


**CENTRE FOR
ECONOMIC
POLICY
RESEARCH**

CEPR PRESS



COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

ISSUE 23
28 MAY 2020

SOCIAL STRUCTURE MATTERS

Jean-Philippe Platteau and
Vincenzo Verardi

WHO LOOKS AFTER THE KIDS?

Almudena Sevilla and Sarah Smith

HOW TO RESTART THE UK ECONOMY?

Anton Pichler, Marco Pangallo, R. Maria
del Rio-Chanona, François Lafond and J.
Doyne Farmer

OPTIMAL POLICIES

Stefan Pollinger

HOUSEHOLD EXPECTATIONS

Tobin Hanspal, Annika Weber and
Johannes Wohlfart

**FLATTENING THE CURVE:
WHAT WORKS?**

Nikos Askitas, Konstantinos Tatsiramos
and Bertrand Verheyden

**PREVENTIVE BEHAVIOUR AND
MENTAL CONDITIONS IN JAPAN**

Eiji Yamamura and Yoshiro Tsutsui

**LOCKDOWN EFFECTS AROUND
THE WORLD**

Jean-Philippe Bonardi, Quentin Gallea,
Dimtriya Kalanoski and Rafael Lalive

Covid Economics

Vetted and Real-Time Papers

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

Founder: Beatrice Weder di Mauro, President of CEPR

Editor: Charles Wyplosz, Graduate Institute Geneva and CEPR

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

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

The Centre for Economic Policy Research (CEPR)

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

Chair of the Board

Sir Charlie Bean

Founder and Honorary President

Richard Portes

President

Beatrice Weder di Mauro

Vice Presidents

Maristella Botticini

Ugo Panizza

Philippe Martin

Hélène Rey

Chief Executive Officer

Tessa Ogden

Editorial Board

Beatrice Weder di Mauro, CEPR

Charles Wyplosz, Graduate Institute
Geneva and CEPR

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

Abi Adams-Prassl, University of Oxford
and CEPR

Jérôme Adda, Bocconi University and
CEPR

Guido Alfani, Bocconi University and
CEPR

Franklin Allen, Imperial College Business
School and CEPR

Oriana Bandiera, London School of
Economics and CEPR

Michele Belot, European University
Institute and CEPR

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

Nick Bloom, Stanford University and
CEPR

Tito Boeri, Bocconi University and CEPR

Alison Booth, University of Essex and
CEPR

Markus K Brunnermeier, Princeton
University and CEPR

Michael C Burda, Humboldt Universitaet
zu Berlin and CEPR

Luis Cabral, New York University and
CEPR

Paola Conconi, ECARES, Universite
Libre de Bruxelles and CEPR

Giancarlo Corsetti, University of
Cambridge and CEPR

Fiorella De Fiore, Bank for International
Settlements and CEPR

Mathias Dewatripont, ECARES,
Universite Libre de Bruxelles and CEPR

Jonathan Dinkel, University of Chicago
Booth School and CEPR

Barry Eichengreen, University of
California, Berkeley and CEPR

Simon J Evenett, University of St Gallen
and CEPR

Maryam Farboodi, MIT and CEPR

Antonio Fatás, INSEAD Singapore and
CEPR

Francesco Giavazzi, Bocconi University
and CEPR

Christian Gollier, Toulouse School of
Economics and CEPR

Rachel Griffith, IFS, University of
Manchester and CEPR

Timothy J. Hatton, University of Essex
and CEPR

Ethan Ilzetzki, London School of
Economics and CEPR

Beata Javorcik, EBRD and CEPR

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

Erik Lindqvist, Swedish Institute for
Social Research (SOFI)

Tom Kompas, University of Melbourne
and CEBRA

Miklós Koren, Central European
University and CEPR

Anton Korinek, University of Virginia
and CEPR

Philippe Martin, Sciences Po and CEPR

Warwick McKibbin, ANU College of
Asia and the Pacific

Kevin Hjortshøj O'Rourke, NYU Abu
Dhabi and CEPR

Evi Pappa, European University Institute
and CEPR

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

Richard Portes, London Business School
and CEPR

Carol Propper, Imperial College London
and CEPR

Lucrezia Reichlin, London Business
School and CEPR

Ricardo Reis, London School of
Economics and CEPR

Hélène Rey, London Business School and
CEPR

Dominic Rohner, University of Lausanne
and CEPR

Paola Sapienza, Northwestern University
and CEPR

Moritz Schularick, University of Bonn
and CEPR

Paul Seabright, Toulouse School of
Economics and CEPR

Flavio Toxvaerd, University of Cambridge

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

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

Jan C. van Ours, Erasmus University
Rotterdam and CEPR

Thierry Verdier, Paris School of
Economics and CEPR

Ethics

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

Submission to professional journals

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

<i>American Economic Review</i>	<i>Journal of Economic Growth</i>
<i>American Economic Review, Applied Economics</i>	<i>Journal of Economic Theory</i>
<i>American Economic Review, Insights</i>	<i>Journal of the European Economic Association*</i>
<i>American Economic Review, Economic Policy</i>	<i>Journal of Finance</i>
<i>American Economic Review, Macroeconomics</i>	<i>Journal of Financial Economics</i>
<i>American Economic Review, Microeconomics</i>	<i>Journal of International Economics</i>
<i>American Journal of Health Economics</i>	<i>Journal of Labor Economics*</i>
<i>Economic Journal</i>	<i>Journal of Monetary Economics</i>
<i>Economics of Disasters and Climate Change</i>	<i>Journal of Public Economics</i>
<i>International Economic Review</i>	<i>Journal of Political Economy</i>
<i>Journal of Development Economics</i>	<i>Journal of Population Economics</i>
<i>Journal of Econometrics*</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 23, 28 May 2020

Contents

How to exit Covid-19 lockdowns: Culture matters <i>Jean-Philippe Platteau and Vincenzo Verardi</i>	1
Baby steps: The gender division of childcare during the Covid-19 pandemic <i>Almudena Sevilla and Sarah Smith</i>	58
Production networks and epidemic spreading: How to restart the UK economy? <i>Anton Pichler, Marco Pangallo, R. Maria del Rio-Chanona, François Lafond and J. Doyne Farmer</i>	79
Optimal case detection and social distancing policies to suppress Covid-19 <i>Stefan Pollinger</i>	152
Exposure to the Covid-19 stock market crash and its effect on household expectations <i>Tobin Hanspal, Annika Weber and Johannes Wohlfart</i>	188
Lockdown strategies, mobility patterns and Covid-19 <i>Nikos Askitas, Konstantinos Tatsiramos and Bertrand Verheyden</i>	263
Impact of the state of emergency declaration for Covid-19 on preventive behaviours and mental conditions in Japan: Difference in difference analysis using panel data <i>Eiji Yamamura and Yoshiro Tsutsui</i>	303
Fast and local: How lockdown policies affect the spread and severity of Covid-19 <i>Jean-Philippe Bonardi, Quentin Gallea, Dimtriya Kalanoski and Rafael Lalive</i>	325

