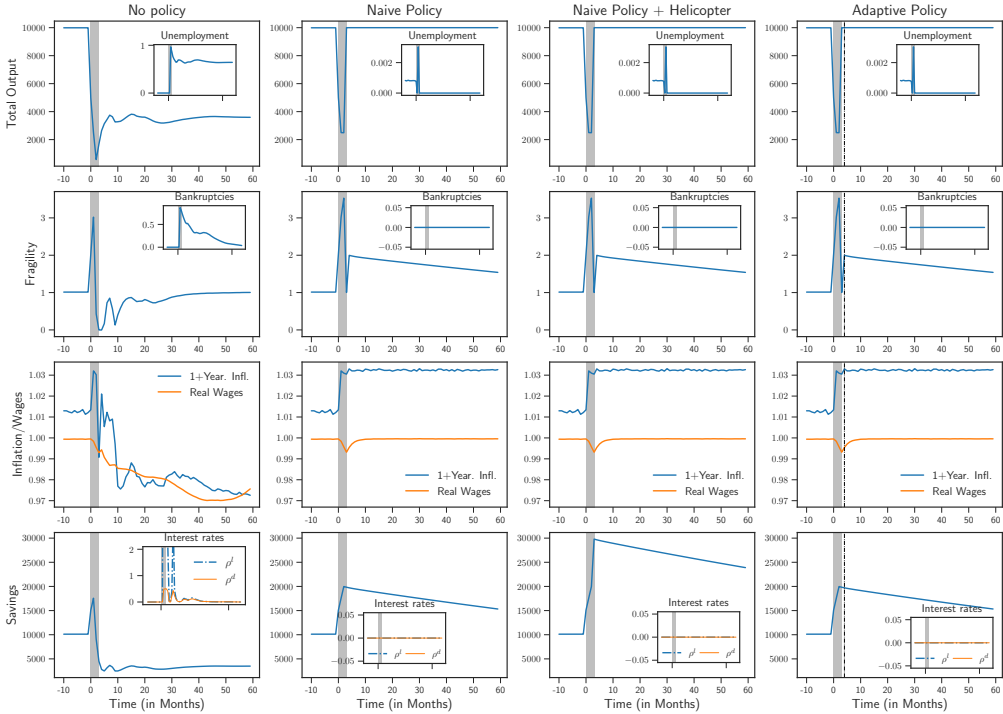


Figure 4: Scenarios for shock length of 3 months (marked in grey) with and without policy. *First Column:* Without a policy intervention, the economy suffers a deep contraction with extremely high rates of unemployment and a subsequent loss in wages. A large number of firms go bankrupt and household savings reduce permanently after a brief increase during the shock. Permanent deflation and drop of real wages are also observed. A rapid increase in the interest rate on loans ρ^l following firms going bankrupt. *Second Column:* A policy of extending the credit limits for all firms is introduced, which lasts the duration of the crisis. This improves the situation of the economy. A temporary contraction in output can not be avoided but the policy is able to prevent bankruptcies and hence keep unemployment during and beyond the crisis very low. *Third Column:* The situation with the naive policy followed by a helicopter drop of money is shown. Since the naive policy by itself was enough to prevent a crisis, the helicopter drop does not change the outcome, apart from increasing the savings of households. Note that the money injected into the economy by the helicopter drop quickly disappears due to inflation. *Fourth Column:* An adaptive policy which reduces the bankruptcy threshold Θ gradually is essentially equivalent to the naive policy in this case. The vertical dotted line marks the end of this policy. Note that in all cases where policies save the day, inflation picks up after the crisis.

fragility and default rate, inflation and wages, savings and interest rates. Recent data indeed points to this scenario bearing out with bankruptcies set to soar in the coming months [17].

Note that whenever the credit policy is successful, inflation shoots up to $\approx 3\%$ post-crisis. We only show the results corresponding to $T = 3$ months and $T = 9$ months. For the shock amplitude that we have chosen, the case $T = 6$ months is qualitatively similar to the case $T = 3$ months and is therefore not shown.

Another possibility, that we call “adaptive”, is to reduce Θ progressively, in a way that is adapted to firms’ average fragility. We assume that the government measures the instantaneous value of $\langle \Phi \rangle$ over firms still in activity, weighted by production, $\langle \Phi \rangle = \sum_i \Phi_i Y_i / \sum_i Y_i$, and sets Θ as:

$$\Theta = \max(\theta \langle \Phi \rangle, 3), \quad (t > T), \tag{9}$$

where θ is some offset that we chose to be $\theta = 1.25$. This means that only the most indebted firms, whose fragility exceeds the average value by more than 25%, will go bust as the effective threshold Θ is progressively reduced. As shown in Figs. 4, 5 fourth column (“Adaptive Policy”), this scheme is very

successful: the economy recovers 100% of its pre-shock output at the end of the shock, for all three durations $T = 3, 6, 9$ months. As can be seen from the plot of the average fragility, this comes at the price of $\langle \Phi \rangle$ reaching very high values for a while (for example $\langle \Phi \rangle \approx 6$, i.e. six times its pre-shock value, when $T = 9$ months), and, again, a much higher post-crisis inflation ($\approx 3\%$). But the slow removal of the easy credit policy allows the economy to smoothly revert to its pre-shock state, with a limited number of bankruptcies.

Consumption + Production Shock - Shocklength = 9, $\Delta c/c = 0.3$, $\Delta \zeta/\zeta = 0.5$

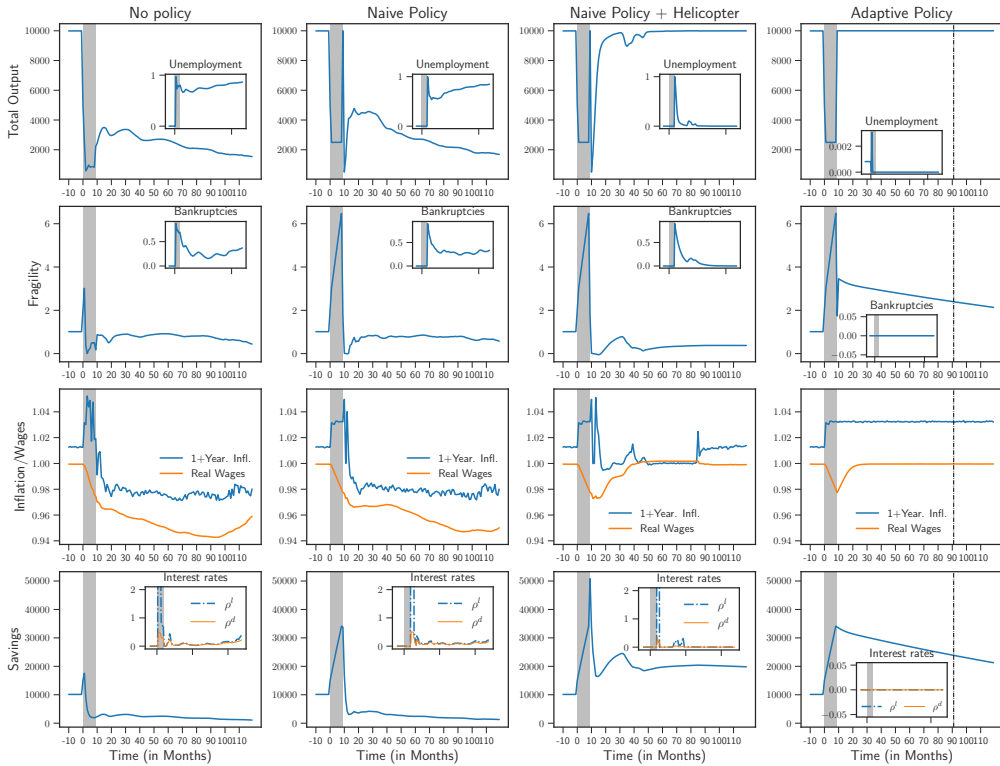


Figure 5: Scenarios for shock length of 9 months (marked in grey) with and without policy. *First column:* Similar to Fig. 4, the economy undergoes a severe and prolonged contraction. However, given the length of the shock, there is a deeper fall in the level of real wages with firms continuing to go bankrupt far after the shock has occurred. *Second Column:* The presence of the naive policy in this case is unable to rescue the economy and in turn exacerbates the situation. Given the already fragile nature of the firms, removing the easy credit policy abruptly leads to a further spate of bankruptcies. This leads to wages being depressed further and unemployment remaining high. *Third Column:* The introduction of helicopter money improves upon the naive policy intervention at the expense of the economy undergoing another endogenous crisis. This leads to the W-shaped scenario from Fig. 1. *Fourth Column:* The adaptive policy in this situation drastically improves the economic outcomes. The contraction in output is inevitable but by providing firms the support they need for as long as possible (for more than 6 years here), the policy is able to keep unemployment low and prevent any bankruptcies due to the shock. The vertical dotted line marks the end of this policy.

Consumption Shock - Shocklength = 9, $\Delta c/c = 0.7$

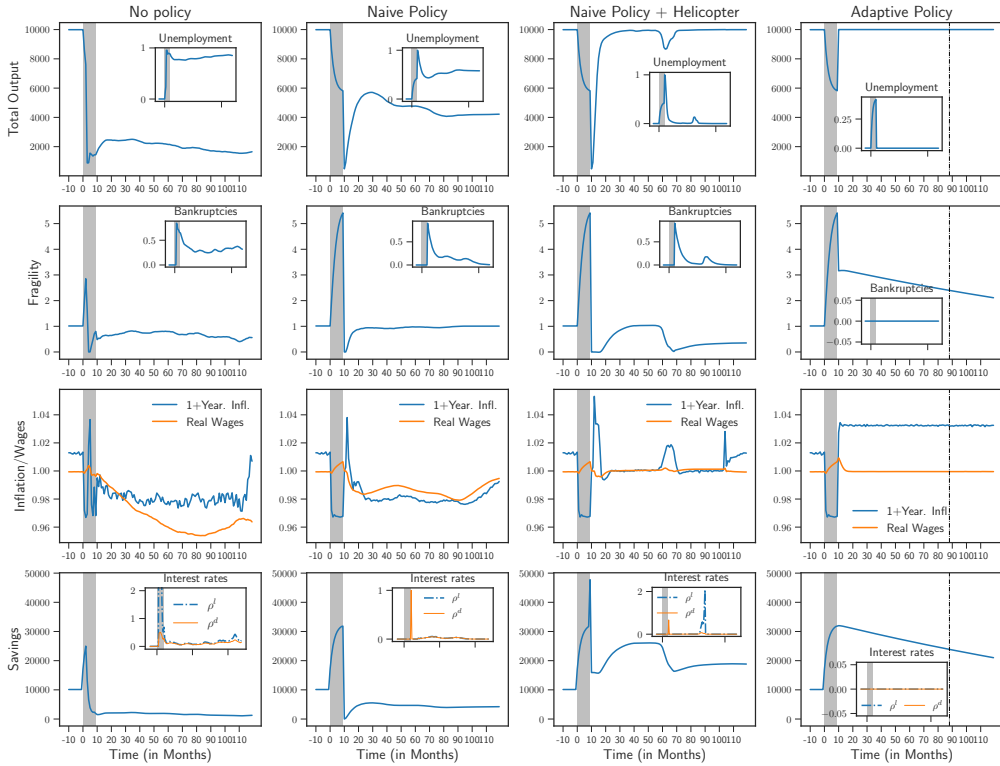


Figure 6: Scenarios for a severe consumption shock $\Delta c/c = 0.7$ of length of 9 months (marked in grey) with and without policy for a pure consumption shock. *First column:* A prolonged crisis with a deep contraction is observed similar to the situation shown in Fig. 5. *Second column:* The naive policy is not sufficient to mitigate the crisis. In fact, removing the policy as the shock ends leads to further contraction and higher rate of bankruptcies. *Third column:* With the presence of helicopter money to boost spending, we observe a rapid recovery. However, several short-lived crises is observed after the initial shock-induced crisis (W-shape recovery). *Fourth column:* With an adaptive policy, we are able to prevent bankruptcies and keep unemployment in control as well.

Note that at the end of the shock, when c returns to its original value, households start to over-spend with respect to the pre-crisis level, because their savings increase during the shock (mirroring the increase of firms' debt) and they want to spend a fixed fraction of them, see Eq. (2). However, this over-spending can be insufficient to drive back the economy to its pre-crisis state.

Another possible policy is thus to inject cash in the economy to boost consumption and facilitate recovery. This is often nicknamed “helicopter money”. This involves the expansion of the money supply by the central bank and has multiple transmission channels: the central bank transfers cash directly to its citizens or it can transfer it directly to the government which in turn would spend it on healthcare or infrastructure projects. This policy has been considered radical due to the fear that an expansion in money supply might lead to runaway inflation. In normal times, there might be support for such a view, but it has been shown that a helicopter drop may not always be inflationary [18]. Given the enormity of the crisis, there have been calls from all corners for central banks to break “taboos” [7, 19, 20] and do what is necessary.

In this work, we implement a helicopter-money drop by assuming that the government distributes money to households multiplying their savings by a certain factor $\kappa > 1$: $S \rightarrow \kappa S$. The distribution takes place at the end of the shock, and we study here how the “naive policy” (for which Θ goes back to its baseline value immediately after the shock) can be improved by some helicopter money.

Results for $\kappa = 1.5$ are shown in Figs. 4, 5 in the third column (“Naive Policy + Helicopter Money”). We indeed see that in the case $T = 9$ months, for which the naive policy was not sufficient to prevent a prolonged recession, increasing the consumption budgets of households does allow the economy to recover. However, a quite interesting effect appears, in the form of a W-shape, or relapse of the economy, *even in the absence of a second lock-down period*. This “echo” of the initial shock is due to financially fragile firms that eventually have to file for bankruptcy when credit has tightened. This second blip is however temporary and the economy manages to settle back on an even keel. This experiment shows the importance of boosting consumption when the shock is over. A similar effect would be obtained if instead of the savings S , the consumption propensity c was increased post-lock-down. This echoes pleas from policy makers, wooing households into over-spending once the shock period is over. A combination of the two might indeed lead the economy to a faster recovery as shown in the U-shape recovery in Fig. 1.

We also studied the case of a pure, rather severe consumption shock $\Delta c/c = 0.7$ lasting $T = 9$ months in Fig. 6. We observe that a prolonged drop in consumption, without any loss in production, can still lead to long-lived crisis. The “naive” policy in this case is not enough to hasten the recovery. Direct cash transfer to households via helicopter money drop helps the economy recover faster but leads to a slow, W-shape recovery. Finally, the “adaptive” policy again works best in keeping unemployment low and ensures a rapid recovery.

5 Discussion & Conclusion

In this paper, we have discussed the impact of a Covid-like shock on the toy economy described by the Mark-0 Agent-Based Model developed in [6, 9, 10]. We have shown that, depending on the amplitude and duration of the shock, the model can describe different kind of recoveries (V-, U-, W-shaped), or even the absence of full recovery (L-shape). Indeed, as we discussed in [9], the non-linearities and heterogeneities of Mark-0 allow for the presence of “tipping points” (or phase transitions in the language of physics), for which infinitesimal changes of parameters can induce macroscopic changes of the economy. The model display a self-sustained “bad” phase of the economy, characterized by absence of savings, mass unemployment, and deflation. A large enough shock can bring the model from a flourishing economy to such a bad state, which can then persist for long times, corresponding to decades in our time units⁵.

We have then studied how government policies can prevent an economic collapse. We considered

⁵Whether such a bad phase is truly stable forever or would eventually recover (via a nucleation effect similar to metastable phases in physics) is an interesting conceptual point. It could also have practical implications because if recovery happens via nucleation, one could imagine triggering it via the artificial creation of the proper “nucleation droplet”, in the physics parlance. We leave this discussion for future studies.

two policies that are currently being implemented in several countries: helicopter money for households and easy credit for firms. We find that some kind of easy credit is needed to avoid a wave of bankruptcies, and mixing both policies is effective, provided policy is strong enough. We also highlight that, for strong enough shocks, some flexibility on firm fragility might be needed for long time (a few years) after the shock to prevent a second wave of bankruptcies [17]. Too weak a policy intervention is not effective and can result in a “swoosh” recovery or no recovery at all. Again, a threshold effect is at play, with potentially sharp changes in outcome upon small changes of policy strength.

Our results then suggest that governments should try to be on the safe side and do “whatever it takes” to prevent the economy to fall in a bad state, and stimulate a rapid recovery. However, we find that when policy is successful, inflation post-crisis is significantly increased compared to the pre-crisis period.

There are, however, some major limitations of our study. For example, in our model money is conserved and is essentially equal to zero (in real value) in the good phase of the economy, because any initial amount of money is washed away by inflation. Hence, total savings equal total debt (unless some helicopter money is injected). The interest rates on both deposits and loans are determined by a central bank via a zero-profit rule, in order to absorb the costs of defaults [6, 10]. Mark-0 thus correctly describes the firm bankruptcies due to excessive debt, and the resulting increase of the interest rate on loans (and decrease of the interest rate on deposits). However, in Mark-0 there is no splitting of debt into a “public” and a “private” sector, hence no competition between investments in corporate bonds and in government bonds. As a result, within the current framework we cannot model possible “panic” effects that would result from a ballooning public debt, which could potentially lead to an increase of the yield of government bonds, possibly resulting in runaway public debt, confidence collapse and hyperinflation. This is indeed the major objection currently being raised against a stronger governmental response.

Similarly, there is no coupling, in the current version of Mark-0, between a firm financial fragility and the interest rate on its debt (i.e. no extra risk premium for fragile firms). This could again lead to a run away mechanism and a collapse of the corporate sector. Modeling all these effects is possible within Mark-0, but we leave these important extensions for future work. We note that in any case our results show that the excess debt accumulated during the crisis decays (via excess inflation) over the scale of a few years, provided economic recovery is achieved.