How to exit Covid-19 lockdowns: Culture matters¹

Jean-Philippe Platteau² and Vincenzo Verardi³

Date submitted: 24 May 2020; Date accepted: 26 May 2020

A key question is how countries can gradually exit the covid-19 lockdown in order to re-open their economies and mitigate the huge economic costs that the lockdown is imposing. Answering this question is the first step of the analysis proposed in this paper. Using a benchmark country known to be severely hit by the virus (Belgium), it compares the epidemiological effects of different stereotyped exit strategies. It concludes that, in order to avoid a rebound in infections and follow a relatively quick path toward ending the epidemic, the re-opening of the economy and the society must be very cautious and strict measures of social distancing and an ambitious and effective testing programme must be implemented. The second step, and main point of the paper, consists of exploring the role of a country's culture, more particularly the prevailing contact habits and norms. This is done by substituting the pattern of inter-individual interactions of two other countries for the pattern observed in the benchmark country. The results are striking: differences in the way people interact, and more specifically the frequencies of their contacts within and between age groups, seem to (partly) explain variations in the incidence of the virus and performances in battling against it. More precisely, if Belgium inherited the interaction pattern of Germany

- 1 During a seminar held at Namur and during many interpersonal discussions, we benefitted from useful comments from both economists and biologists or epidemiologists. Among the former, special thanks are due to François Bourguignon, Mathias Dewatripont, Catherine Guirkingier, Mathias Hungerbühler, Christian Kiedasch, André Sapir, Thierry Verdier, and Rainer von Sachs. Among the latter, Pierre Courtoy, Koen Deforche, Nicolas Franco, Nicole Moguevsky, and Eric Muraille are owed our gratitude. Of course, the authors are solely responsible for eventual errors or omissions.
- 2 Centre for Research in Economic Development (CRED), University of Namur.
- 3 Centre for Research in Economic Development (CRED), University of Namur. Vincenzo Verardi is associate researcher of the FNRS and gratefully acknowledges their financial support.

Copyright: Jean-Philippe Platteau and Vincenzo Verardi

when exiting the lockdown, it could achieve the objective of (partial) re-opening of the economy with more moderate policies than the ones it actually needs. And, conversely, if it inherited the social structure of Italy, it would have to take even more stringent measures lest the cost to bear as a result of economic re-opening should be (much) heavier. In addition to differences in the effectiveness of public health policies and in the genetic make-up of population groups, cultural specificities thus appear to play a significant role in explaining international and inter-regional variations in the incidence of the virus and the impact of public interventions.

How To Exit Covid-19 Lockdowns: Culture Matters

Jean-Philippe Platteau^{*†} and Vincenzo Verardi[‡]

May 27, 2020

Abstract

A key question is how countries can gradually exit the covid-19 lockdown in order to re-open their economies and mitigate the huge economic costs that the lockdown is imposing. Answering this question is the first step of the analysis proposed in this paper. Using a benchmark country known to be severely hit by the virus (Belgium), it compares the epidemiological effects of different stereotyped exit strategies. It concludes that, in order to avoid a rebound in infections and follow a relatively quick path toward ending the epidemic, the re-opening of the economy and the society must be very cautious and strict measures of social distancing and an ambitious and effective testing programme must be implemented. The second step, and main point of the paper, consists of exploring the role of a country's culture, more particularly the prevailing contact habits and norms. This is done by substituting the pattern

^{*}Corresponding author. e-mail:jean-philippe.platteau@unamur.be. Centre for Research in Economic Development (CRED), University of Namur. Rempart de la Vierge, 8. B-5000, Namur (Belgium).

[†]During a seminar held at Namur and during many interpersonal discussions, we benefitted from useful comments from both economists and biologists or epidemiologists. Among the former, special thanks are due to François Bourguignon, Mathias Dewatripont, Catherine Guirking, Mathias Hungerbühler, Christian Kiedasch, André Sapir, Thierry Verdier, and Rainer von Sachs. Among the latter, Pierre Courtoy, Koen Deforche, Nicolas Franco, Nicole Moguilevsky, and Eric Muraille are owed our gratitude. Of course, the authors are solely responsible for eventual errors or omissions.

[‡]Centre for Research in Economic Development (CRED), University of Namur. Rempart de la Vierge, 8. B-5000, Namur (Belgium). Vincenzo Verardi is associate researcher of the FNRS and gratefully acknowledges their financial support.

of inter-individual interactions of two other countries for the pattern observed in the benchmark country. The results are striking: differences in the way people interact, and more specifically the frequencies of their contacts within and between age groups, seem to (partly) explain variations in the incidence of the virus and performances in battling against it. More precisely, if Belgium inherited the interaction pattern of Germany when exiting the lockdown, it could achieve the objective of (partial) re-opening of the economy with more moderate policies than the ones it actually needs. And, conversely, if it inherited the social structure of Italy, it would have to take even more stringent measures lest the cost to bear as a result of economic re-opening should be (much) heavier. In addition to differences in the effectiveness of public health policies and in the genetic make-up of population groups, cultural specificities thus appear to play a significant role in explaining international and inter-regional variations in the incidence of the virus and the impact of public interventions.

1 Introduction

The coronavirus pandemic is probably the most severe crisis that has affected the world since the second world war, and the damage that it creates to health, wealth, and well-being are enormous and much more severe than those caused by the 2007-2009 financial crisis. It is estimated that the loss of world output caused by covid-19 will be of the order of 15% by the end of the year. Without entering into the details of the “grim calculus” involved in addressing “the stark choices between life, death and the economy” (Economist, April 4-10), it is apparent that the implicit value accorded to human life under the present crisis is very high, so high as to easily exceed any benchmark value used so far.¹ It is because of this almost absolute priority given to rescuing human lives that so many countries have adopted more or less severe forms of lockdown and decided that no relaxation of the associated discipline can be considered unless the peak of the epidemic has been overshoot and the

¹Thus, the threshold used by the British National Health Agency (United Kingdom) to decide whether or not to reimburse a treatment is of the order of £25.000 per QALY (or Quality-Adjusted Life Year). The life of an infant with a ‘full-health-equivalent’ life expectancy of 80 years would then be worth 2 million £, or 2.5 million Euros at the exchange rate of £1 = 1.25 Euro. In many countries, the amount of money spent per life saved under the coronavirus neatly exceeds that threshold value.

decline in infections has been sustained. In other words, the point of the lockdown was to accept short-term economic pain at the price of getting the virus under control. Now that the perspective of the post-peak phase has got closer, an increasing number of countries are studying, and in a few cases even experimenting with, different ways of de-escalation. The immense economic costs of the lockdown and their reflection in the growing pressures exerted by the most affected economic groups force the politicians to confront the trade-off between human lives and economic growth. Hence the rapidly growing interest of the economics profession in the present pandemic.

The objective is to restart the economy in such a manner that the public health cost is minimized or at least mitigated. A critical question that springs to mind is whether a particular exit strategy is likely to cause a rebound or, instead, a simple deceleration in the descending portion of the infection and the mortality curves. A comparatively effective exit strategy is one which, for a given economic benefit, would succeed in terminating the propagation of the virus, and mitigating the increase in death-toll, in as short a time as possible. But we are not only, and not even mainly, interested in searching for the best suited strategies to exit a lockdown. Building on this knowledge, our central contribution is to explore the impact of a specific cultural trait, the frequency with which people visit each other, on the epidemiological effectiveness of different lockdown strategies. Contact habits are especially important in the Covid-19 pandemic because the SARS-CoV-2 virus is highly contagious and transmitted through the emission of expiratory aerosol particles (as is typical in respiratory disease transmission).

To predict the covid-19's propagation and its lethal consequences, there are various epidemiological models available in the literature, such as the models of Bernoulli, Reed-Frost and the SIR or SEIR models. One key issue, however, is knowledge about the values to be ascribed to the parameters of the model chosen. In the present case, we are particularly helpless owing to paucity of country-specific data, which itself results from lack of testing for the presence of the virus in the population. Many predictions are therefore broad approximations that need to be continuously updated as more reliable data become available. In this paper, the data problem is not too serious because we do not actually intend to make predictions for a particular country but are instead interested in studying the epidemiolog-

ical effects of different lockdown exit strategies, as well as the way in which these effects are influenced by prevailing social norms of conduct. In other words, we want to compare the covid-19 pandemic trajectories when a country starts from a lockdown situation and contemplates various ways of exiting it. Therefore, we are not too much concerned by the imprecise nature of the figures yielded by the model used (say, infection or mortality rates): what matters is how a number of possible exit scenarios produce different effects on the course of the epidemic. For those interested in predictions for Belgium, Deforche (2020) is a better reference.

We use a model of the so-called SEIR type in which different age classes are embodied. We are thus able to specify age-specific values for a number of parameters that are obviously not uniform across age classes. If we believe that the model provides a coherent and relevant structure to depict the propagation of covid-19, which obviously requires that the parameters take plausible values, we can be confident that the differences observed between the effects of the different scenarios examined are not mathematical artefacts. More precisely, changing (marginally) the values of this or that parameter will affect the numbers associated with the simulated variables (number of people susceptible to the virus, or exposed to it, number of infectious cases) for each scenario, but will not significantly modify the inter-scenario differences for these variables. As a consequence, we can dispense with the tedious task of calculating prediction intervals for a range of plausible values of several key parameters, which has the unavoidable effect of overburdening figures and blurring the lessons learned. In short, since we use a comparative approach, we rely on simulation of paths obtained with given a set of assumed parameters. We nevertheless do a short sensitivity analysis to be reassured that the results are not driven by a unique combination of parameters.

There is one exception to the above rule of fixing the model's parameters, and it is directly related to the second objective assigned to this paper. As hinted at above, we want to have an idea about the way a country's social interaction pattern impinges on the effects produced by different exit strategies. By social interaction pattern, we mean the frequencies of contacts between people belonging to the same age group and between people belonging to different age groups. These frequencies reflect social norms and habits guiding people's behaviour in a given society, with specific reference to inter-individual physical interactions

and visits at the workplace, at school, at home, and at other places. At home and in other activities such as leisure and shopping, frequencies of contacts are partly determined by family systems which are known to differ significantly between countries and even between regions inside particular countries (Todd, 2011; Guirkinger and Platteau, 2020). Contact frequencies are thus taken as a culturally determined. Note that, because of the particular objective pursued in this paper, it has not been possible to follow the common practice among economists addressing other facets of the epidemic. This practice consists of grafting behavioural equations into the SEIR model with a view to endogenizing contact frequencies (implied in the choice of social activities) and the infection rates, which vary between age groups.

Since comparisons between countries are plagued by the presence of numerous confounding factors, we have chosen the following approach: we analyze a benchmark country case on which the model is built, namely Belgium, and we then proceed by substituting the interaction matrices of two countries presumed to be culturally different, Italy and Germany. Three countries will thus be explored, two of which are pseudo-countries: real Belgium (the benchmark), Belgium *as though* it was endowed with Italian contact habits (from the day of exiting the lockdown while keeping every thing else unchanged) and Belgium *as though* it was endowed (*ceteris paribus*) with German contact habits. To the latter two (pseudo-) countries, we apply the same analysis aimed at bringing out the effects of different exit strategies. In this way, we are able to isolate the effects of culture since all other things are assumed equal. Naturally the results of this paper cannot be understood as guidelines for lockdown exit strategies in Germany and Italy. Many other factors (such as population density, international connectivity, genetic characteristics, health infrastructure, etc.) are such that a policy seen as ineffective for pseudo-Germany or pseudo-Italy could actually be quite effective in real Italy or real Germany (and vice-versa).

The choice of Belgium as our benchmark is not coincidental. On the one hand, Belgium has been hard hit by the coronavirus crisis. This is partly due to its central geographical position at the heart of Western Europe and its internally strong connectivity, and partly to the return of many Belgian citizens, including people of Italian origin, from holidays spent during the season of carnivals and skiing in Italy. On the other hand, from a cultural

and social standpoint, Belgium is situated mid-way between the strongly knit societies of southern Europe in which the family plays a central role, and the relatively loose social structures of northern and eastern Europe.

The government is supposed to have the following policy instruments available. First, it can choose the extent to which the economy is being re-opened. Second, it can specify social-distancing and protection measures of varying intensity. Third, it can test the population with a view to detecting infected people and isolating them if found positive. Since the supply of tests is constrained by the existence of various bottlenecks (availability of reagents and/or skilled personnel to administer the tests, limited logistical capacity for sample analysis, in particular) that are (partly) actionable by the government, the scope of testing is itself a decision variable. We therefore use a standard epidemiological model to make it speak to policy-makers and social scientists. Setting up this bridge requires that we interpret certain parameters of the model as so many policy instruments.

Two key findings need to be brought into limelight. First, in order to exit lockdown, there is no escape from a gradual opening of the economy. Complete re-opening would have disastrous effects on public health. If the government could strictly enforce stringent social-distancing measures, the scope of testing would not wield a large influence on the course of the epidemic. However, if only moderate social distancing can be applied, testing appears to be an essential substitute for insufficient discipline lest the relaxation of the lockdown should cause a severe rebound of the epidemic and certainly lead to a high death-toll. Second, the required measures would be much more moderate if Belgium was “lucky enough” to inherit a German-like pattern of social interactions when exiting lockdown, but much more severe if it was endowed with an Italian-like pattern. The latter result calls into question the current view according to which the success in fighting covid-19 of Germany and other countries such as Austria and Norway, is essentially the result of more effective policies based on testing, tracing and quarantining. Other factors that have nothing to do with policy-making seem to be at play, and they do not only involve genetic variations (a biological given), but also variations in norms of conduct (a social given).

The remainder of the paper is structured as follows. In Section 2, we present in three successive steps the epidemiological model chosen to simulate the effects of eleven strategies

of lockdown exit. Then, in Section 3, we explain how we are going to use it. First, we make explicit the underlying assumptions by assigning values to the different parameters of the model and, second, we expound our simulation approach in detail. Section 4 discusses the results obtained by comparing the different exit strategies as applied to the benchmark country (Belgium), while Section 5 examines how the effects of these strategies are affected when the social interaction patterns of Italy and Germany are successively substituted for the Belgian one. Section 6 is more than a conclusion: not only does it summarize our main results but it also puts them in perspective and draws important policy implications.

2 The epidemiological model

2.1 The SIR model

The SIR model divides the population into three groups (compartments) of individuals: S , I and R (with $S + I + R = N$, where N is the population size). Group S is the group of susceptible individuals (i.e. those individuals that are at risk of being contaminated). For the case of Covid-19, at the beginning of the epidemic S is the entire population given that nobody has anti-bodies (it is indeed a new virus for which no vaccine is available). Group I is the group of individuals that have been contaminated recently and that are infectious. Finally, R is the group of individuals that were contaminated but that had an outcome (either a recovery or death). They are not infectious anymore.

The sizes of these groups evolve over time as the virus spreads. The size of S decreases when people get contaminated and move into the infectious group I . When individuals recover or die, they then move from the infectious group I to the removed group R . The evolution of the sizes of these groups can be modeled by a system of 3 differential equations:

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

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

$$dR/dt = \gamma I$$

The first equation states that the size of S decreases by the number of newly contaminated individuals, which is simply the transmission rate (β) multiplied by the number of susceptible individuals (S) who came into contact with infectious individuals (I). More precisely, each susceptible person contacts β people per day, a fraction I/N of which are infectious.