Despite these limitations, we believe that our model is flexible enough and captures enough of the basic phenomenology to be used as an efficient “telescope of the mind” [8]. One can play with the parameters to investigate qualitatively the different scenarios that can arise under different policies and shocks. For instance, one could impose a different shock on the economy by reducing the value of R (the ratio of firms’ hiring and saving propensity), either by increasing the speed at which firms fire their employees, or by reducing the hiring rate.⁶ A low enough value of R indeed drives the Mark-0 economy to a bad state [9]. A different policy, which we did not consider here, is for the government to pay the wages directly, allowing firms to maintain their financial health unscathed. This can also easily be implemented in Mark-0, but should be roughly equivalent to an increase of Θ .

In this work, we did not investigate the feedback channels modeled by the awareness parameter Γ [see Eqs. (6) and (8)], which was set to zero throughout our work. A positive Γ means that firms’ hiring/firing propensity and wage policy depend on their financial fragility, which could lead to interesting effects. We again leave this for a future investigation.

As we have emphasized above, we have not investigated in this study the standard monetary policy tool, namely interest rate cuts. Whereas such cuts are not expected to play a major role in the short term management of the crisis, their effect on the long-term fate of the economy (in particular when the recovery is L-shaped) can be important and should be examined as well. Readers interested in this issue can use the code available on-line [here](#).

Yet another direction for future investigations would be to consider the effect of successive lockdowns due to subsequent spikes of Covid infections. It would be interesting to study the different recovery patterns that can arise in this case and assess which policy strategy is the most effective,

⁶This actually happened during the lock-down, as most hires were frozen.

perhaps coupling Mark-0 with SIR-like models.⁷

Finally, we believe that one of the most needed extension of Mark-0 is to allow the role of inequalities (of firm sizes and of household wealth and wages) to be discussed. One of the peculiarities of the Covid shock has been the asymmetry in the way the crisis has affected households, with lower spending by high-income households compounding the situation for low-income households [21]. Taking into account heterogeneities in income and effects of the shock would bring our ABM closer to reality, while addressing one of the most pressing issues of our current times.

References

- [1] Paul Hannon and Saabira Chaudhuri. Why the Economic Recovery Will Be More of a ‘Swoosh’ Than V-Shaped. *Wall Street Journal*, May 2020. Retrieved 12 June 2020.
- [2] Martin Eichenbaum, Sergio Rebelo, and Mathias Trabandt. The Macroeconomics of Epidemics. *National Bureau of Economic Research, Working Paper Series*, Mar 2020.
- [3] Warwick J. McKibbin and Roshen Fernando. The Global Macroeconomic Impacts of COVID-19: Seven Scenarios. *SSRN Electronic Journal*, pages 1–43, March 2020.
- [4] Luca Fornaro and Martin Wolf. Covid-19 Coronavirus and Macroeconomic Policy. *Barcelona GSE Working Paper Series*, March 2020.
- [5] Veronica Guerrieri, Guido Lorenzoni, Ludwig Straub, and Iván Werning. Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages? *National Bureau of Economic Research Working Paper Series*, 53(9):1689–1699, April 2020.
- [6] Stanislao Gualdi, Marco Tarzia, Francesco Zamponi, and Jean-philippe Bouchaud. Monetary policy and dark corners in a stylized agent-based model. *Journal of Economic Interaction and Coordination*, 12(3):507–537, Oct 2017.
- [7] Richard Baldwin and Beatrice Weder di Mauro. *Mitigating the COVID Economic Crisis: Act Fast and Do Whatever It Takes*. Vox CEPR Policy Portal, 2020.
- [8] Mark Buchanan. This Economy does not compute. *New York Times*, 2008. Retrieved 12 June 2020.
- [9] Stanislao Gualdi, Marco Tarzia, Francesco Zamponi, and Jean Philippe Bouchaud. Tipping points in macroeconomic agent-based models. *Journal of Economic Dynamics and Control*, 50:29–61, 2015.
- [10] Jean-Philippe Bouchaud, Stanislao Gualdi, Marco Tarzia, and Francesco Zamponi. Optimal inflation target: insights from an agent-based model. *Economics: The Open-Access, Open-Assessment E-Journal*, pages 1–19, Sep 2018.
- [11] Edoardo Gaffeo, Domenico Delli Gatti, Saul Desiderio, and Mauro Gallegati. Adaptive Microfoundations for Emergent Macroeconomics. *Eastern Economic Journal*, 34(4):441–463, Oct 2008.
- [12] Domenico Delli Gatti, Saul Desiderio, Edoardo Gaffeo, Pasquale Cirillo, and Mauro Gallegati. *Macroeconomics from the Bottom-up*. Springer Science & Business Media, 2011.
- [13] Frank Smets and Raf Wouters. An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area. *Journal of the European Economic Association*, 1(5):1123–1175, Sep 2003.
- [14] R. Maria del Rio-Chanona, Penny Mealy, Anton Pichler, Francois Lafond, and Doyne Farmer. Supply and demand shocks in the COVID-19 pandemic: An industry and occupation perspective. *Covid Economics*, (6), Apr 2020.

⁷Note again, however, that W-shape recoveries can be observed in our model even in absence of a second wave of infections.

- [15] Mark Carney and Ben Broadbent. Bank of England Monetary Policy Report May 2020. Technical Report May, Bank of England, 2020.
- [16] Pierre Aldama, Raphaël Cancé, Marion Cochard, Marie Delorme, William Honvo, Yannick Kalantzis, Guy Levy-Rueff, Emilie Lor-Lhommet, Jean-Francois Ouvrard, Simon Perillaud, Béatrice Rouvreau, Paul Sabalot, Katja Schmidt, Antoine Sigwalt, Camille Thubin, Youssef Ulgazi, Paul Vertier, and Thao Vu. Macro-economic projections June 2020. Technical report, Banque de France, June 2020.
- [17] Mary Williams Walsh. A tidal wave of bankruptcies is coming. *New York Times*, 2020. Retrieved 19 June 2020.
- [18] Willem H. Buiter. The Simple Analytics of Helicopter Money: Why It Works — Always. *Economics: The Open-Access, Open-Assessment E-Journal*, 8(1), 2014.
- [19] Sony Kapoor and Willem H. Buiter. To fight the COVID pandemic , policymakers must move fast and break taboos. *Vox CEPR Policy Portal*, 2020. Retrieved 12 June 2020.
- [20] Jordi Gali. Helicopter Money: the time is now. *Vox CEPR Policy Portal*, 2020. Retrieved 12 June 2020.
- [21] Emily Badger and Alicia Parlapiano. The rich cut their spending. that has hurt all the workers who count on it. *New York Times*, 2020. Retrieved 19 June 2020.

Spreading the disease: The role of culture

Ioannis Laliotisi¹ and Dimitrios Minos²

Date submitted: 18 July 2020; Date accepted: 20 July 2020

This paper investigates the “cultural” transmission of the SARS-CoV-2 outbreak. Using data from Germany we observe that in predominantly Catholic regions with stronger social and family ties, the spread and the resulting deaths per capita were much higher compared to non-Catholic ones at the local NUTS-3 level. The result is strengthened with Difference-in-Difference estimates at the regional NUTS-1 level. This finding could help explain the rapid spread and high death toll of the virus in some European countries compared to others. Looking at differences within a specific country in a well identified setting eliminates biases due to different social structures, health care systems, specific policies and measures, and testing procedures for the virus that can confound estimates and hinder comparability across countries. Further, we use individual level data as well as mobility data from mobile devices to investigate potential mechanisms. The results highlight the cultural dimension of the spread and could suggest the implementation of targeted mitigation measures in light of disease outbreaks.

- 1 Lecturer in Economics, Programme Director MSc Health Economics & MSc Economic Evaluation in Healthcare, Department of Economics, City University of London.
- 2 Teaching Fellow in Economics, King's Business School, King's College London.

Copyright: Ioannis Laliotisi and Dimitrios Minos

1. Introduction

Since the recent onset of the SARS-CoV-2 outbreak, which resulted into the COVID-19 pandemic in November 2019 in Wuhan, China the virus has spread rapidly across the globe. Italy and Spain are the countries in Europe that were hit the earliest and hardest by April 2020, with France following suit, whereas Northern European countries seemed to have performed better at least in the initial phase, with the exception of the UK. As a result, and despite the relatively homogeneous containment measures (e.g. school closures, lockdowns, social distancing etc.), the incidence and death toll vary substantially across countries and it is difficult to uncover the reasons behind this. The timing of both the onset of the pandemic and the various policies to constrain it could clearly play a role; for example in Sweden, Belgium or the UK. However, recent evidence suggests that the virus was present in Europe long before it was initially believed, e.g. as early as December 2019 in France (Deslandes et al. (2020)). Yet, it spread faster in some countries compared to others, at least in its initial stage, setting of major outbreaks and resulting in vastly overburdened health systems and excess death rates, before most governments enacted their mitigation policies.

An explanation for this phenomenon could be that spatial variation in social norms and culture, e.g. lower levels of social interactions and culturally inherent “social distancing” (Remland et al. (1995); Sorokowska et al. (2017)) may have slowed the spread of the virus to vulnerable groups, e.g. the elderly (Bayer and Kuhn (2020)) in some countries compared to others. The emerging COVID-19 literature highlights that for as long as a medical treatment is not available, social capital and human behaviour are important in containing the transmission of the disease (Bartscher et al. (2020); Durante et al. (2020); Van Bavel et al. (2020)). For example, Bargain and Aminjonov (2020) argue that compliance to containment policies was higher in European regions where people trust their politicians more.

Recent empirical evidence has demonstrated that variations in disease transmission patterns are crucially characterised by the type and frequency of interactions among people, especially in early stages (Alfaro et al. (2020); Liu et al. (2020); Platteau and Verardi (2020)). For example, Dowd et al. (2020) show that in countries like Italy and Spain, co-residence and intergenerational social interaction can partially explain variation in COVID-19 incidence. However, substantial differences in social structures, demographics, health systems, testing, social distancing policies and various types of policy response at different time periods, and event data collection processes, render cross-country comparisons problematic.¹ Recent preliminary evidence has also demonstrated how genetic

¹Differences in the data collection process might arise from the fact that different countries can lie on different points of their epidemic curve, hence not applying the same standards in testing and reporting

differences affect the cross-country variation of the virus spread (Delanghe et al. (2020)). Moreover, such cross-country differences did not allow for a coordinated virus mitigation and lockdown exit strategy at the EU level (Platteau and Verardi (2020)). Therefore, empirical research shifts towards within-country analyses in order to identify the role of social structures on the virus spread. Bayer and Kuhn (2020) observed a country-level positive correlation between the virus spread and the share of people 30-49 years old living with their parents, arguing that social integration might be the factor behind the death toll in Italy. However, following research by Belloc et al. (2020) using Italian regions as their units of observation revealed that the correlation has the opposite sign, suggesting policy recommendations based on cross-country results should be taken with caution. Even though European governments implemented nation-wide measures during the early phases of the pandemic, virus spread and death toll varied substantially at the local level within countries.

For these reasons, we choose to focus on a single country; Germany, where such concerns are rather mitigated. Some other early research attempts to examine whether this variation at the local level is due to differences in compliance to common mitigation measures. Specifically, two studies for Switzerland have focused on the cultural and behavioural gradient of the COVID-19 spread. Mazzona (2020) argues that elderly people living in Latin-speaking (French or Italian) cantons were more severely affected by the virus spread, relative to those in cantons where the German language dominates. This was attributed to the fact that mobility in Latin-speaking cantons was reduced earlier as people in those areas complied more strictly to nation-wide mitigation policies. These results point to the same direction with Brodeur et al. (2020) and Durante et al. (2020) who reported that mobility declines were higher in regions with higher levels of civic capital and trust. However, they contradict with the findings of Deopa and Fortunato (2020) who demonstrate that high-trust German-speaking cantons reduced their mobility less than the French-speaking cantons, arguing that mobility becomes less relevant for the virus spread due to higher compliance to infection prevention and control norms.

Nevertheless, this growing early literature emphasises on the implications of culture on the virus spread. Mazzona (2020) and Deopa and Fortunato (2020) used language as a proxy for culture in 26 cantons in Switzerland. In this paper we proxy culture through religious denomination in 312 NUTS-3 West Germany regions. Guiso et al. (2006) defined culture as those customary values and beliefs being unchangeably transmitted from generation to generation within ethnic, religious and social groups. Speaking language is largely influenced by individual settlement decisions and often voluntarily accumulated. Cultural and behavioural dimensions shaped by religion are inherited by previous COVID-19 cases (Belloc et al. (2020)).

erations, they are not easily modified within the span of a lifetime and individuals have less control on them compared to other social capital; hence they can be treated as being lifetime invariant (Becker (1996); Guiso et al. (2006)). Therefore, our paper is the first to contribute to the current literature by shedding light on how the COVID-19 spread within a country is affected by cultural variations that can be considered as exogenous and to a large extent unclouded by unobserved heterogeneity.

We observe a substantial discrepancy even at the local authority (NUTS-3) level in per capita incidence and number of deaths between Catholic and non-Catholic regions using daily data from the Robert Koch Institute during the initial stages of the outbreak. This could hint at a type of “cultural transmission” which remains strong after partialling out socioeconomic characteristics, geographical proximity to Northern Italy and local fixed effects at the NUTS-3 level, as well as mobility trends per religion group before and during the lockdown. Moreover, we confirm our findings by using difference-in-differences estimates at the NUTS-1 level to eliminate any unobserved heterogeneity. Catholic regions that arguably exhibit stronger social and family ties seem to experience a wider spread of the virus in the general -as well as the more vulnerable (elderly) population- as indicated by their higher death toll. Previous research in economics has established that Catholics are more bound to close social circles and networks, e.g. family and friends, and they have different patterns of social interactions (Arruñada (2009); Glaeser and Glendon (1998); Ekelund et al. (2002); Satyanath et al. (2017)). These cultural differences in behaviour and social ethics can trigger a differentiated transmission of the disease in the population during early stages. We demonstrate that the COVID-19 incidence varies systematically with religion at the regional level. The results suggest that culture leads to different transmission rates within groups. The proposed mechanism is supported by results using individual-level data from the European Social Survey (ESS), the European Values Survey (EVS) and mobility data from mobile phone users that help us investigate how Catholics are systematically and culturally different in their social and family ties, and hence uncover a potential disease transmission channel within this group.

Our results could support epidemiologists and public health policy makers to better understand how cultural factors can largely influence the initial spread of pathogens in a society, and account for differences in such characteristics when designing response policies as suggested by Platteau and Verardi (2020). The remainder of the paper is organised as follows: In Section 2 we present an overview of the pandemic and policy response in Germany. Section 3 describes our identification strategy, data and results, whereas Section 4 discusses and concludes.

2. Background

According to data released from the Robert Koch Institute (RKI), the first reported cases of COVID-19 in Germany were recorded on 27 January 2020 in Starnberg, Bavaria and by March all German states had reported cases. The first deaths were reported on 9 March in North Rhine-Westphalia and by the end of the month all states confirmed hospitalised virus-related deaths. State and federal government response was swift (Stafford (2020)). In early March the federal government and the RKI issued the National Pandemic Plan to be carried out across the country. The states themselves were given some autonomy in handling the pandemic, although the response was coordinated and crucial to avoid inter-state travel. All states in Germany enacted strict social distancing measures and closures of schools, shops and workplaces. Based on the Coronavirus Government Response Tracker developed by the University of Oxford, these social distancing measures were implemented on February 29 (public events cancellation), March 16 (school closings), March 22 (shops and workplaces closings) and April 09 (public transport closing). The number of daily cases peaked in early April according to the RKI and started receding from that point on. Some of the nationwide restrictions such as the closure of small shops were lifted about three weeks later on 22 April.

In terms of testing, widespread PCR testing was made available on 25 March and guidelines to test severe cases only were lifted, so that more than 2 million people received a test by late April. Data on cases and deaths are collected by local authorities (NUTS-3), reported to the Federal Ministry of Health and published by RKI after validation. As of 23 April 2020, 148,046 confirmed cases were reported as well as 5,094 deaths. Figure A.1 displays how the cumulative numbers of reported cases and deaths per 100,000 population were scattered across the country on April 20.

This is the period we will focus on, as the initial stages of the pandemic largely determine the speed of the spread (Zhao et al. (2020)). Given the relatively long incubation period of up to two weeks (Lauer et al. (2020)) and the large number of asymptomatic carriers (Gudbjartsson et al. (2020)), the virus can remain undetected and spread faster in the general population during the early stage before mitigating strategies can come to full effect.

3. Empirical analysis

3.1. Identification Strategy

Given the large number of factors that influence the spread and toll of the disease, it is difficult to compare across countries. However, policy responses have been rather

similar across Europe resulting in arguably varying degrees of efficiency. This can be to some extent explained by differences in societal and individual behaviour and attitudes rather than different approaches in testing or capabilities of health care systems. Yet, our purpose is to determine whether cultural aspects played a role.