The second equation states that the number of infectious individuals (I) will be increased by the newly contaminated individuals minus the previously infectious individuals that had an outcome and moved to group R (i.e. the removal rate γ multiplied by the infectious individuals I).

Finally, the last equation states that the removed group increases by the number of individuals that were infectious that had an outcome (γI). In the case of Covid19, before the beginning of the epidemic the size of S is the entire population (as nobody is immune to the new virus). Then, once a first individual is contaminated, S decreases by one unit and I increases by one unit. This is the beginning of the dynamic of the epidemic. After some time, this infectious individual contaminates new individuals before recovering (or dying). In the meantime, the newly contaminated individual start spreading the virus and the epidemic starts.

2.2 The SEIR model

Though very simple, SIR gives a good idea of how an epidemic is likely to evolve. However, there is sometimes a significant incubation period for infections that is neglected and that should be accounted for. Indeed, during the latency period, individuals have been infected but are not yet infectious. The SIR model is then generally augmented of group E (for the exposed) to take this into account. The system of dynamic equations hence becomes:

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

$$dE/dt = \beta SI/N - \sigma E$$

$$dI/dt = \sigma E - \gamma I$$

$$dR/dt = \gamma I$$

which is very similar in dynamics to SIR except that the incubation rate σ (i.e. the rate of latent individuals becoming infectious) is added.


The lifelong immunity assumption in this model could be relaxed since for many diseases the immunity after infection wanes over time. The model only changes marginally since a given percentage of individuals in group R will return in group S (results could, however, be dramatically different). Similarly, vital dynamics considering both natality and mortality rates could be introduced, yet these improvements are generally more worthy to analyze cyclical endemic infections than a fast evolving epidemic like the one we are concerned with here. Note that since we assume a short time-frame of the epidemic, we do not consider migratory movements and we assume that nobody changes age-class.

2.3 The age-structured SEIR model

A very specific peculiarity of the Covid-19 epidemic is that it does not seem to affect individuals of different age categories in the same way. In the bar diagram here below we present the number of detected cases in Belgium (until april 6th 2020) by age category.


It is striking that only a very limited number of detected cases are younger than 20. This could mean that the young are less susceptible to be contaminated and/or that they only experience light symptoms when affected and hence go unnoticed. In terms of deaths, the picture is even more striking showing that mortality and the age structure of the population (for the same period as the previous graph) are highly correlated:

Figure 1: Positive cases by age groups



Data source: Sciensano, <https://epistat.wiv-isp.be/covid/>, early april

Figure 2: Deaths cases by age groups



Data source: Sciensano, <https://epistat.wiv-isp.be/covid/>, early april

We therefore believe that age should be incorporated in SEIR to better model Covid-19. The possible asymptomatic cases should be considered as well since they are likely to be

present, especially among the young ones. Toward that purpose, we slightly modify the model proposed by Prem et al. (2020) by allowing for age-specific contamination rates while keeping the assumption that the probability of exhibiting symptoms is the same for each age category². We decided to consider β as age-specific because unequal representation of tested positives among age classes (with respect to their share in the population) is an empirical fact. Finally, our model is specified in such a way as to incorporate the possibility of screening based on testing. The extent of testing is represented by parameter τ , which stands for the percentage of infected individuals identified in each period, quarantined, and then moved to the R category. The system of differential equations becomes (for each age group i):

$$dS_i/dt = -\beta_i S_i \sum_j C_{ij} I_j^S / N_j - \alpha \beta_i S_i \sum_j C_{ij} I_j^A / N_j$$

$$dE_i/dt = \beta_i S_i \sum_j I_j^S / N_j + \alpha \beta_i S_i \sum_j I_j^A / N_j - \sigma E_i$$

$$dI_i^S/dt = \rho \sigma E_i - (\gamma_S + \tau) I_i^S$$

$$dI_i^A/dt = (1 - \rho) \sigma E_i - (\gamma_A + \tau) I_i^A$$

$$dR_i/dt = (\gamma_S + \tau) I_i^S + (\gamma_A + \tau) I_i^A$$

As previously β_i is the infection rate and σ is the incubation rate. In this model, to better mimic Belgian data, we assume that the infection rate is age-specific³. A more thorough explanation will come in the parameter section.

Parameter σ is $\frac{1}{d_L}$, where d_L is the latency period in days before becoming infectious. In this version of the model, I^S is the number of symptomatic cases, I^A is the number of asymptomatic cases, and N_j is the number of individuals belonging to age class j . Parameter α is the relative probability of being infected by an asymptomatic individual compared to the probability of being infected by a symptomatic individual (also called the discount on transmission). Parameter ρ is the probability for an infected case to have symptoms (even

²This point will be discussed in more detail later

³As in Towers and Feng (2012), the contact matrices concern all contacts, not only those that might transmit infection. These matrices are hence scaled to reflect the “probability of transmission on contact”.

if often mild).⁴ Finally, parameter γ_S is the removal rate for symptomatic cases (which is $\frac{1}{d_{IS}}$ where d_{IS} is the infectious period for symptomatic cases) and γ_A is the removal rate for asymptomatic cases (which is $\frac{1}{d_{IA}}$ where d_{IA} is the infectious period for asymptomatic cases). The infectiousness period for asymptomatic individuals is much larger than for the symptomatic ones. This is because while symptomatic individuals are quarantined after a few days (and hence removed from the infectious group), asymptomatic individuals remain much longer in the infectious group, implying that their removal rate is lower.⁵

Turning now to the key social structure variable embodied in the model, it is evident that members belonging to each age group can be infected by members of all age groups, and the exposition to infection depends on the number of contacts between (and within) age groups. The frequency of these contacts are included in contact matrices C (hence C_{ij} is the average number of contacts an individual belonging to age group i has with age group j). As far as contact matrices C are concerned, they are location-specific in the sense that they take into account the number of contacts between and within age groups at home, school, work or in a generic class called other locations (which includes transportation en leisure). The final contact matrix is the sum of all these location-specific matrices. These have been estimated by Mossong et al. (2008) and by Prem et al. (2017) using contact surveys and demographic data. The interest of both researcher teams in contact matrices was born of their belief that heterogeneity in contact networks has a major effect in determining whether a pathogen can become epidemic or persist at endemic levels. Contact matrices for a wide range of countries (including Belgium) are available from Prem et al. (2017). A big advantage of having location-specific contact matrices is that it makes it easy to see what would happen if specific restrictions, such as closing (or re-opening) schools or confining the elderly, were implemented.

It must be stressed that the contact frequencies reported in the matrices are constant, reflecting the fact that transitions from one state to another in epidemiological models are

⁴In our model any symptom, even mild, is considered. Naturally, only a very limited number of symptomatic individuals (especially among the young) will observe severe symptoms and will need to be hospitalized. The model could be modified to include a hospitalization compartment but, given our objective, we believe that this would complexify the model without adding any valuable information.