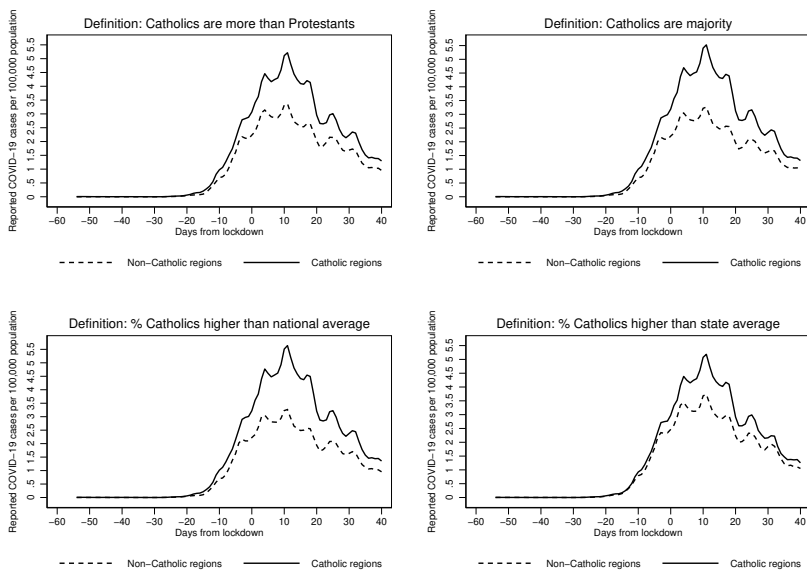
We choose to focus within a single country on a religious divide that has been used in the literature before to explain how cultural differences may affect individual behaviour (Iannaccone (1998); Ekelund et al. (2002); Arruñada (2009); Spenkuch (2017); Spenkuch and Tillmann (2018); Becker and Pascali (2019)). Specifically, we observe that in Germany regions with a higher share of Catholics seem to have been more severely affected by the virus, which could be explained by stronger social and family ties.² For example, when looking Figure A.1 in the Appendix, there is a striking similarity between how the disease has spread and how Catholics are scattered across regions.

Moreover, in Figure 1 the spread of the disease seems to have evolved differently over time in Catholic regions, i.e. where Catholics are the majority, versus non-Catholic ones. We observe a clear discrepancy between Catholic and non-Catholic regions. In robustness checks, this is evident regardless how Catholic regions are defined. More specifically, a NUTS-3 region is defined as a Catholic (vs. a non-Catholic one) if: (a) Catholics are more than Protestants in that region; (b) Catholics are the majority in that region; (c) the share of Catholics in that region is higher than their national average; and (d) the share of Catholics in that region is higher than their share in the respective NUTS-1 area. The spread of the virus at the onset of the pandemic is higher in NUTS-3 regions dominated by Catholics, and this could arguably be the reason behind higher levels of reported COVID-19 cases per 100,000 population in those areas, as well as a higher number of resulting deaths.

To clearly identify this cultural aspect we look as to whether there were any differences between past overall monthly mortality in Catholic and non-Catholic regions similar to a difference-in-difference pre-trends design (Appendix Figure A.2). Time series begin in 2011, considerably after the last major swine flu outbreak in 2009. To our knowledge, no other major threat to public health that can be transmitted via social contacts occurred since then. Both lines are quite close and move in parallel implying that until the COVID-19 pandemic nothing noticeable caused Catholics to pass at a higher rate, relative to non-Catholics. Any systematic differences in genetic predisposition, pre-existing conditions, overall health status, socioeconomic characteristics or behaviours such as risky attitudes would be likely reflected in differences in past mortality rates between groups.

²We choose to focus on West Germany only, as the East part historically exhibits systematic differences and could bias our results (Becker et al. (2020)). Given the very low number of cases/deaths and Catholics it would very likely bias the results in our favour.

Figure 1: Spread of COVID-19 and mobility in Catholic and non-Catholic regions.



Source: Robert Koch Institute; Apple. Horizontal axes are centered at the lockdown date (22 March, 2020).

This would also capture any unobserved heterogeneity that could drive mortality rates. Several demographic and economic characteristics are balanced between Catholic and non-Catholic regions before the pandemic as can be seen in Table A.2 in the Appendix, suggesting that the population we are investigating is rather homogeneous. Some differences are statistically significant, but rather small.³ Moreover, we employ a set of regressions using yearly mortality rates from all causes and our full set of controls for the time period 2011-2017 to formally test the parallel trends assumption at the NUTS-3 level. The results suggest that Catholic regions do not exhibit any differences in mortality in the past using OLS and the Two-step Fixed Effect Estimator by Pesaran and Zhou (2018) (see discussion below for details). This implies that any differences in mortality in 2020 have to be a result of the ongoing pandemic and the spread of COVID-19.

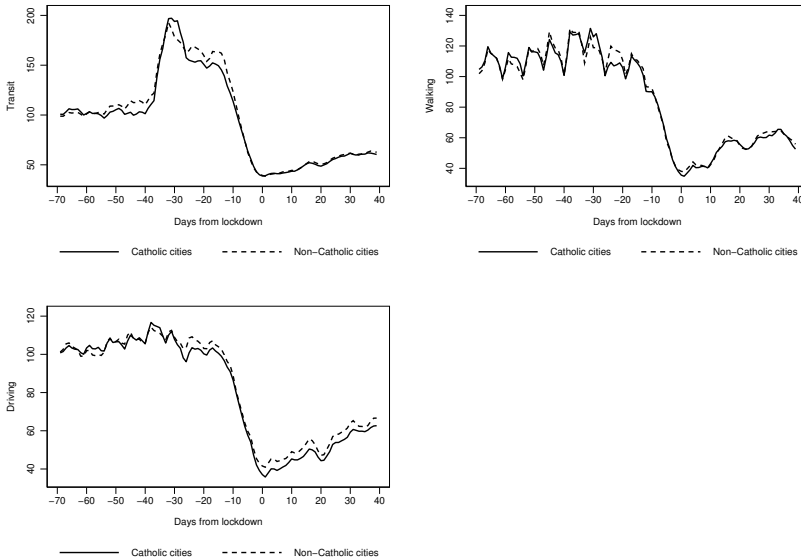
A threat to identification could be that two of the most prominent Catholic states (Bavaria and Baden-Württemberg) are in closer geographical proximity to Northern Italy, which was reportedly the first European area hit the hardest by the pandemic. Even though large airports are the primary entryways into the country and air travel is the largest contributor in spreading infectious diseases (Christidis and Christodoulou (2020)), we also consider land travel as some recent preliminary evidence suggests (Pluemper and Neumayer (2020)). To that end, we use Google Maps data, to control whether driving distance from Milan, Italy, affects our estimates. To further argue about the cultural transmission rather than geographic proximity to Italy, we additionally focus on North Rhine-Westphalia, a state further away from Germany's southern border and one in which there is considerable variation in Catholic and non-Catholic NUTS-3 regions.

Some of the growing COVID-19 literature (Pluemper and Neumayer (2020)) argues that the higher spread in North Rhine-Westphalia is likely due to numerous carnival festivities during February in that state. To address this concern, we look at Apple daily mobility data. These data are obtained through GPS tracking and they are available at the city level (after 13 January 2020, and relative to that date). They provide information regarding requests for directions by transport type, i.e. walking, driving and transit. We classified the 17 cities available for (West) Germany into Catholic and non-Catholic dominated ones and we calculated the mean mobility indicators for each group of cities (weighted by the local population). Trends by transportation type before and after the lockdown (22 March, 2020) are in Figure 2. There is a spike in using public transport about 30 days before the lockdown (upper left Panel), which coincides with

³For example, defining Catholic regions as those where Catholics are the majority, there are no notable differences in characteristics like GDP per capita, share of foreigners, hospital beds per person, and average number of nights spent by tourists. Non-Catholic areas have a higher share of people over 65, more people completed secondary education and they are more densely populated. With respect to mortality, only education would work against us, but again the difference is rather small.

the culmination of carnival festivities in late February. However, this spike is visible not only in Catholic, but also in non-Catholic cities during that period, and the two lines throughout the period are almost indistinguishable from each other, implying that there are no noticeable differences in mobility between Catholic and non-Catholic regions. If the spread of the virus were due to carnival festivities in Catholic regions, we would likely observe substantially higher mobility in these cities compared to non-Catholic ones. It is likely that even though carnival started as a Catholic tradition, it has evolved into a nation-wide celebration for the youth, so that the spike in mobility occurs all over West Germany and thus transmission of an infectious disease is just as likely anywhere in the country.

Figure 2: Spread of COVID-19 and mobility in Catholic and non-Catholic regions.



Source: Robert Koch Institute; Apple. Horizontal axes are centered at the lockdown date (22 March, 2020). Relative Apple mobility volumes are weighted by the local population.

A natural starting point for the analysis is to demonstrate that the reported COVID-19 incidence does vary with religion at the regional level. This is challenging because demographic information at a geographically disaggregated level do not arrive at the same frequency as data on infectious diseases do. Nevertheless, we model the number of COVID-19 cases as follows:

$$Y_{rd} = \alpha Y_{rd-\tau} S_{rd-\tau} + \beta Catholic_r + X_r \gamma + t_r + \lambda_R + \epsilon_{rd} \tag{1}$$

where Y is the number of reported COVID-19 cases in NUTS3 region $r \in \{1, \dots, N\}$ in day d of the outbreak, $Catholic$ is the logged share of Catholics in region r (based on recent data before the virus outbreak), t_r is a linear time trend starting from the day when the first case was reported in each region, λ_R is a set of regional fixed effects (at a more aggregated level so they are not perfectly collinear with the demographic and economic predictors) and ϵ_{rd} is the error term.

The variable S represents the proportion of susceptible individuals in the regional population, where their stock in each region has been approximated as the number of people after removing those reported deceased from the virus in each day (Adda (2016)). The τ parameter represents the incubation period and has been set equal to 14 days, although we results are robust to a wide range of alternative lags (i.e. 3 to 20 days) (Lauer et al. (2020)). We do not consider any spatial variation in the model so the incidence rate in each region is solely determined by its own past realisations, i.e. parameter α should be interpreted as an estimate of the within-region spread. All models control for the size of the local population and for a series for economic and demographic controls, as well as for travelling distance from Milan. Estimating these empirical models will provide an indication of whether the spread of the COVID-19 disease varies with the share of Catholics in the region, conditional on other characteristics, area fixed effects and regional time trends.

A problem with this specification is that regional (at the NUTS-3 level) time-invariant unobserved heterogeneity is not adequately partialled out. Similar to the growing COVID-19 empirical literature, results based on 1 are conditional only to state-level (Länder) fixed effects.⁴ To address this concern, we apply the 2-step method suggested by Pesaran and Zhou (2018) in order to uncover time-invariant effects in a case where N is large and T is small and fixed. More specifically, the predicted residuals from a fixed-effect estimation are averaged for each NUTS-3 region over the entire period. In the second step, they are used as the dependent variable in an OLS regression over the cross-sectional sample of NUTS-3 regions, where models control for the share of Catholics and other local characteristics, e.g. travelling distance from Milan and socio-economic variables. The results are robust even after removing the NUTS-3 level regional fixed effect, indicating a positive relationship between the virus spread and the prevalence of Catholics in the region.

Having demonstrated that the spread is higher in (NUTS-3) Catholic regions, we focus on NUTS-1 regions due to data availability issues discussed in Section 3.3.1 below. We test whether mortality rate is higher in those regions relative to the non-Catholic ones

⁴In fact, most empirical studies report estimates that either condition on broader-level fixed effects or they are unconditional to fixed differences over space e.g., Sa (2020), and Pluemper and Neumayer (2020).

after the pandemic onset. Having established that the pre-pandemic mortality trends were similar across regions (Figure A.2 in the Appendix), we empirically test for this by adopting the following difference-in-differences framework:

$$m_{rt} = \beta_0 + \beta_1 \text{Catholic}_r + \beta_2 \text{Onset}_t + \beta_3 \text{Catholic}_r \times \text{Onset}_t + u_{rt} \quad (2)$$

where m_{rt} is the mortality rate in the r -th region on day t , *Catholic* is an indicator of a predominantly Catholic region and *Onset* is a dummy variable switched on after the first COVID-19 related death in the region was reported. The coefficient of interest in Equation 2 is β_3 and indicates whether differences in mortality rates between Catholic and non-Catholic regions have changed after the onset of the pandemic and as a result of it. Based on our hypothesis about greater transmission of the disease in regions where social ties are stronger, we should expect that $\beta_3 > 0$.

3.2. Data

We combine several data sources. Data on COVID-19 cases come from the Robert Koch Institute (RKI) which is a German federal government research institute responsible for disease control and prevention. It has been publishing validated data on reported COVID-19 cases and related deaths since January 28, 2020. These are daily data by gender, age group and administrative district (412 *Landkreise* in total). In this analysis, we focus only on the 324 West Germany districts, because including East Germany could bias the results (Becker et al. (2020)). In order to be matched to a series of regional characteristics, these COVID-19 series were collapsed by NUTS-3 region and date (running from January 28 to May 01), resulting in a balanced panel of more than 32,000 observations.

Our main controls are taken from the Federal Statistical Office (*Statistisches Bundesamt*) of Germany and are validated until 2017. These include a number of demographic and socioeconomic characteristics at the regional level, e.g. population and population density, share of people over 65 years old, share of foreigners, share of males and females, share of those who completed secondary education, GDP, number of hospital beds, and number of nights spend per person as an indicator for tourism-related activity. Moreover, we use Google Maps data in order to calculate the fastest driving distance between Milan, Italy and the major city of each NUTS-3 region. We also use the 2011 German Census in order to calculate the number of Catholics, Evangelicals and other/no denominations in each region. The share of Catholics relative to the total population in the area will be our main variable of interest.

The individual level data stem from the European Social Survey (ESS) 2018 and the European Values Survey (EVS) 2018, two representative data sets of 1,881 and 4,259

adults, respectively, in West Germany, covering a range of questions around values and social norms. The variables of interest proxy for close family and social ties, i.e. the frequency of social interactions, the perceived frequency compared to individual of the same age, the number of people one is comfortable confiding in, the importance of family, the importance of friends, the level of trust towards family members and whether the respondents reside with their parents or parents-in-law. Further, we use individual and household characteristics as controls. These include the age of the respondent, gender, employment status, education level, subjective health status, age of the youngest household member (EVS only), household size, household income, size of the city/town (EVS only) and NUTS-1 fixed effects. Finally, to check for differences in mobility (by means of transportation), we use Apple daily data on relative mobility volume at the city level.

3.3. Results

3.3.1. Baseline results

Table 1 displays the results from Equation 1 using the daily number of new reported COVID-19 cases per NUTS-3 region as outcome. To address any issues regarding proximity to the Northern Italian border, we used Google Maps to calculate the fastest driving distances from Milan, Italy to the major city of each one of the West German NUTS-3 regions. Although this does not completely rule out any geographic heterogeneity, it mitigates any concerns with respect to geographic dispersion of the pandemic in Europe. In Column 1 we use our full set of controls including travel distance to Milan and state fixed effects and we observe a positive and statistically significant relationship between the local share of Catholics and COVID-19 incidence. The 14-day lagged COVID-19 incidence is as expected a strong predictor of today's count.⁵

The same conclusions hold when the cumulative number of reported COVID-19 cases is considered as the outcome, in Appendix Table A.3. Again, the lagged number of cumulative incidence at the regional level is a strong determinant of the current spread, although to a lesser extent. Hence, the results in Column 4 of 1 should be expected, i.e. a positive relationship between the local share of Catholics and the number of deaths, as well as the cumulative number of deaths (Table A.3 Column 4).

⁵The 14-day lagged number of incidence has been multiplied by a factor that represents the fraction of susceptible individuals in the local (NUTS-3) population. This factor is an approximation, and the local population is calculated as the population minus the cumulative number of COVID-19-related deaths in each day. As such, it does not consider population changes due to local-specific fertility and mortality from other causes, and does not address any endogeneity concerns. Nevertheless, results are robust to the inclusion of various lags, ranging from 3 to 20 days.

Table 1: Number of reported COVID-19 cases/deaths and local Catholics share in West Germany.

	Cases [1]	Cases [2]	Deaths [3]	Deaths [4]
% Catholics	.153*** (.045)	.105* (.058)	.339*** (.098)	.321* (.182)
Lagged cases	.414*** (.016)	.334*** (.024)	.656*** (.044)	.623*** (.090)
Daily trend	Yes	Yes	Yes	Yes
Local controls	Yes	Yes	Yes	Yes
State FE	Yes	No	Yes	Yes
Sample ¹	WG	NRW	WG	NRW
Observations	16,952	2,938	16,952	2,938
NUTS-3 regions	312	53	312	53

Source: Robert Koch Institute (RKI). Gamma regression estimates. Standard errors in parentheses are clustered by NUTS-3 region. ¹WG denotes West Germany and NRW denotes North Rhine-Westphalia.

The three states with most cases and deaths as of April 2020 are Bayern, Baden-Württemberg and North Rhine-Westphalia. The former two are predominantly catholic, 56% and 38% Catholics, respectively, whereas North Rhine-Westphalia is split in the middle. However, the two southern states of Bayern and Baden-Württemberg border on Austria and are geographically closer to Northern Italy, the alleged epicenter of the pandemic in Europe. As this geographical proximity might affect our results, we focus part of the analysis on North Rhine-Westphalia where the share of Catholics is high at 42%, but a substantial share of Evangelicals (28%) also live there. When the estimation sample is restricted to that state, in Column 3 of Table 1, a positive relationship still emerges after controlling for time effects and local demographics. This is also the case when the outcome is the number of deaths, in Column 6 of Table 1. With respect to the cumulative number of cases in Column 3 of A.3, the relationship is not statistically significant, but remains positive and in line with what reported when conditioning on all West Germany regions. However, it is again positive and highly significant when considering the cumulative number of deaths in Column 4.

3.3.2. Difference-in-differences results

A concern could be that the higher death toll of COVID-19 in predominantly Catholic regions was not due to the increased virus spread, but because of unobserved factors

that led to systematically increased mortality rates among Catholics. If that was the case, then our baseline estimates would be simply picking up pre-existing differences in mortality trends. However, we already showed that pre-pandemic mortality trends were almost identical in Catholic and non-Catholic regions over several years before the pandemic onset, in Figure A.2. Therefore, we are confident enough that our baseline estimates regarding the higher death toll among Catholics simply reflect the spread of COVID-19, which our baseline estimates also showed to be significantly higher in regions where Catholics are the majority group.

We then proceed with a difference-in-differences (DiD) estimation using the total number of deaths from all causes across regions. This would allow us to eliminate any unobserved heterogeneity and identify the causal impact of the pandemic on mortality for the two groups. A limitation is that only NUTS-1 level counts are provided by the Federal Statistical Office ⁶. Time series span from 01 January to 17 May, 2020. West Germany NUTS-1 regions are split into predominantly Catholic (treated) and non-Catholic ones (control) based on whether Catholics are the majority. To capture the “after” dimension we consider the death count after the first reported fatality in each region.

Table 2: The impact of COVID-19 on total death count: Difference-in-difference estimates for West Germany.

	Deaths	Deaths	Deaths	Deaths	Deaths
DiD estimate	33.686** (16.790)	60.024*** (21.242)	45.623** (21.825)	18.190*** (3.001)	12.783*** (2.374)
Sample until April, 22	No	Yes	Yes	Yes	Yes
Days since 1st death	No	No	Yes	Yes	Yes
Regional fixed effects	No	No	No	Yes	Yes
Days since 1st case	No	No	No	No	Yes
Observations	1,380	1,090	1,090	1,090	1,090

Note: NUTS-1 regions Schleswig-Holstein, Hamburg, Bremen, Niedersachsen and Hessen classified as non-Catholic. Daily number of deaths from all causes since 1 January. Onset indicates the date of the first reported COVID-19 related death in the state. In the last Column 5 onset indicates the first reported case.

The results are in Column 1 of Table 2, indicating an increased death toll in “treated” NUTS-1 regions during the exposure period. This is confirmed when the estimation sample is limited until April, 22, i.e. when some restrictions were lifted, covering roughly one month since the nation-wide lockdown. This would allow to capture fatalities resulting

⁶Ideally, we would like the series broken down at the NUTS-3 level, but these are not available yet

from the initial stage of the pandemic.⁷ In Column 3 we add a time trend (days since first reported death) to potentially capture the already higher spread in some regions compared to others. Further, we add NUTS-1 fixed effects in Column 4. Finally, we change the onset variable to “first reported case” in Column 5 to capture the already existing high case counts before the first fatalities were reported. In every case, the DiD estimate is positive and highly significant providing further strong evidence that both the virus spread and related fatalities were higher in predominantly Catholic regions as a result of the COVID-19 pandemic.

3.3.3. Regional (NUTS-3) fixed effects

However, a remaining concern could be that our baseline (and DiD) results are conditional only to state-level (NUTS-1) fixed effects. The reason behind this is that controlling for NUTS-3 fixed effects in the baseline models would not allow to obtain estimates for the local Catholics share, as they also vary at the same level. However, this could produce biased coefficients, i.e. in the case where the share of Catholics (and the spread of the virus) are correlated with time-invariant unobserved heterogeneity at the local level. This is a common problem in the emerging COVID-19 empirical literature because local economic and demographic covariates are not available at the daily level and during the period under study.

Therefore, as an attempt to address this issue, we use the fixed-effects filtered (FEF) estimator suggested by Pesaran and Zhou (2018). This allows to uncover the effect of a time-invariant covariate when using a large N , small T panel dataset. More specifically, in the first step the (logged) number of COVID-19 cases (and deaths) was regressed on a set of NUTS-3 fixed effects, past COVID-19 incidence and a regional linear time trend. Then, the residual COVID-19 incidence (and residual deaths) was obtained and averaged over the period for each region. In the second step, the mean residual was regressed on the local Catholics share and other local characteristics using the cross-sectional sample of NUTS-3 regions. The results, in Panel A of Table 3, confirm a positive and significant relationship. Moreover, the estimated coefficients remain positive and significant even after controlling for demographic and economic characteristics and driving distance from Milan. The picture remains largely the same when repeating the same exercise for the number of deaths. Therefore, even after controlling effectively for time-invariant unobserved heterogeneity, our results confirm that the virus spread faster in Catholic regions, relative to non-Catholic ones.⁸ Moreover, given the identical pre-

⁷The results are robust to limiting the time-frame until 31 March which would still capture infections prior to the lockdown

⁸The residual obtained from the fixed effects estimation in the first step, also includes the unexplained time-varying component. The estimations in the second step do not only indicate that the spread in

pandemic parallel mortality trends, the subsequent higher death toll in Catholic regions -also confirmed by the 2-step analysis- was solely due to the higher COVID-19 spread in those areas.

Table 3: Residual of fixed-effects regressions and local Catholics share in West Germany: Estimates before and after the lockdown.

Panel A: Total period:	Cases	Cases	Deaths	Deaths	
% Catholics	.156***	.113***	.019***	.013**	
	(.022)	(.020)	(.004)	(.006)	
R-squared	.133	.589	.043	.151	
Local controls	No	Yes	No	Yes	
Observations	312	312	312	312	
Panel B: Sub-periods (cases):	Before lockdown		After lockdown		
% Catholics	.135***	.088***	.202***	.133**	.109***
	(.029)	(.029)	(.028)	(.026)	(.026)
Cumulative cases per population	-	-	-	-	.147***
					(.037)
R-squared	.059	.565	.130	.580	.605
Local controls	No	Yes	No	Yes	Yes
Observations	312	312	312	312	312

Source: Robert Koch Institute (RKI). The dependent variable is the mean residual of a fixed-effects estimation in the first step. Robust standard errors in parentheses. March 22, 2020 defines the before & after the lockdown.

Moreover, we estimate whether the spread of the virus was higher in two phases of the pandemic in Germany, i.e. before and after the general lockdown implemented on March 22, 2020. The results are in panel B, Table 3. In this case, we estimated the first step residuals using observations only until the implementation of the lockdown. The effect of Catholics share remains positive and significant conditional on regional characteristics. Then, we use the COVID-19 data to estimate residual virus incidence in each region during the lockdown period. The local share of Catholics remains a positive and significant predictor even after controlling for local characteristics and the (log) stock of cumulative cases per 100,000 local population at the date of lockdown implementation. This suggests that even after mobility was restricted, and controlling for the already existing stock of cases from the first phase of the pandemic, COVID-19 transmission remained higher among Catholic regions.

What is also worth noticing is that the results are mainly driven by the older ones, i.e. those above 60 years old. The RKI data are provided by age of each case (or fatality) Catholic regions was higher conditional on NUTS-3 fixed effects, but that it also spread faster over time.

therefore we were able to test for non-linearities conditional on age. The spread and death toll of COVID-19 was much greater for the elderly. This result is confirmed using either a daily panel (Panel A, Table 4) or the two-step estimator by Pesaran and Zhou (2018) to eliminate any unobserved local heterogeneity (Panel B). This is particularly the case when looking the estimates for COVID-19-related mortality where the baseline results are mostly driven by those over 60 years old. Even though non-Catholic regions have a higher share of elderly, the spread and toll of the virus among the more vulnerable increases with the local share of Catholics⁹

Table 4: Local share of Catholics and COVID-19 spread: Estimates by age.

	Cases <60 years old	Cases ≥60 years old	Deaths <60 years old	Deaths ≥60 years old
Panel A: Daily panel estimates				
% Catholics	.133*** (.047)	.235*** (.056)	-.164 (.497)	.394*** .096
Observations	16,759	16,759	16,759	16,759
Panel B: Regional cross-section estimates (2-step)				
% Catholics	.068*** (.014)	.089*** (.018)	.001 (.001)	.018*** (.006)
Observations	312	312	312	312
R-squared	.554	.521	.049	.230

Source: Robert Koch Institute (RKI). Panel A: Gamma regression estimates. Standard errors in parentheses are clustered by NUTS-3 region. All models control for the usual set of local characteristics, time trends and state-level fixed effects. Panel B: OLS estimates. The dependent variable is the mean residual of a fixed-effects estimation in the first step. Robust standard errors in parentheses.

3.3.4. Mechanisms

Having established a positive link between COVID-19 incidence and Catholics share at the local level, a possible mechanism could operate through social norms, i.e. tighter social networks and stronger family ties. Our aim is not to discuss religion or Catholicism in Europe in itself, but rather to explore whether cultural differences within an otherwise very similar population could have impacted groups differently.

Historically, it seems that Protestants and non-Catholics have developed more market oriented and individualistic behaviours (Becker and Pascali (2019)), whereas Catholic regions were characterised by rent-seeking behaviours through established social structures (Ekelund et al. (2002)). Further, Satyanath et al. (2017) argue that Catholics were

⁹See Table A.2 in the Appendix.

frequently organizing their social life through the church and much less through independent organizations, which implies that there was a strong focal point to enhance social activity. Moreover, Arruñada (2009) argues that these differences mostly reflect on social dimensions and ethics rather than economic dimensions or work ethic. Furthermore, they argue that social interaction is regarded by Protestants only as an enforcement and control mechanism over peers and society. This suggests that social interaction as an intrinsic behaviour is more likely to be found among Catholics. For these reasons we believe that Catholics in Germany exhibit stronger social and family ties.

To look deeper into these potential mechanisms, we turn to the 2018 European Social Survey (Wave 9) for Germany. The survey contains questions on an individual level on the frequency of social interactions, on the perceived frequency of such interaction compared to ones peers and also a question on the number of individuals with whom one feels comfortable discussing private and individual matters. We code the first two questions as dummies indicating the frequency of social interactions as more than “several times a week” and the perceived relative measure as “at least the same”. The third variable of interest is the number of individuals respondents confide in. Further, we control for age, gender, subjective health status, employment status, household size, household income and NUTS-1 fixed effects. We also code a dummy for identifying as belonging to the Catholic denomination. The results in Panel A of 5 suggest that there is a positive relationship between belonging to the Catholic denomination and our measures for social interactions. The coefficient for frequency is barely not significant at the 10% level (t -statistic = 1.67), but the other two are, indicating that Catholics tend to have more frequent social interactions and exhibit stronger ties.

An issue with these estimations is that unobserved heterogeneity may bias the results. This is often the case with cross-sectional analyses using survey data. To partially correct for this and obtain unbiased estimates, we follow the methodology proposed by Oster (2019), that has been widely used in the economics literature for this purpose (e.g. Alesina et al. (2016); Satyanath et al. (2017); Tabellini (2020)). The idea is to simultaneously observe the co-movement of the coefficient of interest and the R^2 when including the full set of controls. This allows to approximate the selection on unobservables relative to the selection on observables and obtain an unbiased coefficient. We follow the standard assumptions (also referred to as the “conservative” estimate (Tabellini (2020))) that unobservables matter as much as observables and thus set the parameter δ equal to 1 and define the maximum variance that can be explained as 0.3 greater than the one explained when only controlling for observables, i.e. $R_{max}^2 = 1.3 * R^2$.

The 2018 European Values Survey provides another set of questions that could serve as proxies for these attitudes and behaviours. Here we employ a set of regressions using

individual-level data on the importance of family and friends in an individual's life, the level of trust towards members of the family and the likelihood to reside with parents (or parents-in-law) in the same household. Our variable of interest is a dummy on whether the individual belongs to the Catholic denomination. Controlling for the same set of socioeconomic characteristics and regional fixed effects as in Panel A ¹⁰, the results in Panel B of Table 5 indicate that Catholics, relative to non-Catholics, regard their family and friends as very important, exhibit higher trust towards family members and have a higher probability of residing in the same household with their parents or parents-in-law. This is in line with Arruñada (2009) and suggests that Catholics exhibit much stronger social and family ties, which in turn is a channel through which the pandemic spread more in predominantly Catholic regions.

There are valid concerns around unobserved heterogeneity and a biased Catholic dummy coefficient. We correct for this and obtain the unbiased coefficient (Oster (2019)). For all our estimations in Panel A of Table 5, the unbiased β suggests that if omitted variables matter as much as the ones included, the bias would be against us, as the coefficients are rather stable or even become larger and further away from zero implying that there is a true positive relationship between belonging to the Catholic denomination and the frequency of social interactions. The same applies to the coefficients in Panel B with the exception of Column 3 on Trust in Family. Again, we find a true positive effect between Catholics and social and family ties. Most importantly there is indeed a higher probability for Catholics to live in the same household with their parent or parents-in-law. All of this suggests that Catholics exhibit stronger social and family ties, which arguably could have led to faster spread of the virus in general and to the more vulnerable population in particular in the initial phase of the outbreak.

To further strengthen our proposed mechanism, we employ another set of regressions with the individual-level data. The idea is to determine whether cultural differences could affect the differential outcomes of the outbreak in other ways rather than through social contacts. It could be the case that the outbreak triggered other culturally induced behaviours. Specifically, we test whether Catholics have a higher propensity to justify cheating behaviour in avoiding taxes, accepting bribes and avoiding fares in public transport. Further, we examine their confidence in the government, even though recent research suggest that in most European countries confidence in the governments and the implemented measures have increased (Bol et al. (2020)) and that during the recent pandemic individuals tend to prioritize health (Hargreaves Heap et al. (2020)). This should indicate whether the higher spread among Catholics could be a result of overall riskier

¹⁰Additionally we control for the age of the youngest household member and the size of the city, which are only provided in the EVS

Table 5: Importance of social and family ties

Panel A: Social Interactions		Frequency	Relative	Closeness	
Catholic		0.044 (0.027)	0.056* (0.026)	0.156*** (0.016)	
Unbiased β		0.050	0.056	0.161	
Individual Controls		Yes	Yes	Yes	
Regional FE		Yes	Yes	Yes	
R-squared		0.047			
Observations		1,660	1,660	1,660	
Panel B: Social and family ties		Family	Friends	Trust	Parents
Catholic		0.040*** (0.011)	0.092*** (0.022)	0.056*** (0.017)	0.017*** (0.007)
Unbiased β		0.034	0.099	0.014	0.018
Individual Controls		Yes	Yes	Yes	Yes
Regional FE		Yes	Yes	Yes	Yes
R-squared		0.079	0.055	0.050	0.226
Observations		3,247	3,237	2,338	3,247

Note: OLS regressions using robust standard errors clustered at NUTS-1 level. The data in Panel A are from the European Social Survey 2018 for Germany. The independent variable in the first Column is dummy coded one if the individual meets with others more than "several times a week". In the second Column the dummy takes the value one if the individual perceives themselves to meet more than "about the same" compared to their peers. In the final Column respondents indicate the number of people with whom they "can discuss intimate and private matters". The data in Panel B are from the European Values Survey 2018 for Germany. The independent variables are on a 1-4 scale on the importance of family and friends in ones life and the level of trust towards members of the family, as well as a binary variable on whether one resides with their parents or parents-in-law. Unbiased β refers to Oster (2019) under the assumptions of $\delta = 1$ and $R_{max} = R^2 * 1.3$.

behaviour or groups disregarding the rules and the protective measures put into place. This would likely have a differential impact on the spread and mortality of the disease even before the measures came into place. We empirically test this hypothesis using individual-level German data from the European Values Survey.

Broadly speaking, the data reveal that generally adherence to rules is very high in Germany. Moreover, the regression results in Table 6 do not show a higher propensity for Catholics to disobey rules using the same set of individual controls as in Panel A and regional fixed effects. If anything, Catholics are less likely to tolerate even the milder offense (avoiding fare) and exhibit marginally higher confidence in the government. For this reason we believe that the increased spread is unlikely to be a result of risky

Table 6: Adherence to rules and confidence in government

Adherence to rules	Taxes	Bribe	Fare	Government
Catholic	-0.023 (0.058)	-0.062 (0.064)	-0.159** (0.064)	0.064* (0.034)
Unbiased β	-0.010	-0.051	-0.107	0.062
Individual Controls	Yes	Yes	Yes	Yes
Regional FE	Yes	Yes	Yes	Yes
R-squared	0.024	0.021	0.085	0.081
Observations	2,322	2,333	2,333	2,269

Note: OLS regressions using robust standard errors clustered at NUTS-1 level. The data are from the European Values Survey 2018 for Germany. The independent variables are on a 1-10 scale whether it is justified to avoid taxes, to accept a bribe and to avoid fare in public transport, as well as a 4-point scale indicating confidence in the government. We control for perceived health status, gender, age, household income, education level, employment status, age of the youngest household member and size of the city, as well as state fixed effects. Unbiased β refers to Oster (2019) under the assumptions of $\delta = 1$ and $R_{max}^2 = R^2 * 1.3$.

behaviour and disregard for rules or the government's measures and more related to social and family networks. Moreover, the unbiased estimator as proposed by Oster (2019) reveals that accounting for unobservables under the standard assumptions of $\delta = 1$ and $R_{max}^2 = R^2 * 1.3$ moves the coefficients in the first three columns closer to zero. This further suggests that belonging to the Catholic denomination does not suggest any differences in justifying cheating behaviour and complying with rules. The coefficient in Column 4 remains rather stable suggesting that Catholics exhibit indeed more confidence in the government. We are therefore confident that any observed differences in the number of cases and fatalities is a result of stronger social and family ties and not because of riskier behaviour.

To further rule out other behavioural differences between Catholics and non-Catholics, we look again at the Apple mobility data in Figure 2. There is a significant drop in mobility after the lockdown. However, there is no significant differentiation on the basis of religion. We are therefore confident that the higher transmission rates among Catholics are likely to be a results of close social circles and strong family ties and not of higher mobility before or after the lockdown.¹¹

Finally, the Federal Statistical Office provided commuter mobility data at the NUTS-

¹¹The data that Apple provides refer to searches for directions to a specific destination via certain means of transport. We have no means of formally testing it, but one could plausibly argue that most individuals would not necessarily look up the address of family members and close friends before visiting prior to the lockdown.

3 level gathered by Teralytics for 30 million mobile phones in Germany. The data refer to individuals crossing NUTS-3 borders during peak commuter hours between Monday and Thursday. The data is then provided as the monthly average between January and May 2020, compared to the respective month of the previous year. We use these estimates as independent variables for each month separately as well as for the whole period to determine whether there are any differences in work related mobility between Catholics and non-Catholics that could have affected the spread of the virus. We find no evidence of this as the coefficient for the local Catholics share is insignificant (see Table A.5). This implies that there is no differential change in commuter mobility compared to the previous year, which further strengthens our evidence that the spread and higher mortality among Catholics is more likely due to stronger social and family ties.