⁵We could have created a specific compartment for quarantined individuals but we decided to move them in the R compartment directly. The outcome of the model would remain unaffected.

typically determined by aggregates without microeconomic foundations.⁶In economic models, by contrast, transitions are influenced by optimizing behaviour in which individuals weigh the benefits and costs of different types of actions. When individual decision-making is thus taken into account, an equilibrium is reached in which the amount of social distancing is smaller than the amount that a government acting as a social planner would choose (Garibaldi et al., 2020 ; Alfaro et al., 2020 ; Eichenbaum et al., 2020). This is essentially because, when making their decisions, individuals take into account the risk of infection that social activities imply for themselves, yet not the externalities that their behaviour is liable to create. More specifically, they think of the risk of being infected by other people but not of the risk of contaminating them. Moreover, they ignore the congestion externalities that expose the available medical facilities to the risk of acute stress (Ichino et al., 2020).

The (constant) contact frequencies in a basic SEIR model are those prevailing in ordinary circumstances. In the present paper, we will scale them down to a larger or smaller extent depending upon whether the country is under or after lockdown. In considering a partial re-opening of the economy after lockdown, the figures in all cells of the interaction matrix will be thus multiplied by a constant factor smaller than one but closer to unity when compared to the lockdown situation. Clearly, the new contact frequencies are products which can be interpreted in different ways: as physical contacts constrained by a government reluctant to open the whole economy, or as contacts freely chosen by the individuals under the fear of getting infected (as attested by the opposition of several trade union organizations against the re-opening of schools, public transportation, and even business firms where distance work is not feasible), or else as a mixture of the two possibilities.

Whether one interpretation is more valid than the other is an empirical question that we are unable to solve at this stage given the lack of appropriate data. In theory, if the government is able and willing, through strong detection and sanctions, to strictly enforce its policy of (only) partial re-opening, the increased scope of interactions will be mostly determined by the public policy. If, on the other hand, enforcement of the chosen policy is weak, the degree of influence of the government will depend on the proportion of law-abiding people in the


⁶More elaborate models take into account the evolution of interaction patterns, yet to the best of our knowledge, the contacts are not endogenized in the sense of being the outcomes of optimal individual behaviour.

population, or on the civic sense of the citizens. If the civic sense is low, the increase in interactions following a partial re-opening of the economy will be mostly attributable to the decentralized decisions of self-centered individuals.

The same complexity born of the interaction between individual behaviour faced with a serious risk of infection and government protective policies is still more in evidence when public health measures are considered, social-distancing measures in particular. Theory yields ambiguous predictions. On the one hand, strict distancing measures, including the obligation to wear a protective mask, may induce people to increase the frequency of their visits because wearing the mask makes them feel safer. But, on the other hand, the same measures may have the opposite effect if the cost of wearing the mask is high, or if people perceive them as signalling a high risk involved in human encounters. Plausibly, individuals differ in their reactions, and these reactions probably vary between age groups, too. It is therefore hard to know for sure how different distancing policies affect contact habits. We therefore prefer to assume that contact frequencies are not influenced by the degree of social distancing imposed (or recommended) by the government.

Finally, since our central objective is to compare particular exit scenarios not between themselves but as they unfold in real Belgium, pseudo-Germany (that is, Belgium as though it was endowed with the contact matrix of Germany at lockdown exit), and pseudo-Italy (Belgium as though endowed with the contact matrix of Italy), our approach consisting of scaling up contact frequencies after lockdown is rather inconsequential. Indeed, what matters for our purpose are the relative levels of contact frequencies between the three countries and, since we actually deal with Belgium only, whether the absolute number of visits increases or not, the values of the ratios stay unchanged.

Figure 3: Location specific interaction matrices (Belgium)



Covid Economics 23, 28 May 2020: 1-57

3 Methodology

3.1 The parameters

Let us now display the numerical values of the different generic (i.e. non age specific) parameters of the model that we will use for simulations. We must bear in mind that some of them have been set in such a way that our model's predictions square well with the actual empirical figures obtained for Belgium at the baseline;

Generic parameters		
Relative prob. to be infected by an asymptomatic indl	α	= 1/4
Latency period	d_L	= 2
Infectious period for symptomatic indls	d_{IS}	= 3
Infectious period for asymptomatic indls	d_{IA}	= 10
Probability of exhibiting symptoms	ρ	= 0.8

The probability of being contaminated by an asymptomatic individual, α , is assumed to be one-fourth of the probability of being contaminated by a symptomatic individual, which is according to the assumption made by Prem et al (2020). The latency period (the period that elapses before a contaminated individual becomes infectious) is assumed to be 2 days, which is shorter than the period needed for the first symptoms to show up, estimated to be about 4 days. Moreover, for those individuals who will become symptomatic, we set the infectious period to 3 days. Our assumption is that an exposed individual will not be infectious for 2 days. Then, he will become infectious but will not show strong symptoms for 3 more days. After this lapse of 5 days after contamination, the individual will become symptomatic and quarantined (generally at home but sometimes in a hospital). Note that when we set the removal rate for the asymptomatic, we assume that the infectious period for them lasts 10 days since they are not identified and hence not quarantined and removed. There are two effects running here: on the one hand, symptomatic patients are believed to be more contaminant because they cough and sneeze, yet at the same time they are easily identified and quarantined; on the other hand, asymptomatic people are less contaminant, but they are hard to identify (in the absence of testing), and hence they tend to stay longer in contact with susceptible individuals. The probability of exhibiting symptoms is set to 0.8 as estimated in Mizumoto et al. (2020). In several intermediate models, we allowed the ρ coefficient to vary between age classes and, in particular, we assumed it to be 0.4 for the young and 0.8 for the others, except the elderly for whom we set it to 0.9. However, scenarios in which an age-specific β is very different for the most extreme age classes and ρ is age specific with the values defined above yields results very similar to those obtained under the scenario in which ρ is fixed. We therefore decided to give up the idea of age-specific ρ so as to reduce the number of parameters.

Belgian data show that infectiousness varies significantly by age group, and it is therefore difficult to keep the conventional assumption of a uniform infectiousness rate. In particular, the elderly appear to be comparatively infectious, which is probably related to the importance of retirement homes as a place of residence for a relatively large proportion of this age class in the country. As evidence about the propagation of the epidemic has shown, the virus has spread massively in old-age homes. By contrast, the young seem to be less infectious than older age groups, a fact well documented for several countries. While this point is still strongly discussed in the literature, Bunyavanich et al. (2020) have recently published a paper showing the age-dependent expression of ACE2 in nasal epithelium, usually the first point of contact for SARS-CoV-2 and the human body. According to the authors, lower ACE2 expression in children relative to adults may help explain why Covid-19 is less prevalent among the young.

To allow for the most important differences, we consider 16 classes of 5-year intervals ranging from 0 to 75, plus a class of 75 or older.

The age-specific infectiousness parameter is $\beta_i = \beta z$, where z is a vector containing the over- and under-representation factors among tested positives for each age class. Parameter β is gauged on the data. As is commonly done in SEIR models with age-based disaggregation, its value is set by relying on the initial reproduction number R_0 (see Towers and Feng, 2012). More precisely, $\beta = R_0 K$ with $K = \frac{\gamma}{\rho + \alpha(1-\rho)}/\lambda$, where γ is the average removal (or recovery) rate.⁷ Parameter ρ is the probability for a contaminated individual to be symptomatic.⁸ As for α , we know that it represents the discount on transmission for the asymptomatic group. Finally, λ is the largest eigenvalue of the matrix $M_{ij} = C_{ij}N_i/N_j$, where N_i and N_j are the number of individuals in age classes i and j .⁹ The initial date for which we estimate the value of R_0 is March 13th, which corresponds to the first date at which the number of hospitalizations started to be systematically counted. The estimation approach is the time-dependent method proposed by Wallinga and Teunis (2004), and in our case it yields a R_0 value for the first day that is equal to 3.5 approximately.¹⁰ (This

⁷For simplicity, we assume a unique $\gamma = 0.25$, which is a weighted average of γ_c and γ_{sc}

⁸It is set to $\rho = 0.8$ as estimated by Mizumoto et al. (2020)

⁹Dietz (1993) states that R_0 is “the number of secondary cases one case would produce in a completely susceptible population”

¹⁰To estimate the R_0 parameter, the distribution of serial intervals (the times between symptomatic cases

value implies that about five to ten days earlier, the R_0 associated to infections was of the same order). The removal rate is equal to one-third, that is, a value of one divided by the assumed length of the infectiousness period (3 days).

Vector $z = (0.2, 0.8, 1, 1, 4.1)$ for the age categories (0 – 20, 20 – 40, 40 – 60, 60 – 75, 75+), with the z values representing approximately the over- or under-representation of various age groups in the population of people tested positive for the presence of the virus (as of April 17). Using the Belgian data, these values have been calculated as the share of different age groups in the positively tested population relative to their share in the total population. Thus, for example, younger individuals appear to be less easily contaminated (20%) than the median class 40-60 when being in contact with infectious individuals. At the other extreme, the 75+ class is 4.1 times more present among the contaminated than the median class. Since late April, Belgium has continuously increased the number of tests performed in retirement homes. Our estimate dates back to before this increase in testing. Nevertheless, we are in a position to compare the over-representation of the elderly long before this effect could be suspected. What we find is that on April 4, the over-representation of the elderly was broadly speaking of the same order as what we discussed above.

It could be objected that under-representation among the young may reflect selective testing: only those individuals exhibiting the most severe symptoms are being tested while young people tend to present only light symptoms when they are contaminated. However, Boast et al. (2020), relying on studies from South Korea and Iceland (which have undertaken widespread community testing), found significantly less positive cases among children. Similarly, in the Italian town of Vo' where 70% of its population has been tested, no children younger than 10 have been found positive, despite a 2.6% positive rate in the general population. In the light of this evidence, Belgian data do not look anomalous.

We still want to deal with the potential bias noted above when we will conduct a sensitivity analysis of our results. More precisely, we will allow the proportion of positive cases among the young to be as high as 4 times the value observed in the data, fixing the observed value of 0.2 as the minimum. In other words, we will generate the representation of the young

in a chain of transmission) is assumed to be a gamma distribution with location and shape parameters respectively of 3 and 1.75. When relying on Wallinga and Teunis (2004), we use the R library R0 while when relying on Cori et al. (2013), we use the R library EpiEstim.

from a uniform distribution (0.2,0.8). To guarantee that the average β does not change when the under-representation among the young is reduced, the over-representation of the elderly will be reduced accordingly. The intermediate age categories, which are apparently neither under-represented nor over-represented, will not be modified.

Finally, the large over-representation of the elderly could stem from the way in which Belgian care facilities, old-age homes in particular, were organized before the epidemic. It is highly possible (and even desirable) that the transmission rate of the elderly will be much lower after the lockdown. This would imply a lower β_i for the elderly category, and hence a smaller aggregate β too. In this eventuality, the shape of the curves would certainly be affected as the post-lockdown epidemic would be less severe than in the preceding period. In any case, this should not affect the main messages of the paper.

Given the high number of unknowns regarding covid-19, many parameter values, even when they are extracted from the fast-growing specialized literature or based on empirical evidence related to Belgium, could be considered as somehow arbitrary. Hence the importance of confronting the simulations done with the model to real data so as to be able to verify if our simulated paths are in line with real life observations.

First note that the effective reproduction number, R_t , is the number of people in a population who can be infected by an individual at any specific time. For the purpose of estimating the observable R_t , we use the total number of hospitalizations. We believe this is a better indicator of the evolution of the epidemic than the number of people tested positives, since it does not depend on the number of tests performed to detect infected individuals. As for the R_t inferred from our simulation model, it is based on the total number of infections, $I_i^S + I_i^A$ from March 1st till the beginning of the lockdown. We compute the average of the successive daily values of R_t during this period of 20 days. The daily values of R_t are obtained by using the aforementioned time-dependent model of Wallinga and Teunis (2004). The value obtained for the observable R_t is 1.5684. It appears to be rather close to the simulated value, which is equal to 1.6352. We can therefore hope that our model roughly follows the path of the true epidemic.

We still want to be reassured that the path simulated by our model coincides with the observed path both before and during the lockdown period. To achieve this, we use the

Table 1: Observed and estimated number of deceased

Number of deceased by age ¹³		
Age class	Observed	Estimated
[0-24]	1	13
[25-44]	19	103
[45-64]	309	354
[65-74]	653	1259
[75+]	3942	4224

Source: Sciensano, <https://epistat.wiv-isp.be/covid/covid-19.html>

method of Cori et al (2013), which estimates time-varying instantaneous reproduction rates in a backward-looking manner.¹¹ The correlation between the values of these rates as obtained for the observed hospitalizations and the values obtained for our simulated infections is as high as 0.96.¹²

We also need to verify how the model performs in terms of predicted values by age class. We therefore compare the number of deaths coming from the model with those actually observed (by applying infection fatality rates by age class as reported by Ferguson et al. (2020) to all simulated infections). On the 23rd of April 2020, the mortality rates by age class are those shown in table 1). Evidently, these numbers depend on the initial contamination rate assumed in the model, and should therefore be seen as very broad approximations

Our model thus yields slightly more pessimistic figures than the actually observed ones.¹⁴ Given the uncertainty (for Belgium) of the mortality rates by age as used by Ferguson et al. (2020) (to the best of our knowledge, these estimates are not available for Belgium), and given the likely imprecision in the counting of covid-related deaths, we believe that

¹¹We were not able to compute the R_t values for the period after the start of the lockdown by using the forward-looking approach proposed by Wallinga and Teunis (2004).

¹²Note that the hospitalizations have been lagged by 5 days to take into account the latency period and the delay before going to the hospital.

¹³On april 26th, 2020

¹⁴Deaths for which no information was available agewise (they amount to more than 2000 for that data) were not used in the analysis. This seems less arbitrary than distributing them proportionally to age classes. Admittedly, our procedure may cause our model to look a bit too pessimistic regarding the total number of deceased observed by class, although no such bias would exist if the non-accounted deaths are proportionally redistributed among the age classes.

these figures are sufficiently similar not to reject the set of underlying parameters. It is evident, however, that we cannot avoid relatively inaccurate computed values for the numbers infections. However, as stated previously, our objective here is not to predict but to compare scenarios. In short, although we would like to have more precise estimates for the parameters of our model, we think that our approach has sufficient legitimacy to yield meaningful results given our central objective. A sensitivity analysis is conducted to check whether our results significantly depend on initial assumptions.

Finally, relying on seriological tests performed on blood donors, it has been estimated that on April 14 around 4.5% of the population has been infected.¹⁵ With our model-based simulations at that date, we have that approximately 3.5% of the population has been infected. Though not negligible, the difference is not really big, and it must be borne in mind that the figure of 4.5% is not necessarily fully representative to the extent that it was estimated on the basis of a sample of relatively healthy individuals (those who decided to donate blood). However, we cannot rule out the possibility that a rebound simulated by our model is larger than the one that would be obtained if the true immunity rate was captured (a lower infection rate today means a higher rate tomorrow in case a rebound occurs).