4. Discussion

Culture shapes social norms and, to a large extent, dictates societal and individual behaviour. Cultural differences are frequently unobserved and it is difficult to isolate their potential impact (Alesina and Giuliano (2015)). Nevertheless, it is of very high importance to take them under consideration as they can lead to different outcomes. Such heterogeneities need to be taken into account not only in epidemiological modelling and parametrization, but also to design optimal policy responses. Moreover, they do not only exist between, but also within countries and often go beyond standard socioeconomic characteristics.

This is all the more important when facing crises such as the current COVID-19 pandemic, which has already claimed many lives directly or indirectly (Vandoros (2020)). Being able to explain the wildly different trajectories of countries could provide invaluable insight to policymakers to better target mitigation measures and policies around pandemics (Platteau and Verardi (2020)). Previous attempts to highlight the importance of culture and how it can determine social interactions are marred by unobserved heterogeneity and it is difficult to disentangle the number of factors that need to be taken into account. Differences in the timing of the pandemic, health care systems, testing methods to accumulate data and the timing and the nature of policy response, render cross country comparisons rather difficult.

In this paper we used data on daily cases and deaths attributed to COVID-19 in West Germany to overcome some of these issues, to address unobserved heterogeneity and to clearly identify the impact of cultural differences on the spread of the disease and the resulting mortality. Our results, conditional on local characteristics, regional fixed effects, past mortality rates, closeness to the pandemic epicenter in continental Europe,

and mobility patterns, suggest that Catholics that inherently and historically exhibit stronger social and family networks, as is implied by our individual-level analysis, were more severely affected. This is important and indicates that especially in societies with stronger social and family ties (e.g. Spain, Italy), virus outbreaks should be carefully and promptly managed by the authorities in order to avoid rapid spread and, consequently, higher death tolls.

References

- ADDA, J. (2016): “Economic activity and the spread of viral diseases: Evidence from high frequency data,” *Quarterly Journal of Economics*, 131, 891–941.
- ALESINA, A. AND P. GIULIANO (2015): “Culture and Institutions,” *Journal of Economic Literature*, 53, 898–944.
- ALESINA, A., J. HARNOSS, AND H. RAPOPORT (2016): “Birthplace diversity and economic prosperity,” *Journal of Economic Growth*, 21, 101–138.
- ALFARO, L., E. FAIA, L. N., AND S. F. (2020): “Social Interactions in Pandemics: Fear, Altruism and Reciprocity,” *NBER Working Paper*, 27134.
- ARRUÑADA, B. (2009): “Protestants and Catholics: Similar work ethic, different social ethic,” *Economic Journal*, 120, 890–918.
- BARGAIN, O. AND U. AMINJONOV (2020): “Trust and compliance to public health policies in times of COVID-19,” *IZA Discussion Paper Series*.
- BARTSCHER, A., S. S., S. SIEGLOCH, M. SLOTWINSKI, AND N. WEHRHOFER (2020): “The role of social capital in the spread of COVID-19,” *VoxEU*.
- BAYER, C. AND M. KUHN (2020): “Intergenerational ties and case fatality rates: A cross-country analysis,” *IZA Discussion Paper Series*, DP No. 13114.
- BECKER, G. (1996): “Preferences and values,” in *Accounting for Taste*, Gary Becker, ed. Cambridge: Harvard University Press.
- BECKER, S. O., L. MERGELE, AND L. WOESSMANN (2020): “The Separation and Reunification of Germany: Rethinking a Natural Experiment Interpretation of the Enduring Effects of Communism,” *Journal of Economic Perspectives*, 34, 143–171.
- BECKER, S. O. AND L. PASCALI (2019): “Religion, Division of Labor, and Conflict: Anti-Semitism in Germany over 600 Years,” *American Economic Review*, 109, 1764–1804.
- BELLOC, M., P. BUONANNO, F. DRAGO, R. GALBIATI, AND P. PINOTTI (2020): “Cross-country correlation analysis for research on COVID-19,” *VoxEU*.
- BOL, D., M. GIANI, A. BLAIS, AND P. J. LOEWEN (2020): “The effect of COVID-19 lockdowns on political support: Some good news for democracy?” *European Journal of Political Research*.

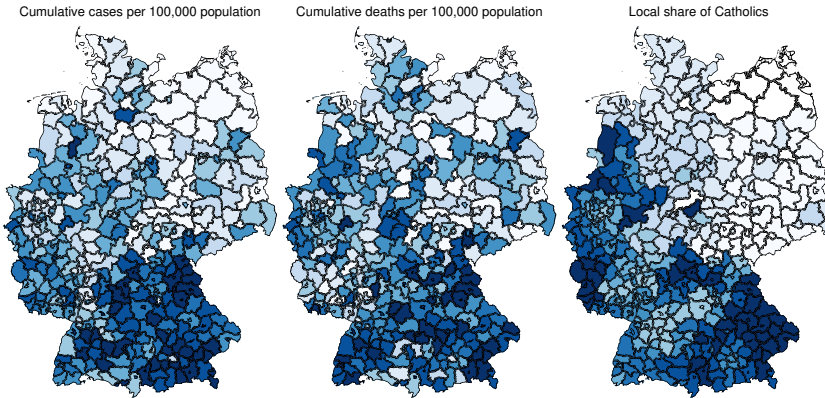
- BRODEUR, A., I. GRIGORYEVA, AND L. KATTAN (2020): “Stay-at-home orders, social distancing and trust,” *Working Paper*.
- CHRISTIDIS, P. AND A. CHRISTODOULOU (2020): “The Predictive Capacity of Air Travel Patterns during the Global Spread of the COVID-19 Pandemic: Risk, Uncertainty and Randomness,” *Int. J. Environ. Res. Public Health*, 17.
- DELANGHE, J., M. SPEECKAERT, AND M. DE BUYZERE (2020): “The host’s angiotensin-converting enzyme polymorphism may explain edipemiological findings in COVID-19 infections,” *Clinica Chimica Acta*, 505, 192–193.
- DEOPA, N. AND P. FORTUNATO (2020): “*Coronagraben*. Culture and social distancing in times of COVID-19,” *COVID Economics, CEPR Press*, 39, 186–216.
- DESLANDES, A., V. BERTI, Y. TANDJAOUI-LAMBOTTE, C. ALLOUI, E. M. CARBONNELLE, J. ZAHAR, S. BRINCHLER, AND Y. COHEN (2020): “SARS-COV-2 was already spreading in France in late December 2019,” *International Journal of Antimicrobial Agents*, ., .
- DOWD, J., V. ROTONDI, L. ANDRIANO, D. BRAZEL, P. BLOCK, X. DING, Y. LIU, AND M. MILLS (2020): “Demographic science aids in understanding the spread and fatality rates of COVID-19,” *Proceedings of the National Academy of Sciences*, 117, 9696–9698.
- DURANTE, R., L. GUISO, AND G. GULINO (2020): “Asocial capital: Civic structure and social distancing during COVID-19,” *Working Paper*.
- EKELUND, R., R. HÉBERT, AND R. TOLLISON (2002): “An Economic Analysis of the Protestant Reformation,” *Journal of Political Economy*, 110, 646–671.
- GLAESER, E. AND S. GLENDON (1998): “Incentives, predestination and free will,” *Economic Inquiry*, 36, 429–443.
- GUDBJARTSSON, D., A. HELGASON, H. JONSSON, O. MAGNUSSON, P. MELSTED, G. NORDDAHL, J. SAEMUNSDOTTIR, A. SIGURDSSON, P. SULEM, A. AGUSTSDOTTIR, B. ERIKSDOTTIR, AND R. E. FRIDRIKSDOTTIR (2020): “Spread of SARS-CoV-2 in the Icelandic Population,” *New England Journal of Medicine*.
- GUISO, L., P. SAPIENZA, AND L. ZINGALES (2006): “Does culture affect economic outcomes?” *Journal of Economic Perspectives*, 20, 23–48.

- HARGREAVES HEAP, S., C. KOOP, K. MATAKOS, A. UNAN, AND N. WEBER (2020): “COVID-19 and People’s Health-Wealth Preferences: Information Effects and Policy Implications,” *SSRN Working paper*.
- IANNACCONE, L. (1998): “Introduction to the Economics of Religion,” *Journal of Economic Literature*, 36, 1465–1495.
- LAUER, S., K. GRANTZ, O. BI, F. JONES, Q. ZHENG, H. MEREDITH, A. AZMAN, N. REICH, AND J. LESSLER (2020): “The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: Estimation and application,” *Annals of Internal Medicine*, ., .
- LIU, Y., Z. GU, S. XIA, B. SHI, X. ZHOU, Y. SHI, AND J. LIU (2020): “What are the underlying transmission patterns of COVID-19 outbreak? An age-specific social contact characterization,” *EClinicalMedicine*, ., .
- MAZZONA, F. (2020): “Cultural differences in COVID-19 spread and policy compliance: Evidence from Switzerland,” *COVID Economics, CEPR Press*, 33, 163–185.
- OSTER, E. (2019): “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business & Economic Statistics*, 37, 187–204.
- PESARAN, M. AND Q. ZHOU (2018): “Estimation of time-invariant effects in static panel data models,” *Econometric Reviews*, 37, 1137–1171.
- PLATTEAU, J.-P. AND V. VERARDI (2020): “How To Exit Covid-19 Lockdowns: Culture Matters,” *COVID Economics, CEPR Press*, 23, 1–57.
- PLUEMPER, T. AND E. NEUMAYER (2020): “The COVID-19 pandemic predominantly hits poor neighborhoods, or does it? Evidence from Germany,” *medRxiv 2020.05.18.20105395*.
- REMLAND, M., T. JONES, AND B. H (1995): “Interpersonal distance, body orientation, and touch: Effects of culture, gender, and age,” *Journal of Social Psychology*, 135, 281–297.
- SA, F. (2020): “Infections and mortality in England and Wales,” *COVID Economics, CEPR Press*, 22, 47–58.
- SATYANATH, S., N. VOIGTLÄNDER, AND H.-J. VOTH (2017): “Bowling for Fascism: Social Capital and the Rise of the Nazi Party,” *Journal of Political Economy*, 125, 478–526.

- SOROKOWSKA, A., P. SOROKOWSKI, P. HILPERT, K. CANTARERO, AND ET AL. (2017): "Preferred Interpersonal Distances: A Global Comparison," *Journal of Cross-Cultural Psychology*, 48, 577–592.
- SPENKUCH, J. AND P. TILLMANN (2018): "Elite Influence? Religion and the Electoral Success of the Nazis," *American Journal of Political Science*, 62, 19–36.
- SPENKUCH, J. L. (2017): "Religion and work: Micro evidence from contemporary Germany," *Journal of Economic Behaviour and Organization*, 135, 193–214.
- STAFFORD, N. (2020): "Covid-19: Why Germany's case fatality rate seems so low," *BMJ*, Opinion.
- TABELLINI, M. (2020): "Gifts of the Immigrants, Woes of the Natives: Lessons from the Age of Mass Migration," *Review of Economic Studies*, 87, 454–486.
- VAN BAVEL, J., K. BAICKE, AND P. T. BOGGIO (2020): "Using Social and Behavioural Science to Support Covid-19 Pandemic Response," *Natural Human Behaviour*, 4, 460–471.
- VANDOROS, S. (2020): "Excess Mortality during the Covid-19 pandemic: Early evidence from England and Wales," *Social Science and Medicine*, 258.
- ZHAO, S., Q. LIN, J. RAN, S. MUSA, G. YANG, W. WANG, Y. LOU, D. GAO, L. YANG, D. HE, AND M. WANG (2020): "Preliminary Estimation of the Basic Reproduction Number of Novel Coronavirus (2019-nCoV) in China, From 2019 to 2020: A Data-Driven Analysis in the Early Phase of the Outbreak," *Int J Infect Dis*, 92, 214–217.

Appendix

Figure A.1: COVID-19 spread and local share of Catholics at the NUTS-3 level.




Source: Robert Koch Institute (RKI); Federal Statistical Office of Germany. Notes: Data on COVID-19 spread are as at April 20, 2020. Data on the local share of Catholics refer to 2017.

Table A.1: Mortality rate in Catholic and non-Catholic regions, 2011-2017

	Deaths [1]	Residual [2]	Deaths [3]	Residual [4]
% Catholics	-.008 (.055)	-.142 (.156)	-	-
Catholic dummy	-	-	-.058 (.063)	-.061 (.063)
Local controls	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
R-squared	.521	.518	.522	.518
Observations	2,195	2,195	2,196	2,196
NUTS-3 regions	317	317	318	318

Source: Federal Statistical Office. Standard errors in parentheses are clustered by NUTS-3 region.

Figure A.2: Mortality rate in Catholic and non-Catholic regions, 2011-2017.



Source: Federal Statistical Office of Germany. Mortality rate is defined as the number of deaths per 100,000 regional population. Catholic regions are defined as those where Catholics are the majority.

Table A.2: Means of control variables

	Catholic			Non-Catholic			p-value
	Mean	SD	n	Mean	SD	n	
GDP	11.159	.134	162	11.166	.157	161	.681
Population	11.922	.584	162	12.076	.718	162	.035
Pop. Den.	5.554	.929	162	5.962	1.159	162	.000
Share over 65	3.035	.091	162	3.071	.106	162	.001
Foreigners	2.310	.353	162	2.374	.466	162	.167
Comp. Sec.	3.375	.303	160	3.489	.254	162	.000
Hospital beds	1.589	.705	161	1.687	.587	159	.181
Nights spent	.883	.2896	160	.890	.313	156	.835

Note: All variables have been transformed into natural logarithms. Catholic regions are defined as regions where Catholics are the majority.

Table A.3: Cumulative number of reported COVID-19 cases/deaths and local Catholics share in West Germany.

	Cases [1]	Cases [2]	Deaths [3]	Deaths [4]
% Catholics	.146*** (.045)	.111 (.085)	.399*** (.111)	.458* (.266)
Lagged cases	.368*** (.016)	.317*** (.032)	.520*** (.027)	.482*** (.059)
Daily trend	Yes	Yes	Yes	Yes
Local controls	Yes	Yes	Yes	Yes
State FE	Yes	No	Yes	No
Sample	WG	NRW	WG	NRW
Observations	16,952	2,938	16,952	2,938
NUTS-3 regions	312	53	312	53

Source: Robert Koch Institute (RKI). Gamma regression estimates. Standard errors in parentheses are clustered by NUTS-3 region. WG denotes West Germany and NRW denotes North Rhine-Westphalia.

Table A.4: Means of control variables in ESS and EVS

Panel A: ESS	Catholic			Non-Catholic			p-value
	Mean	SD	n	Mean	SD	n	
Age	52.677	19.062	521	48.784	18.969	1,833	.000
Gender	.481	.500	521	.487	.499	1,837	.826
Health	2.299	.878	521	2.316	.898	1,836	.692
Education	4.317	1.748	521	4.257	1.753	1,828	.489
Employment	.923	.266	521	.914	.279	1,836	.520
HH Size	2.641	1.235	521	2.564	1.291	1,835	.225
HH Income	6.441	2.707	458	5.968	2.834	1,630	.001
Panel B: EVS	Catholic			Non-Catholic			p-value
	Mean	SD	n	Mean	SD	n	
Age	52.492	17.791	1,306	50.564	17.606	2,906	.001
Gender	.525	.499	1,313	.490	.500	2,928	.036
Health	2.695	.844	1,305	2.684	.890	2,889	.695
Education	4.246	1.782	1,308	4.365	1.841	2,894	.051
Employment	.615	.486	1,322	.618	.485	2,937	.839
HH Size	2.557	1.239	1,299	2.541	1.281	2,884	.716
HH Income	5.802	2.754	1,147	5.804	2.879	2,599	.983
Age Youngest	37.758	25.552	1,168	35.932	24.979	2,626	.039
City Size	2.649	1.174	1,322	2.901	1.252	2,937	.000

Note: ESS and EVS data for West Germany, 2018. p-values refer to t-testing. Wording of questions and coding of variables differs between surveys.

Table A.5: Falsification test for Commuter Mobility in West Germany, Jan-May 2020

	Mobility [1]	Mobility [2]	Mobility [3]	Mobility [4]	Mobility [5]	Mobility [6]
% Catholics	.009 (.545)	-1.003 (.604)	.859 (.673)	-.207 (.846)	.285 (.901)	-.011 (.566)
Month	Jan	Feb	Mar	Apr	May	All
Month FE	No	No	No	No	No	Yes
Local controls	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	.158	.158	.416	.498	.466	.871
Observations	310	310	310	310	310	1,550

Source: OLS regression estimates. Standard errors in parentheses are clustered by NUTS-3 region. Catholic is a dummy on whether Catholics are a majority in the region.

Growth forecasts and the Covid-19 recession they convey¹

Javier G. Gómez-Pineda²

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

With the help of growth forecasts and a simple structural model, we build a likely forward-looking account of the depth, length and shape of the recession as well as of the demand and supply shocks that are driving it. The results point to a V-shaped recession with partial recovery in advanced economies and to an L-shaped recession in emerging and developing economies. In addition, the projected shapes likely involve, in advanced economies, an output level shock, and in emerging and developing economies, an output growth shock. The depth and shape of the recession in output is important for fiscal debt sustainability analysis; in this matter the results are robust to the model parameters and assumptions. In turn, the depth and length of the recession in the output gap is critical for monetary and fiscal policies; in this matter we had to appeal to an assumption about the extent of the demand shock. The simple structural model does not have the problem of univariate filters that can misleadingly attribute to demand shocks a large part of the variability of output that is actually originated in supply shocks.

- 1 This chapter is a short version of the paper “The depth, length and shape of the recession conveyed in mid-2020 growth forecasts,” Banco de la República (the central bank of Colombia), Borradores de Economía, 2020, forthcoming. The author thanks Sofía Salamanca and David López for excellent research assistance. The findings, recommendations and interpretations expressed in this paper are those of the author and do not necessarily reflect the view of the Banco de la República or its Board of Directors.
- 2 Senior Economist (Investigador Principal) at Banco de la República and part time professor (profesor de cátedra) at Universidad del Rosario.

Copyright: Javier G. Gómez-Pineda

1. Introduction

Growth forecasts can help articulate a forward-looking account or story of the depth, length, shape of the covid-19 recession as well as of the shocks that are implicit, with the help of a structural model. Using a simple structural model of the decomposition of output between potential output and the output gap, as well as between demand and supply shocks, the purpose of the paper is to try to uncover the forward-looking story of the unfolding covid-9 recession that is implicit in mid-2020 growth forecasts.

The structural model is a simple decomposition of output between the output gap and potential output. The former explained by demand shocks; the later, by supply shocks. More elaborated versions of the model would include a set of Phillips curves and policy rules to close a standard new-Keynesian model. In this light, the simple structural model in the paper is the real block of a standard model in the New Neoclassical Synthesis (NNS) tradition.

A standard approach is to use a structural model and the demand and supply shocks to forecast a projected path of output. Examples are Mckibbin and Fernando (2020) and Stannard et al. (2020). Counter to this method, we use the structural model and the growth forecasts to infer backwards from the forecasts to uncover the underlying demand and supply shocks. The outcome is a likely story of the projected covid-19 recession in terms of the depth, length and shape of the projected covid-19 recession as well as of the extent, type and mix of the underlying demand and supply shocks.

The chapter has five sections including this introduction. Section 2 explains the structural model. Section 3 describes the data. Section 4 describes the results. Section 5 presents the conclusions and policy implications.¹

¹ Detail about the calibration, estimation of the model as well as the robustness exercises are available in Gómez-Pineda (2020a).

2. The model

The model is a decomposition of output into output gap and potential output, with processes for the output gap and potential output driven by demand and supply shocks.

Output Y_t is split between the output gap \hat{Y}_t and potential output \bar{Y}_t as follows:

$$Y_t = \hat{Y}_t + \bar{Y}_t. \quad (1)$$

The output gap is a stationary, autoregressive process driven by a demand, or output gap shock $\varepsilon_t^{\hat{Y}}$. The output gap shock aggregates a variety of demand shocks, such as to confidence (global uncertainty and risk aversion),² government expenditure, an income and employment-related shock and an income and the terms of trade-related shock. The output gap equation is

$$\hat{Y}_t = \alpha \hat{Y}_{t-1} + \varepsilon_t^{\hat{Y}}. \quad (2)$$

In turn, potential output is a nonstationary process driven by a supply output level shock $\varepsilon_t^{\bar{Y}}$ and by the potential-output growth rate γ_t as follows:

$$\bar{Y}_t = \bar{Y}_{t-1} + \gamma_t + \varepsilon_t^{\bar{Y}}. \quad (3)$$

The potential-output growth rate $\bar{Y}_t - \bar{Y}_{t-1}$ is nonstationary as the shock $\varepsilon_t^{\bar{Y}}$ has permanent effects on the output level. In contrast with the measure of the potential-output growth rate $\bar{Y}_t - \bar{Y}_{t-1}$, the measure γ_t is stationary as it converges to the long-term, potential-output growth rate γ as follows:

$$\gamma_t = \theta \gamma_{t-1} + (1 - \theta) \gamma + \varepsilon_t^{\gamma}. \quad (4)$$

The stationary potential-output growth rate γ_t is driven by the supply output growth shock ε_t^{γ} .

We add detrended output Y_t^{Det} as the metric to gauge the depth and shape of the recession. It is defined as the sum of the output gap and detrended potential output \bar{Y}_t^{Det} as follows:

$$Y_t^{Det} \equiv \hat{Y}_t + \bar{Y}_t^{Det}, \quad (5)$$

where detrended potential output \bar{Y}_t^{Det} is equal to potential output minus trend potential output \bar{Y}_t^{Trend} ,

$$\bar{Y}_t^{Det} \equiv \bar{Y}_t - \bar{Y}_t^{Trend}, \quad (6)$$

² For a paper that incorporates global uncertainty and risk aversion as drivers of the output gap see Gómez-Pineda (2020b).

and trend potential output is the hypothetical path of potential output were all supply shocks zero, as in

$$, (7)$$

where γ is the nonstochastic potential-output growth rate that follows:

$$\gamma. (8)$$

In the absence of shocks, equations 7 and 8 mean that y_t is simply a time trend. In addition, detrended output is normalized to zero in the base year, for our purposes $t=0$, so that trend potential output is equal to potential output in that base year,

In equation 5, detrended output is the sum of the output gap and detrended potential output, or in other terms, $(y_t - \hat{y}_t) = (y_t - \hat{y}_t) + (\hat{y}_t - \hat{y}_t)$. Trend output rises at the pace of the long-term, potential-output growth rate. As potential output is not stationary, it drifts away from trend output as permanent shocks affect its level. Detrended potential output cannot be the subject matter of demand policies such as monetary and fiscal policies; in contrast, the output gap is an important input in the formulation of monetary and fiscal policies; in addition, it is also affected by them. Detrended potential output, in turn, is the outcome of the containment, social-distancing and de-escalation policies implemented to deal with the covid-19 pandemic.

Finally, a definition of trend output can be provided as the sum of the output gap and trend potential output

$$. (9)$$

Using this definition, detrended potential output can be calculated either as in equation 6 or as

The simple structural model in equations 1 to 4 can be extended in several ways. On the demand side, the output gap equation can be enhanced to include confidence variables such as global uncertainty and global risk aversion. On the supply side, the potential output block in equations 2 and 3 can be augmented using information on containment measures such as quarantines, establishments and school closures and restrictions on local and international transportation, an example is Stannard, Steven, and McDonald (2020) and data is available in Thomas et al. (2020).

Both the output gap and detrended potential output can be obtained as the sum of current and past demand and supply shocks, respectively. Importantly, with a stationary process, as the one that defines the output gap in equation 2, demand shocks have effects that are transitory. It then makes sense to talk about the length of the recession in demand or in the output gap because in the absence of shocks the output gap eventually converges to zero. The story is different when considering supply shocks. As equation 3 is not stationary, supply shocks have permanent effects on output. In this light, there cannot be a length of the recession in detrended potential output. A similar rationale applies to detrended output. Along this line, supply shocks have permanent effects on output so there cannot be a length of the recession in detrended output.

In order to build the combined shock, we made identifying assumptions about the relative demand and supply shocks. Using equations 1 to 4, the model shocks can be written as a function of the growth forecast and information available at time t as follows:

$$\varepsilon_t^{\bar{Y}} + \varepsilon_t^{\bar{Y}} + \varepsilon_t^{\bar{Y}} = (Y_t - Y_{t-1}) + k_t, \quad (13)$$

where $Y_t - Y_{t-1}$ is the growth forecast and $k_t = (1 - \alpha)\hat{Y}_{t-1} - \theta\gamma_{t-1} - (1 - \theta)\gamma$ is the information available at time t .

According to equation 13, a given growth forecast is consistent with multiple combinations of supply and demand shocks. We then make assumptions about the relative demand and supply shocks to find the model shocks as

$$\varepsilon_t^{\bar{Y}} = \frac{\eta^{\bar{Y}}}{(1 - \eta^{\bar{Y}})(1 - \eta^{\bar{Y}})} \eta [(Y_t - Y_t) + k_t], \quad (14)$$

$$\varepsilon_t^{\bar{Y}} = \frac{\eta^{\bar{Y}}}{(1 - \eta^{\bar{Y}})} \eta [(Y_t - Y_t) + k_t], \quad (15)$$

$$\varepsilon_t^{\bar{Y}} = \eta [(Y_t - Y_t) + k_t], \quad (16)$$

where parameter $\eta^{\bar{Y}}$ is the relative demand shock, defined as $\eta^{\bar{Y}} \equiv \varepsilon_t^{\bar{Y}} / (\varepsilon_t^{\bar{Y}} + \varepsilon_t^{\bar{Y}} + \varepsilon_t^{\bar{Y}})$; parameter $\eta^{\bar{Y}}$ is the relative supply level shock, defined as $\eta^{\bar{Y}} \equiv \varepsilon_t^{\bar{Y}} / (\varepsilon_t^{\bar{Y}} + \varepsilon_t^{\bar{Y}})$; and $\eta = (1 - \eta^{\bar{Y}})(1 - \eta^{\bar{Y}})[\eta^{\bar{Y}}\eta^{\bar{Y}} + \eta^{\bar{Y}}(1 - \eta^{\bar{Y}}) + \eta^{\bar{Y}}(1 - \eta^{\bar{Y}}) + (1 - \eta^{\bar{Y}})(1 - \eta^{\bar{Y}})]^{-1}$.

The identification of the combined shock is as follows. The output gap takes the restriction given by equation 14 on impact, afterwards, it follows equation 2. In turn, supply shocks can be obtained either by making supply level shocks endogenous to the growth forecast while supply

growth shock follow equation 16. Alternatively, supply output growth shocks are made endogenous to the growth forecast and supply level shock follow equation 15.

These identifying assumptions about the combined shock can enable us to build a forward-looking story about how the covid-19 recession is to unfold. Such a story is, in principle, relevant for monetary policy, at least because the depth of the recession in the output gap is an input in Taylor rules and in models with forecast rules. It is also relevant for fiscal policy as the depth of the recession in the output gap is important for estimating fiscal cyclical revenue as well as the structural balance, if fiscal rules are at all binding during the covid-19 recession. Furthermore, the projected recession in detrended output should be relevant for overall fiscal revenues. The identifying assumptions cannot be true; they are only assumptions subject to model uncertainty. In turn, growth forecasts are not true, they are subject to additive uncertainty.

In the future, as more data becomes available and also with the benefit of hindsight, a story about the role of supply and demand shocks in the recession can be estimated; however, that would only be a historical account of the covid-19 recession. The research strategy in the paper is to make use of the identifying assumptions about the relative demand shock and supply level shocks to find the implicit forward looking story about the projected recession conveyed in mid-2020 growth forecasts. We then analyze the robustness of the results to the identifying assumptions.

3. The data

Given the large amount of uncertainty currently surrounding the growth forecast we use yearly data. Growth forecasts at yearly frequency are available for a forecasting horizon of five years from source Focus Economics.³ The sample of economies includes the economies available in to us in the Focus Economics database; that is, 65 economies, 29 advanced and 36 emerging and developing. In 2019, the economies in the sample accounted for 83.5 percent of world output, evaluated at PPP exchange rates, 38.7 for advanced economies, 44.8 emerging and developing economies and 25.6 for emerging and developing economies excluding China.^{4 5}

³ In contrast with yearly forecasts, quarterly growth forecasts are typically available for a 2-year forecast horizon.

⁴ The advanced economies in the sample are Australia, Austria, Belgium, Canada, Cyprus, Estonia, Finland, France, Germany, Greece, Hong Kong SAR, Ireland, Italy, Japan, Korea, Latvia, Lithuania, Luxembourg, Netherlands, New

In the Focus Economics database, available growth forecasts for each economy are the median of a number of panelists. The panelists include organizations such as investment banks, universities, research institutions and one credit rating agency. The number of panelists in each economy ranges from 7 to 55, with a median of 22. The panelists are interviewed monthly for a total of 12 forecasting vintages each year.⁶

Data for the output gap is available for most of the advanced economies in the World Economic Outlook (WEO) database, October 2019. For the 6 advanced economies without output gap data in the WEO, October 2019 database as well as for all the emerging and developing economies in the sample, the output gap was estimated.⁷

4. Results

The first result is about the depth of the recession. Current growth forecasts convey recessions of -9.5 and -6.7 percentage points in advanced and emerging and developing economies, respectively. At the end of the forecasting horizon the depth of the recession is smaller in advanced economies.

Zealand, Portugal, Singapore, Slovak Republic, Slovenia, Spain, Switzerland, Taiwan Province of China, United Kingdom and United States. The emerging and developing economies in the sample are Argentina, Bangladesh, Belize, Bolivia, Brazil, Brunei Darussalam, Cambodia, Chile, China, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Haiti, Honduras, India, Indonesia, Jamaica, Lao P.D.R., Malaysia, Mexico, Mongolia, Nicaragua, Pakistan, Panama, Paraguay, Peru, Philippines, Russia, Sri Lanka, Thailand, Trinidad and Tobago, Uruguay and Vietnam.

⁵ We used all the economies available in our Focus Economics service but excluded Puerto Rico because it was not in the IMF database for the period of the global financial recession and Myanmar and Malta because these countries are small. Also, we did not include the countries in the Middle East, North Africa and Sub-Saharan Africa because they are not in our available Focus Economics service.


⁶ Each forecast vintage contains information up to the end of the last month. This information is available at the beginning of each month. For example, the July forecast vintage has information up to end-June. In contrast, for the countries in Latin America and the Caribbean each forecast vintage has information available up to the middle of the month. This information is available at the middle of the month. In the example, the July forecast vintage as information up to mid-July and is available by mid-July. We use as time convention the information that is available at the beginning of the month. Hence, for countries in Latin America and the Caribbean we used the information available up to the middle of the earlier month. In the example, in Latin America and the Caribbean the July forecast vintage uses information available up to mid-June. The 2020 growth forecasts vintages available at the time of writing the paper are those from January through July.

⁷ Output level data for the period 1995–2019 was constructed using output growth data from source the World Economic Outlook database, April 2020. (Log) output level data was constructed accumulating growth rates; that is, using $Y_t = Y_{t-1} + g_t$, where Y_t is log output and g_t is output growth. The percent growth figures were transformed into logarithmic growth with the expression $g_t = 100 \times \log(1 + G_t/100)$, where G_t is percent growth.

As for the depth of the recession across different forecasting vintages, in advanced economies the depth of the recession is zero in the March 2020 forecast vintage and gradually increases to about -9.5 percentage points in the July forecast vintage (Figure 1). In emerging and developing economies the depth of the recession is about null in the March 2020 forecast vintage. Then it gradually increases to about -6.7 percentage points in the July forecast vintage (Figure 1). The revision in the depth of the recession can be quite large from month to month so further revisions can be expected during 2020.

Figure 1. The projected depth and shape of the covid-19 recession

Median detrended output in the Focus Economics March through July 2020 forecast vintages



We now deal with the distribution of the depth of the recession, in Figure 2. The interquartile range of the depth of the recession is 2.7 percentage points in advanced economies. At the end of the forecasting horizon this range narrows to 1.6 percentage points.⁸ In emerging and developing economies, the interquartile range of the depth of the recession is 4.5 percentage points. In turn, at the end of the forecasting horizon it narrows to 3.4 percentage points although it broadens when accounting for the quartile distribution.

The second result is about the length of the recession in the output gap. Figure 3 shows the quartile distribution of the projected output gap. If we take the length of the recession as the time necessary for the output gap to shrink by 4, the length of the recession is about 4 years for the standard economies accounted for within the interquartile range.

⁸ These figures are in approximate percent. Besides, the graphs depict numbers in logarithmic growth.

Figure 2. The projected depth and shape of the covid-19 recession

Quartile distribution of detrended output in the Focus Economics July 2020 forecast vintage


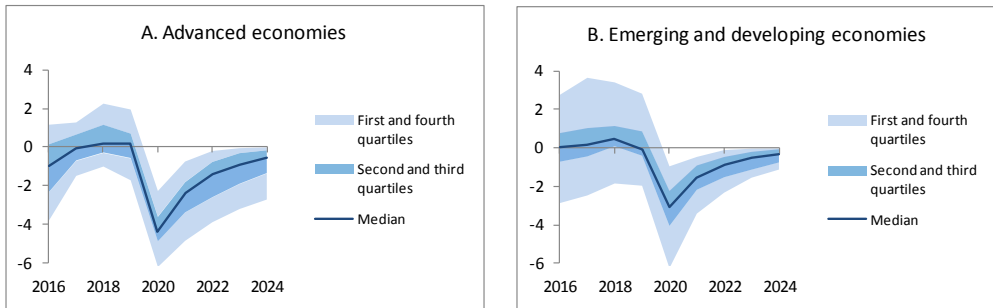


Figure 3. The assumed depth and estimated length of the covid-19 recession

Quartile distribution of the output gap in the Focus Economics July 2020 forecast vintage

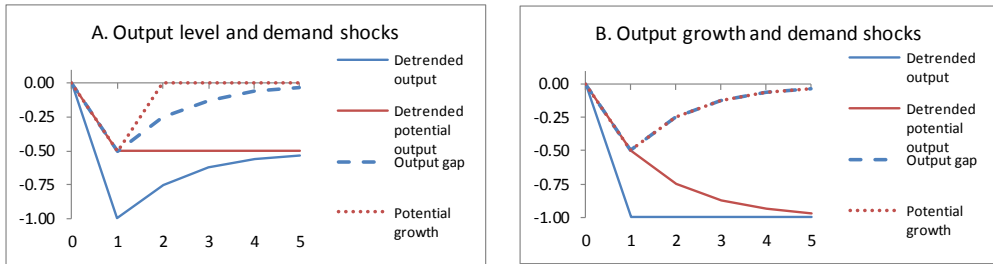


The third result is about the shape of the recession, presented in Figures 1 and 2. The recession is V-shaped with partial recovery in advanced economies and L-shaped in emerging and developing economies. Across forecasting vintages, the shape can vary as the forecast are revised; nonetheless, Figure 1 shows that the current projected shape emerged in May forecast vintage.