3.2 The simulation approach

The model estimates daily values for S, E, I and R for all age categories. To simplify the reading of graphs we create four categories. The “young”, a category that ranges from 0 to 25 years, the “young adults” ranging from 25 to 45, the “middle-aged adults” ranging from 45 to 60, the “old adults” ranging from 60 to 75, and the “elderly” going from 75 and beyond.

The beginning of the epidemic in Belgium roughly coincided with the return of people from Italy after skiing (and attending carnival festivals) on February 29th, 2020. We assume that at this date 300 infectious individuals appeared in each age category. There were certainly several cases already active in the country at that point, but the significant shock was

¹⁵<https://www.sciensano.be/fr/coin-presse/43-de-la-population-belge-a-developpe-des-anticorps-contre-le-coronavirus>

caused by Belgian residents coming back from holidays. This assumption is important for the beginning of the epidemic, but should not really matter in the long-run.¹⁶

In the following, for all the scenarios considered, the starting point is the lockdown imposed by the government around mid-March. To recall, on March 12, it banned all public events involving large meetings. On March 14, the closure of schools was announced while social-distancing measures were recommended. Finally, on March 18, the lockdown was imposed. To represent this lockdown in the terms of the model, we make the following assumptions: while the interpersonal contacts at home are maintained at a level of 80% of what they were before the lockdown, they are maintained at 30% in the workplace, at 30% in other activities (including transportation, shopping and leisure activities), and at 5% in the schooling environment where only few children continue to be admitted when it is needed to relieve critical workers (in the health sector more particularly) of childcare duties.

As pointed out earlier, the government has three policy instruments at its disposal: the extent of re-opening of the economy, the degree of strictness of social-distancing measures, and the scope of testing.

Degree of re-opening of the economy and society

In terms of our model, modifying contacts at re-opening is represented by changes in the appropriate coefficients of the interaction matrix. During the lockdown, and except for interactions at home, interpersonal contacts are assumed to have been considerably restricted compared to what they were before (see above).. The small levels of interactions occurring at the workplace, at schools, and in other activities can be modified at will to consider different de-confinement strategies. It is also possible to intervene on specific rows and/or columns of the interaction matrix, if one is interested in policies targeted at specific age groups. This will not be done in this paper. Here, we will consider only two cases: com-

¹⁶This number was set by looking at the Influenza Monitoring of Sciensano (<https://epistat.wiv-isp.be/influenza/>). According to epidemiologist Marius Gilbert (heard in the news), the infection curve exhibits a second peak that is unlikely to be caused by the flu and is most probably related to Covid. The incidence per week (per 100,000 inhabitants) for week 10 (the week starting on March 2) was 313 cases of infection, implying that the incidence per day was approximately $313 \cdot 110/7 \simeq 4900$ (the multiplying factor 110 is obtained by dividing the total population of 11 Million by 100,000) . We propose a reasonably good approximation since we arrive at 300 individuals for the end of the previous week in each class, which leads to a total of 4800 infectious individuals.

plete and partial re-opening. While the meaning of complete re-opening is clear, partial re-opening implies that contacts at home are maintained at 80% of what they were before the lockdown, contacts at the workplace and in other activities are scaled up to 60% of the same, and those at schools are raised to 30%. As pointed out earlier, the new contact frequencies are thus products which can be interpreted in different ways: as physical contacts constrained by the government (compared to the pre-lockdown situation), as contacts freely chosen by the individuals under the fear of getting infected, or as a mixture of the two possibilities.

Imposing general social-distancing measures

General social-distancing measures are aimed at reducing the general transmissibility of the virus SARS-CoV-2. In terms of our model, the absence of any public health measures amounts to keeping the infectious rate β_i at the level that prevailed before the lockdown. Remember that $\beta_i = \beta z$, where $\beta = R_0 K$ and $R_0 = 3.5$. If, everything else being constant, moderate social distancing is imposed, β is assumed to decrease from $3.5K$ to $3K$. If strict distancing is imposed, there is a further decrease in β , from $3K$ to $2.5K$. Bearing in mind that this is equivalent to assuming that, if we were at the very start of the epidemic, the reproduction number would be reduced from 3.5 to 3.0, and then to 2.5. In other words, the effect of distancing is conceptualized as though it were causing a fall in the initial severity of the epidemic.

In choosing our incremental steps to describe the rising levels of distancing, we wanted to avoid the two extreme cases in which the impact of social distancing on the evolution of the epidemic would be either unrealistically small or unrealistically big. Thus, if we had chosen to decrease β to a value much smaller than $2.5K$ to describe strong distancing, complementary measures, testing in particular, would be almost redundant.

Carrying out testing

Changes in the scope of testing are modeled as variations in parameter τ . We consider three possible scenarios: i) no testing, ii) moderate testing that would allow to detect and quarantine 1% of the infectious individuals every day, and iii) intensive testing that would allow to detect and quarantine 5% of the these individuals again every day. The assumption

Table 2: *Parameter configurations corresponding to different exit scenarios*

N.	Distancing	Maintained contacts at re-opening				Tests
		Home	Work	Other	School	Screening
1	No	100%	100%	100%	100%	No
2	No	100%	100%	100%	100%	Strong
3	Strong	100%	100%	100%	100%	No
4	Strong	100%	100%	100%	100%	Strong
5	No	80%	60%	60%	30%	Strong
6	Moderate	80%	60%	60%	30%	No
7	Moderate	80%	60%	60%	30%	Moderate
8	Moderate	80%	60%	60%	30%	Strong
9	Strong	80%	60%	60%	30%	No
10	Strong	80%	60%	60%	30%	Moderate
11	Strong	80%	60%	60%	30%	Strong

behind intensive testing is obviously very requiring, yet it has the advantage of representing a polar case against which milder degrees of testing could be measured.

We are now ready to explore the impact of several types of exit strategies whereby the government relaxes the initial lockdown with a view to restarting the economy and mitigating the economic costs of people's confinement. Each of them is initiated on May 3, by assumption. Table 1 below depicts the configurations of the model's parameters that correspond to the exit strategies we wish to consider.

In the first four scenarios, we assume that the economy, the schools, and the society are fully re-opened. Scenario 1 is a wild scenario in which this complete re-opening is unaccompanied by any public health policy or safety measure. It is not quite credible either since the underlying assumption is that people go back to the interaction patterns that existed before the appearance of the coronavirus as if nothing ever happened. In scenario 2, the government opts for an ambitious testing program that allows for 5% of the infectious people to be detected and immediately isolated on each day. Yet, no social-distancing measure is imposed (or self-adopted by the people) in contrast to scenario 4 in which strong measures of that type are imposed (and/or self-adopted) in parallel to the testing of the population. Scenario 3 is the inverse of scenario 2 in the sense that no testing is undertaken but strong social

distancing is imposed (and/or self-adopted).

In all other scenarios, the economy, the schools, and the society are only partly re-opened. As for the other two policy variables, they are varied in several manners. Scenario 5 allows us to see what happens when the absence of social-distancing measures is combined with strong testing while scenario 6 considers moderate social distancing with zero testing. In the two subsequent scenarios (7 and 8), testing is increased to a low (1%) and a high level (5%).¹⁷ The last three scenarios repeat the same sequence of increasing τ from 0 to 1 to 5%, but this time in combination with strong social-distancing measures.