The fourth result is that the recession can better be characterized in advanced economies as the result of an output level shock and in emerging and developing economies as a result of an output growth shock. Figure 4 presents the response to a combined demand and supply shock amounting to 1 percentage point of potential output. In Panel A the relative demand shock is $\frac{1}{2}$ while the supply shock is an output level shock, that is, the relative supply shock is 1. In Panel B the relative demand shock is $\frac{1}{2}$ while the supply shock is an output growth shock, that is, the relative supply shock is 0.

When the supply shock is a level shock, Panel A, the recession in detrended output is V-shaped with partial recovery. Recall that, from equation 5, detrended output is equal to detrended potential output plus the output gap. The V-shaped recession with partial recovery is the result of a V-shaped recession in the output gap, the dashed blue line; plus an L-shaped recession in detrended potential output, the solid red line. In turn, the L-shaped recession in detrended potential output is the result of a shock to equation 2 in period 1. From period 2 onwards detrended potential output is constant as the potential-output growth rate returns to zero immediately after the shock, the red dotted line.

Figure 4. Response to a combined supply and demand shock

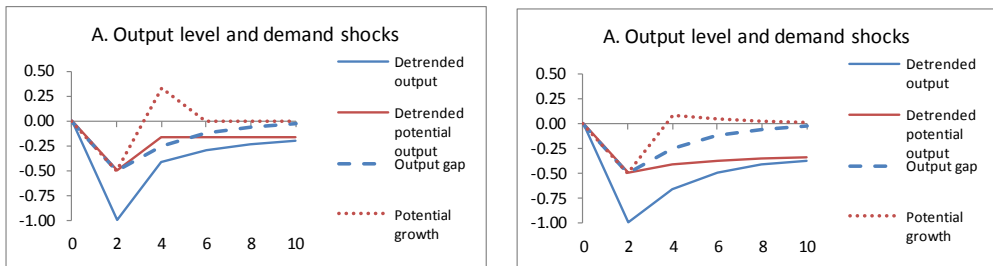


When the supply shock is an output growth shock, Panel B, detrended output follows an L-shaped recession. Again, note that according to equation 5 detrended output is equal to detrended potential output plus the output gap. The L-shaped recession is the result of a V-shaped recession of the output gap, the dashed blue line, plus a curved-shaped recession in detrended potential output, the solid red line. In turn, the curved-shaped recession in detrended potential output is the result of a shock to equation 3 in period 1. From period 2 onwards, the potential-output growth rate begins a gradual convergence to zero, the red dotted line, diving below zero throughout the recession and so pulling detrended potential output along a curved-shaped recession. Detrended potential output converges to -1 as fast as the output gap converges to zero. While detrended output and the output gap diverge, they add up to 1 , an L-shaped recession in detrended output ensues.⁹

⁹ The recession in detrended output is L-shaped if the persistence of the potential-output growth rate is maintained in 0.5 and output gap persistence is 0.5 . If output gap persistence raises above 0.5 , the L-shaped recession is concave upward. If output gap persistence decreases below 0.5 , the L-shaped recession is convex upwards.

In the impulse responses considered, the scarring effects in advanced economies are twice as large as those in emerging and developing economies. The scarring is the sequel or permanent effect of supply shocks and can be measured by detrended potential output. Figure 4 shows that the scarring is the same on impact but twice as large five years after the shock when the supply shock is an output growth shock. As detrended potential output is not stationary, the scarring diverges with new supply shocks. For incomplete scarring an upward supply shock must follow the downward impact shock. Figure 5 presents the response to a combined demand and supply shock with incomplete scarring; that is, adding a subsequent supply shock with opposite sign. To obtain a scarring of $1/3$ the extent of the supply shock in period 1 is $-2/3$ times the supply shock in period 0. Panel A shows incomplete scarring with output level shocks; in turn, Panel B shows incomplete scarring with output growth shocks. The result is that the scarring is $1/3$ by construction. On one hand, when supply shocks are output level shocks, in Figure 5, Panel A, the scarring is $-1/2 \times 1/3 = -1/12$, compared with -0.5 in Figure 4 without scarring. On the other hand, when the supply shocks are output growth shocks, in Figure 5, Panel B, the scarring is $-1/3$, compared with -1 in Figure 4, without scarring. Furthermore, no matter the type of supply shock, whether output level or growth shock, both detrended potential output and detrended output undergo a V-shaped recession with partial recovery.

Figure 5. Response to a combined demand and supply shock with incomplete scarring effects




We now turn to the important issue of the shock identification strategy. A look at the impulse responses in Figure 4 and at the shape of the recessions in Figure 2 reveals that, taking the growth forecasts as given, and with the help of the model, supply shocks may better be characterized as output level shocks in advanced economies and output growth shocks in

emerging and developing economies.¹⁰ We then turn this observation into the assumption that in advanced economies supply shocks are output level shocks while in emerging and developing economies they are output growth shocks. The story is easy to understand: a one-time output level shock in advanced economies and a one-time output growth shock in emerging and developing economies.

As mentioned above, we work backwards from the growth forecasts to the shocks that may give rise to them. Using the assumption about the type of shock by type of economy, and using the growth forecasts and the backward induction method, the implicit supply shocks appear in Figure 6. The implicit output level shocks in advanced economies, in Panel A, are not entirely stylized or V-shaped as those that give raise to the responses in Figure 4. Indeed, small and decreasing supply shocks, opposite in sign to the impact shock appear since 2021 indicating incomplete scarring effects, although incipient. Still, an almost stylized V-shaped path of the output level shock arises in the July forecast vintage. The implicit output growth shocks in emerging and developing economies, in Panel B, are not entirely stylized and V-shaped either, but a V-shaped path of the output growth shock arises starting the May forecast vintage, although with a small positive output growth shock in 2021 indicating somewhat incipient scarring effects.

Figure 6. The projected output level and output growth shocks in the covid-19 recession
Median output level and growth shocks in the Focus Economics March through July 2020 forecast vintages




The distribution of the implicit supply shocks appears in Figure 7. The interquartile range, including the more standard economies, shows the almost stylized V-shaped path with a small shock in 2021 in both advanced and emerging and developing economies.

¹⁰ In other terms, the relative supply level shock in advanced economies is 1 and in emerging and developing economies is 0.

Covid Economics 40, 30 July 2020: 211-228

Figure 7. The projected output level and output growth shocks in the covid-19 recession
 Quartile distribution of output level and growth shocks in the Focus Economics July 2020 forecast vintage



In conclusion, the implicit forward-looking story seems to be one of a combined demand and output level supply shock with almost complete scarring in advanced economies and a combined demand and output growth supply shock with almost complete scarring in emerging and developing economies.

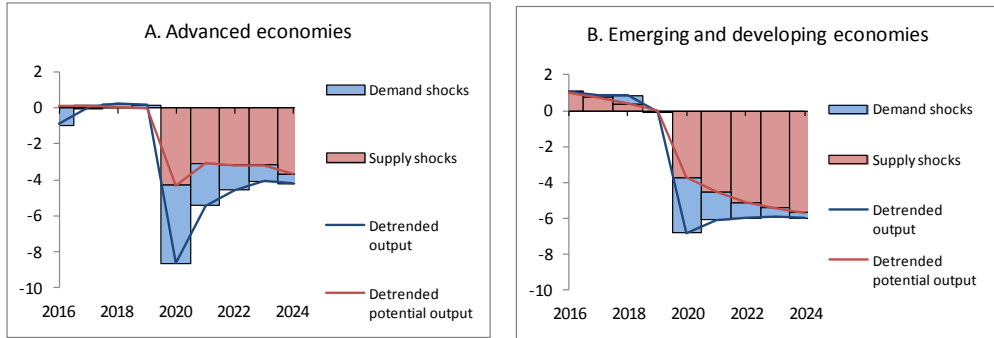
We now turn to the shock decomposition of the recession. The shock decomposition of median detrended output in the latest available, July 2020 forecasting vintage appears in Figure 8. In the shock decomposition the recession in detrended output is explained by cumulative demand and supply shocks, the recession in the output gap by cumulative demand shocks and the recession in detrended potential output by cumulative supply shocks. Explained by cumulative supply shocks, the recession in detrended potential output is V-shaped with partial recovery and curved-shaped in emerging and developing economies. In turn, explained by cumulative demand and supply shocks, the recession in detrended output is V-shaped with partial recovery in advanced economies and L-shaped in emerging and developing economies.

In summary, contained in current growth forecasts are recessions of about -9.5 and -6.7 percentage points in advanced and emerging and developing economies, respectively. The recession in detrended output is V-shaped with partial recovery in advanced economies and L-shaped in emerging and developing economies, distributed in a narrower and narrowing band in advanced economies and along a broader and broadening range in emerging and developing economies. The recession can be better characterized as the result of an output level shock in advanced economies and as a result of an output growth shock in emerging and developing economies.

Covid Economics 40, 30 July 2020: 211-228

Figure 8. The assumed demand and supply shocks during the covid-19 recession

Shock decomposition of median detrended output in the Focus Economics July 2020 forecast vintage



The results are derived from the growth forecasts and the model; they do not necessarily mean that the growth forecasts can be attached high probabilities, particularly given the unprecedented level of uncertainty currently surrounding the growth forecasts.¹¹

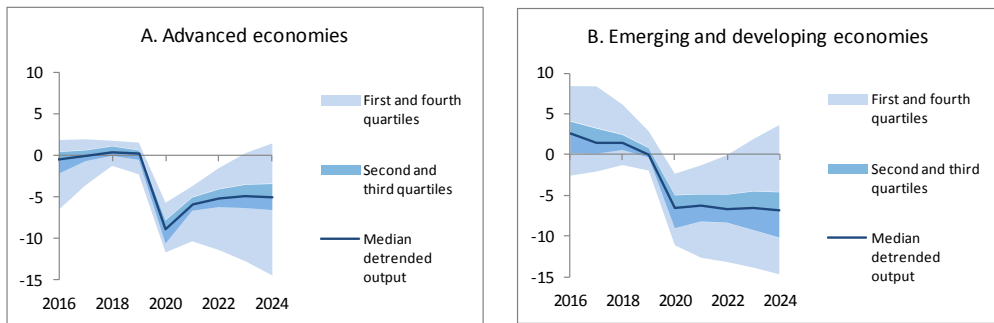
The results are robust to changes in the long-term, potential-output growth rate for standard economies within the interquartile range. We have assumed the long-term, potential-output growth rate is equal to the forecasted-output growth rate at the end of the forecast horizon, in 2024. Consider, alternatively, that the long-term, potential-output growth rate is the average of 2010–2019. The new parameter can change detrended potential output by the end of the forecasting horizon particularly in nonstandard economies outside interquartile range.

The results about the shape of the recession are robust for the median long-term, potential-output growth rate. In advanced economies, the median long-term, potential-output growth rate raises 0.4 percentage points to 2.4 percent, from 2.0 in the base case. As a result, at the end of the forecasting horizon, in 2024, detrended potential output drops 1.0 percentage point to -4.4 percent, from -3.4 in the base case. In emerging and developing economies, median long-term, potential-output growth rate raises 0.6 percentage points to 3.9 percent, from 3.3 in the base case. As a result, at the end of the forecasting horizon detrended potential output drops 0.6 percentage points to -7.0 percent, from -6.4 in the base case.

¹¹ In addition, Gómez, Guillaume and Tankeri (20) show that a semi-structural model can produce better forecasts than those of the analysts.

The dispersion of the results is robust within the interquartile range (Figure 9). In advanced economies the interquartile range increases to 3.5 percentage points, from 1.5 in the base case. In emerging and developing economies, it increases to 5.9 percentage points, from 4.3 in the base case.

Figure 9. Robustness: Change in the long-term, potential-output growth rate
Quartile distribution of detrended output in the Focus Economics July 2020 forecast vintage



In contrast, the results are not robust to changes in the long-term, potential-output growth rate in countries outside the interquartile range. In these non-standard countries, the results are not robust because these economies had highly stylized rates of growth during 2010–2019.¹² All in all, the results are robust to the calibration of the long-term, potential-output growth rate for standard economies within the interquartile range.

Another robustness concern is whether the depth of the recession is invariant to different assumptions about the relative demand and supply shocks. The depth of the recession in detrended output does not depend on the relative demand shock and so the results about the recession in detrended output are robust to this assumption. In contrast, the results about the depth of the recession in the output gap and detrended potential output are naturally not robust; these results depend on the identifying assumption about the relative demand shock. In this light, research on the effect of containment and de-escalation measures on the potential output block of the model can help in the estimation of the output gap.

¹² Examples of countries outside the interquartile range are Greece, Portugal, Iceland and Singapore. Output growth in these countries exhibits a highly idiosyncratic behavior during the estimation period 2010–2019.

5. Conclusions

We use growth forecasts and a simple structural model to build backwards from the growth forecasts to the implicit forward-looking story about the extent of the demand and supply shocks likely explaining the unfolding covid-19 recession.

The forward looking story is as follows: the recession in advanced economies –9.2-percent deep, 4-years long and V-shaped with partial recovery, with almost complete scarring effects and the result of combined demand and supply shocks including output level shocks. The recession in emerging and developing economies is –6.5 percent deep, 4-years long and L-shaped, also with almost complete scarring effects and the result of combined shocks likely including an output growth shock.

The results are derived from the growth forecasts and the model; they do not necessarily mean that the growth forecasts are not surrounded with a great deal of uncertainty, particularly given the unprecedented level of uncertainty during the unfolding covid-19 recession.

The results of the paper about the depth and shape of the recession in detrended output are robust to changes in the long-term, potential-output growth rate, for standard economies within the interquartile range. The results of the paper about the depth of the recession in the output gap are naturally the consequence of the identifying assumption about the relative demand shock.

In the policy implications, a forward looking estimate of the depth of the recession in output and in the output gap is critical for monetary and fiscal policies, yet there is plenty of uncertainty about the depth of the recession, particularly in the output gap. Concerning monetary policy, Taylor and forecast rules require an estimate of the output gap. The former because the output gap enters the rule, the later because the inflation forecast depends on the output gap, among other factors. As for fiscal policy, debt sustainability analysis requires output growth projections and the structural balance also needs a forward looking projection of the output gap. The former depends critically on the projected shape of the recession, on this matter we have used the growth forecasts; the later depends on the projection of the output gap, on this matter we have used the structural model and an identifying assumption about the relative demand shock. In this light, an important research agenda for the estimation of the output gap is the enhancement of the potential output block of the model with information on containment

measures such as quarantines, establishments and school closures and restrictions on local and international transportation.

Detrended potential output cannot be the subject matter of demand policies only the output gap can be. Nonetheless, detrended output and output itself do depend on other set of policies, namely, the containment and de-escalation policies that have been implemented to deal with the pandemic. In the future, as more output data is available, the depth, the length and shape of the recession can be estimated more easily; however, that would be a historical account of the recession, not a forward-looking story that can inform demand policies at the time they have to be formulated. The research strategy in the paper is to make use of identifying assumptions about the relative demand and supply level shocks to find the implicit forward-looking story contained in the current growth forecasts. We then analyze the robustness of the results to these identifying assumptions and find the results robust except, naturally, for the assumption about the extent of the demand shock. Further research on the effect of containment and de-escalation policies on potential output can shed light on the estimation of the output gap.

Also in the policy implications, structural models such as the one used here can incorporate a separate role for supply and demand shocks. Nonetheless, nonstructural univariate filters give a smooth path of potential output thereby misleadingly attributing to demand shocks a large part of output variability that is actually originated in supply shocks. In this light, structural estimations of the output gap can be helpful in the formulation of demand policies such as monetary and fiscal policies.

References

- Carabenciov, I., Freedman, C., García-Saltos, R., Laxton, D., Kamenik, O., & Manchev, P. “GPM6--The Global Projection Model with 6 Regions.” *IMF Working Paper* No 13/87. Washington DC. 2013.
- Gómez-Pineda, J. G. “The depth, length and shape of the covid-19 recession conveyed in mid-2020 growth forecasts,” Banco de la República, *Borradores de Economía*, 2020a, forthcoming.
- Gómez-Pineda, J. G. “Volatility spillovers and the global financial cycle across economies : Evidence from a global semi-structural model.” *Economic Modelling*, 90, pp. 331–373. 2020b.

- Gómez-Pineda, J. G., Guillaume, D., & Tanyeri, K. “Systemic risk, aggregate demand, and commodity prices.” *IMF Working Paper* No. 15/85. Washington DC. 2015.
- Mckibbin, W., & Fernando, R. “The Global Macroeconomic Impacts of COVID-19: Seven scenarios.” *CAMA Working Paper*, No. 19/2020. 2020.
- Stannard, T., Steven, G., & McDonald, C. “Economic impacts of COVID-19 containment measures.” *Reserve Bank of New Zealand Analytical Note*, No. 4. 2020.
- Thomas, H., N. Angrist, B. Kira, A. Patheric T. Phillips, S. Webster, “Variation in government response to covid-19,” *Blavatnik School of Government Working Paper* No. 2020/032. 2020.

School disruption and pupil academic outcomes - evidence from the 2001 foot and mouth disease epidemic in England

Will Cook¹

Date submitted: 22 July 2020; Date accepted: 23 July 2020

The Covid-19 crisis has led to disruption to schooling across the world. Though it is recognized that pupils are suffering immediate learning loss, there exists a lack of understanding as to how this disruption might affect longer-term educational outcomes. This study considers this issue by examining the effect of school disruption in England due to restrictions put in place to manage the Foot and Mouth Disease epidemic in cattle in 2001. Using a difference in difference approach, I analyze whether primary schools that had been significantly disrupted by the epidemic experienced lower performance in standardized tests for pupils aged 11 in English, maths and science in the year of the outbreak and in subsequent years. I find that primary schools that had been significantly disrupted by the measures to contain the epidemic exhibited achievement falls in the year immediately after the outbreak, driven by sizeable falls in maths performance. The negative effects weaken in subsequent years suggesting that the effects of school disruption had faded out to some extent by the time that cohorts that were younger at the time of exposure took the age 11 tests.

¹ Senior Lecturer, Future Economies Research Centre, Manchester Metropolitan University.

1 Introduction

1.1 Background and related literature

In response to the Covid-19 pandemic, schools have been fully or partially closed in an effort to control the spread of the virus. This pause in education provision and the reliance on schooling at home has led to concerns about the long-term effect of foregone learning and the potential for these effects to exacerbate existing educational inequalities. Evidence from the UK and USA suggests that many pupils are not engaging in their schools' efforts to maintain education provision (Lucas et al, 2020; Cullinane and Montacute, 2020) and that there exists a substantial difference in home learning between the highest and lowest income households (Chetty, 2020; Andrew et al, 2020; Anders et al, 2020). However, current studies of the impact of Covid-19 on learning can only measure the short-term impacts of school disruption on learning *inputs*; the long-term impact on pupil attainment is a matter of conjecture. Burgess and Sievertsen (2020) provide an estimate of the learning loss due to missing 12 weeks of schooling of up 10% of an S.D. unit, based on studies that have estimated the effect of varying instruction time in schools (Carlsson et al, 2015; Lavy, 2015).

There are a couple of historical parallels with the current situation of widespread school disruption due to pandemics. Meyers and Thomasson (2017) find that likely school disruption due to measures to control the 1916 polio pandemic in the USA resulted in pupils obtaining less years of education compared to individuals who had left school in the years just prior to the pandemic. Goulas and Megalokonomou (2020) study the impact of relaxing school attendance requirements in Greece in response to the H1N1 swine flu epidemic. They find that when school attendance rules are relaxed it is high attaining pupils that tend to attend school less and their academic performance suffers.

This paper aims to provide further evidence to aid the assessment of long-term effects of the disruption to schooling due to Covid-19. It does so by estimating the effect of school disruption related to school closures, pupil movement restrictions and the psychological distress of pupils that was caused by measures to contain an epidemic in cattle, the 2001 foot and mouth disease (FMD) outbreak in the UK. The modelling estimates the effect of the outbreak on schools' test score

performance on standardized tests for 11 year olds in maths, English and science in the year of, and those following, the crisis. It is important to note that I am not able to estimate whether pupils were affected beyond the age 11. However, examining these test scores across cohorts allows an assessment of whether any effects of disruption are present in the immediate aftermath of the crisis and whether these effects are still detectable in cohorts who take the age 11 tests sometime after the disruption had occurred.

1.2 The FMD Outbreak

FMD is a highly infectious viral disease that causes blistering on the hooves and the mouth of cattle and other livestock, affecting movement and feeding; it is potentially fatal to young cows. There is however no significant threat to human health. The FMD epidemic in the UK lasted from February 2001 until the last case was declared in September 2001, although measures to contain the disease lasted in some places into 2002. A number of measures were implemented to reduce the transmission of the virus, including widespread culling of livestock, movement restrictions on livestock and the closure of countryside rights of way, and, of relevance to this study, the closure of schools and other public institutions. Even when schools were open, pupils and staff from some locations were advised or simply chose not to travel for fear of spreading the virus¹; in some cases, pupils were off school for months and schools resorted to delivering education via post.

Along with the effect on school pupils from missing school, through closure or otherwise, the FMD outbreak also had serious psychological effects on farming communities that lost livestock due to the mass culling and burning of the animal carcasses (Mort et al, 2001). The widespread closure of the countryside and the images from the outbreak also had the effect of deterring tourism from the affected areas, which further compounded the economic effect of the outbreak. Taken together, it is hypothesized that the disruption to school opening and attendance, along with the stressful environment that the outbreak created in the affected communities affected the learning of pupils living in the affected areas.

¹ The official report of the Cumbria FMD task force provides details of the restrictions imposed and the effect on local communities in the UK's worst affected area:

<https://www.cumbria.gov.uk/eLibrary/Content/Internet/538/716/37826163827.pdf>

2 Data & Method

2.1 Data description

The data used is publically available school performance data covering all mainstream primary schools in England from 1997/1998 – 2005/2006. For this period, pupils took standardized tests (“KS2 tests”) in English, maths and science² during the May of the school year in which they turn 11 years old. These tests are used for published school performances measures and for setting the baseline for pupil progress over secondary school. During the period in question, the main school performance measure was the percentage of pupils reaching the ‘expected level’ for each subject and the average of these percentages over the three subjects. These measures are used as the outcome measures in this study. Schools were identified by their local education authority as being ‘significantly affected’ by FMD³ in order that this could be reported alongside their school accountability measures published in 2001. Unfortunately data on the exact nature of the disruption does not exist, however school absence data indicates that the number of days that pupils were absent in affected schools was 30% higher in 2001 compared to the years after the outbreak⁴, in contrast to 6% higher in control schools. Given that the outbreak affected just a third of the academic year this suggests that the FMD outbreak had a substantial effect on pupil absence.

The sample is restricted to schools that had a full set of attainment data for this period⁵, this leaves 52 FMD affected schools and 11,024 control schools. Descriptive statistics are shown in table 1. These show that FMD affected schools had, on average higher attainment prior to the outbreak and that this was maintained post crisis. In order to control for these and other⁶ pre-existing differences, the estimate of the effect of the outbreak is made using a differences in differences

² Science tests were discontinued after 2009.

³ This was decided upon by the Local Education Authority with guidance from central government that schools should be designated as significantly affected by FMD if schools performance is likely to have been affected ‘because pupils have been absent for prolonged periods or because they have been affected by the stress and trauma surrounding the disease’ DfES (2002).

⁴ Authors analysis; school absence data was only collected from 2001 so a before/after comparison cannot be made.

⁵ The main reason for schools having incomplete data is that school performance data is suppressed when the number of pupils taking the tests is less than ten, for privacy reasons.

⁶ E.g. FMD affected schools are more likely to be located in less densely populated areas.

design including school fixed effects. In addition, I use leads in an event study analysis to examine whether estimated effects may be driven by pre-existing trends.

Table 1. Descriptive statistics

		FMD affected schools		Control Schools	
		Before	After	Before	After
<i>Age 11- proportion reaching expected level</i>					
3-Subject Average	<i>Mean</i>	0.793	0.854	0.730	0.804
	<i>SD</i>	0.131	0.090	0.155	0.121
Maths	<i>Mean</i>	0.754	0.806	0.683	0.752
	<i>SD</i>	0.159	0.126	0.174	0.146
English	<i>Mean</i>	0.779	0.832	0.716	0.781
	<i>SD</i>	0.141	0.109	0.161	0.139
Science	<i>Mean</i>	0.847	0.924	0.790	0.880
	<i>SD</i>	0.129	0.073	0.162	0.109
Observations		156	312	33,072	66,144

Note: 'Before' refers to observations prior to 2001; 'After' refers to observations from 2001 onwards.

2.2 Empirical approach

To estimate the effect of the disruption caused by the FMD outbreak on affected cohorts a simple difference in difference model is estimated:

$$y_{st} = \alpha_s + \tau_t + \gamma FMD_{st} + \epsilon_{st} \quad (1)$$

Where y is the percentage of pupils passing the expected level of attainment for English, Mathematics or Science at age 11 and the average over the three subjects, in school s for year t . α_s are school fixed effects and τ_t are a set of time fixed effects that control for time invariant school characteristics and common cohort shocks respectively; there are no explanatory variables in the model due to data availability. FMD is a dummy variable indicating if a school was identified as being significantly affected by FMD from 2001 onwards; γ is the difference in difference estimate of the FMD school disruption averaged out over all affected cohorts. However as the outbreak occurred at a specific point in time it may be the case that the effect of the outbreak was most apparent for those cohorts that took the tests during or immediately after the crisis. At the same time, while younger cohorts may have been affected by the outbreak, by the time they came to take their age 11 tests, they may have managed to ‘catch-up’ any lost learning. To test this idea, an event study framework is used:

$$y_{st} = \alpha_s + \tau_t + \sum_{t=-2}^{t=6} \gamma_t (FMD * T)_{st} + \epsilon_{st} \quad (2)$$

Where FMD*T represents the interaction between a dummy variable for a school being identified by its local education authority as being significantly affected by FMD and a set of time dummies, with the year before the outbreak, 2000, as the reference category (i.e. $t=0$). The set of coefficients γ_t provide estimates of the effect of the FMD disruption for $t>0$ for each cohort as each cohort passes through the end of primary school tests. For $t<0$, γ provides a test of the parallel trends assumption. All results are weighted by the number of pupils in each school-cohort⁷, and standard

⁷ In practice this makes little difference to the results.

errors are clustered at the level of the school. The models are estimated for each outcome: maths, English and science, and the average over the three subjects).

3 Results

The simple difference in difference estimates as per (1) are shown in table 2. A negative effect of the FMD disruption is found over each of the outcomes. For the 3-subject average the estimated effect is a reduction in the proportion of pupils reaching the expected level on 1 percentage point (approx. 10% of a SD), with the largest estimated effect being on maths performance and the smallest on science. These estimates are however not statistically significant.

Table 2. Difference in difference estimates

<i>Outcome</i>	
3-Subject Average	-0.01 (0.01)
Maths	-0.015 (0.011)
English	-0.01 (0.011)
Science	-0.004 (0.013)
Observations	99331

Notes: The table reports estimates of γ for each of the outcome measures. Standard errors clustered by school reported in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Regressions are weighted by the number of pupils in each school-cohort.

The estimates from the event study estimation are shown in table 3, with corresponding graphical results in figures 1-4. Over the three subjects, average there is a sharp and statistically significant fall in results for FMD affected schools in the year after the crisis, 2002, with results recovering somewhat afterwards. The estimates and figure 1 suggest that after the initial shock in 2002, the performance of subsequent cohorts was 1 percentage point lower than it would have been, though these estimated effects are not statistically significant. The estimates for the treatment leads (1998 and 1999) are very close to zero indicating that pre-existing trends are not driving the results and that the parallel trends assumption is reasonable.

Results in all subjects exhibit a fall in 2002; however, the main effect appears to be on maths, where the effect of the FMD outbreak is estimated to be a reduction of 5.4 percentage points (34% of an SD) in 2002 with negative but declining effects estimated for subsequent cohorts. Across all outcome measures, it is clear that the FMD outbreak did not appear to have any effect on test performance in 2001, i.e. when test were taken during the outbreak.

Table 3. Difference in difference estimates by individual cohort (event study)

	3-Subject Average	Maths	English	Science
1998*FMD	-0.003	-0.008	-0.02	0.019
(pre)	-0.017	-0.022	-0.02	-0.02
1999*FMD	-0.001	-0.027	0.012	0.011
(pre)	-0.014	-0.017	-0.016	-0.02
2001*FMD	-0.002	-0.009	-0.007	0.009
(post)	-0.013	-0.018	-0.017	-0.01
2002*FMD	-0.027 **	-0.054 ***	-0.018	-0.01
(post)	-0.012	-0.015	-0.016	-0.01
2003*FMD	-0.01	-0.032 **	-0.017	0.018
(post)	-0.012	-0.016	-0.018	-0.02
2004*FMD	-0.008	-0.026	-0.01	0.011
(post)	-0.014	-0.016	-0.021	-0.02
2005*FMD	-0.009	-0.025	-0.008	0.007
(post)	-0.015	-0.017	-0.018	-0.02
2006*FMD	-0.01	-0.011	-0.015	-0
(post)	-0.016	-0.017	-0.02	-0.02
Observations	99,331	99,331	99,331	99,331

Notes: The table reports estimates of γ_t from (2) for each of the outcome measures. Standard errors clustered by school reported in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Regressions are weighted by the number of pupils in each school-cohort.

Figure 1. Average of 3 subjects (maths, English, science): event study estimates




Figure 2. Maths: event study estimates




Figure 3. English: event study estimates





Figure 4. Science: event study estimates



4 Conclusion

4.1 Summary of results

The analysis presented here provides evidence that disruption to schools due to the FMD outbreak affected the test performance of affected schools. The main result is that affected schools experienced a drop in performance in the year after the outbreak (i.e. 2002), with some indication that the effect of the outbreak persisted into following years. These estimates for later years are however smaller and not statistically significant and suggests that those who were younger during the outbreak were either less affected or managed to catch up in their learning by the time they took the age 11 tests. That the effect is most apparent for the cohort that took the tests in 2002 is consistent with evidence that pupils make the most learning gains between the ages of 9 and 11 in English primary schools (Gray et al, 2004) and therefore would be most affected by school disruption. The overall fall in performance appears to be driven by the effects on maths performance. A potential reason for this is that attainment in mathematics may be more reliant on pre-requisite knowledge obtained in earlier schooling compared to English and science, which means that the effect of learning loss will be more persistent and that pupils are less likely to catch up. A puzzling result is the lack of any effect detected in the year of the outbreak (2001). This may be because the 2001 tests occurred near the beginning of the outbreak and thus the disruption to learning was limited. It may also be because pupils that were particularly affected did not take the tests – something that cannot be tested with the data at hand.

4.2 Policy implications

There are two tentative policy implications from this work in terms of the current Covid-19 crisis. The first is that pupil attainment is likely to be affected by the school lockdowns, particularly in mathematics, and as such remedial provision will be required to make up for lost learning. This effect however may fade out over time and therefore remedial interventions should be carefully targeted and evaluated rather than as a blanket response to missed schooling. The results suggest that pupils who are taking high stakes assessment in the next future should be prioritized for support. The second policy implication is that school accountability measures will be affected by the Covid-19 crisis even when schools return to normal and these results provide some support for

the argument that test-based accountability requirements may need to be relaxed next year to recognize the disruptive effect of the crisis.

4.3 Limitations and further work

The main limitations of these results are twofold. The first is that the analysis is based on publicly available school level data. This means that the analysis is limited in terms of choice of research design, control variables and tests for heterogeneity of effects between different types of pupils and schools. The restriction of the sample to schools with more than ten pupils in their test-taking cohort means that small schools (mainly in rural areas) are not included in this analysis. Given their rurality, it is potentially these schools that were most affected by the FMD outbreak and so the effects of the outbreak may be underestimated. On the other hand, the results may overestimate the effect of the outbreak if it induced differential effects on pupil migration such that higher attaining pupils were more likely to leave affected schools after the outbreak. The second limitation is that the treatment variable, the indicator of whether a school was significantly affected by FMD, was a subjective designation by local government and provides little information as to the nature and intensity of the disruption for each individual school. Future research could usefully exploit pupil level datasets⁸ to solve these issues. This may include the matching of pupil level records to infected farm addresses to get a clearer idea of the effect of the outbreak on affected families.

⁸ The access to such datasets is currently challenging due to temporary closure of research facilities that allow access to this data.

5 References

Andrew, A., Cattan, S., Costa-Dias, M., Farquharson, C., Kraftman, L., Krutikova, S., Phimister, A., Sevilla, A. (2020) Learning during the lockdown: real-time data on children's experiences during home learning, *IFS Briefing Note BN288*

(Available at <https://www.ifs.org.uk/publications/14848>)

Anders, J., Macmillan, L., Sturgis, P. & Wyness, G. (2020) Homeschooling during lockdown deepens inequality, *LSE Covid-19 Blog*

(Available at <https://blogs.lse.ac.uk/covid19/2020/06/05/homeschooling-during-lockdown-will-deepen-inequality/>)

Burgess, S. & Sievertsen, H. (2020) The long-term consequences of missing a term of school, *IZA Commentary*

(Available at <https://wol.iza.org/opinions/the-long-term-consequences-of-missing-a-term-of-school>)

Carlsson, M., Dahl, G., Öckert, B., Rooth, D. (2015). The Effect of Schooling on Cognitive Skills, *The Review of Economics and Statistics*, MIT Press, vol. 97(3), pages 533-547

Chetty, R. (2020) The Economic Impacts of COVID-19 Real-Time Evidence From Private Sector Data, *Princeton University seminar*

(Available at <https://bcf.princeton.edu/wp-content/uploads/2020/06/chetty.pdf>)

Cullinane, C. and Montacute, R. (2020) Covid-19 and social mobility, *Sutton Trust Research Brief*

(Available at <https://www.suttontrust.com/wp-content/uploads/2020/04/COVID-19-Impact-Brief-School-Shutdown.pdf>)

DfES (2002) Implications of foot & mouth disease for schools, DfES guidance note,

(Available at

<https://webarchive.nationalarchives.gov.uk/20020601225953/http://www.dfes.gov.uk:80/footandmouth/#perform>)

Gray, J., Schagen, I., & Charles, M. (2004) Tracking pupil progress from Key Stage 1 to Key Stage 2: how much do the 'route' taken and the primary school attended matter?, *Research Papers in Education*, 19:4, 389-413

Goulas, S. & Megalokonomou, E. (2020) School attendance during a pandemic, *Economics Letters*, Volume 193

Lavy, V. (2015), Do Differences in Schools' Instruction Time Explain International Achievement Gaps? Evidence from Developed and Developing Countries. *Economic Journal*, 125: F397-F424.

Lucas, M., Nelson, J. and Sims, D. (2020). Schools' Responses to Covid-19: Pupil Engagement in Remote Learning. Slough: NFER.

(Available at

https://www.nfer.ac.uk/media/4073/schools_responses_to_covid_19_pupil_engagement_in_remote_learning.pdf)

Meyers, K. & Thomasson, M. (2017) Paralyzed by Panic: Measuring the Effect of School Closures during the 1916 Polio Pandemic on Educational Attainment, *NBER working paper 23890*

(Available at <https://www.nber.org/papers/w23890.pdf>)

Mort, M., Convery, I., Baxter, J. & Bailey, C. (2005) . Psychosocial effects of the 2001 UK foot and mouth disease epidemic in a rural population: qualitative diary based study, *BMJ* 2005; 331 :1234