Four outcomes are of interest: (1) the peak of the infection measured by the global maximum of infected people; (2) the total death-toll of the pandemic, and its disaggregation by age class; (3) the time interval required to control the epidemic; and (4) the time interval required to build herd immunity, defined as the point at which only 30% to 40% of the population is still susceptible to the virus. Since there is a lot of controversy about the extent to which immunity is built against the covid-19 after recovery from infection, we leave (4) aside in our discussion. An agreement seems nevertheless to be slowly emerging that patients with covid-19 are actually building acute antibody responses against the virus. Thus, a study of 285 patients in China reveals that within 19 days after symptom onset, 100% of them tested positive for antiviral immunoglobulin-G (IgG) (Long et al., 2020). Incidentally, such evidence would validate the assumption behind the SEIR model according to which, once recovered, an individual cannot be contaminated again (at least in the short run). We will offer some comments on (2) but, given the uncertainty on the mortality rate of Covid-19 in Belgium, we will not show precise numbers in order to avoid unnecessary misinterpretations.

¹⁷It is worth noting that Belgium shows a comparatively high degree of testing when compared to other countries by the end of April: 41,000 tests per million inhabitants. In this respect, Belgium (and Spain) lag behind a few countries only: Israël (48,000), Portugal (46,000), and Ireland (43,000). These figures can also be compared to Italy (37,000), Germany (30,400), the United States (23,000), the United Kingdom (20,000) and France (16,000). Even countries that have effectively resorted to testing-cum-tracing, and started to do it very early, have tested on a comparatively small scale (e.g., South Korea with a ratio of only 12,000).

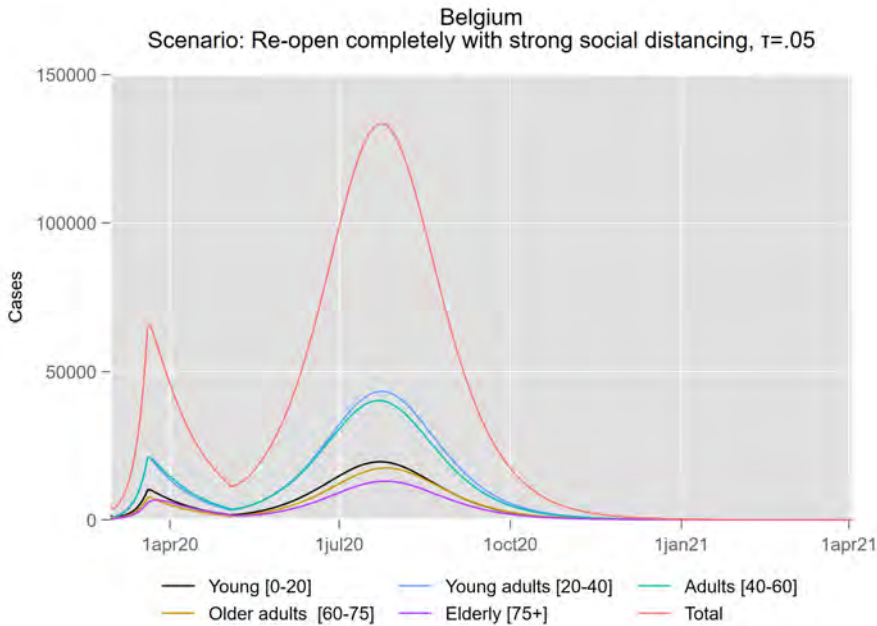
4 Results I: Comparing different exit strategies

Our results are crystallized in the infection curves drawn for each and every scenario, and those curves are also disaggregated for each broad age group. Each figure is simply labeled Fig. Si , where i corresponds to the number of the corresponding scenario. However, in order to save space and ensure readability, we will present in the main text only the figures that relate to four particularly meaningful scenarios: scenarios 4, 7, 10, and 11. Figures relating to all the other scenarios are displayed in Appendix 1. We can now bring out the main results that emerge from our numerical simulations.

4.1 Scenarios of complete re-opening of the economy and society

Clearly, the option of abruptly relaxing the lockdown on May 3, as envisaged under scenario 1, would be a disaster. There is a huge rebound that would lead to a peak very much higher than the peak reached under the lockdown, and the mortality would certainly shoot up to intolerable levels, especially for older people. Does systematic testing or strong social distancing allow the country to avoid a rebound while re-opening the whole economy, schooling and society (see scenarios 2 and 3)? The answer is negative: if the magnitude of the rebound is mitigated, and the death-toll is not expected to rise as much as before, the situation remains extremely worrying. When the government resorts to both measures simultaneously instead of separately (under scenario 4), the epidemiological effects of the complete re-opening of the economy are somewhat improved, yet not enough to suppress a still serious rebound. In sum, any scenario based on a fast and complete re-opening is out of question.

Figure 4: S4



4.2 Scenarios of partial re-opening of the economy and society

4.2.1 Scenarios with strong testing and varying severity of social-distancing measures

Here, we are concerned with scenarios 5, 8, and 11, in which $\tau = 0.05$, and the severity of social-distancing measures goes from zero to strong through moderate.¹⁸ With scenario 5, a rebound continues to exist but the second peak is now smaller than the initial peak. When we compare this scenario with scenario 2, we see that the extent to which the government

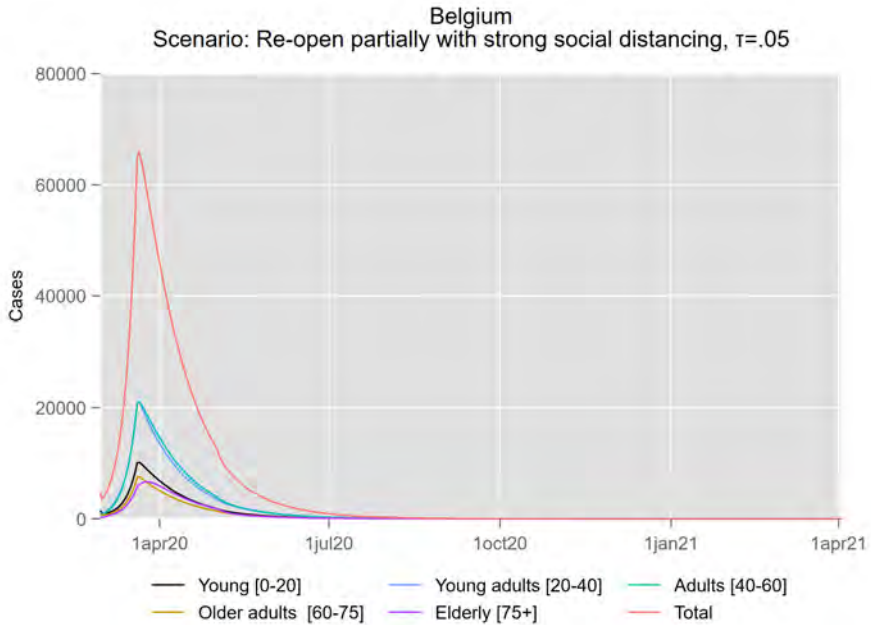
¹⁸Note that identifying and quarantining 5% of the infectious every day corresponds to a very ambitious scenario.

chooses to re-open the economy has a major impact: under scenario 5, the number of infections at the second peak is only a small fraction (less than 10%) of what it is under scenario 2, and the total number of expected deaths is also much lower.

It is evident that the government cannot rely exclusively on a strong testing programme to box the epidemic, even assuming that it chooses to re-open the economy only partially. Let us now assume, as in scenario 8, that it concomitantly imposes moderate social-distancing measures. What we find is that the rebound has vanished and that, compared to scenario 5, the number of deaths should be considerably reduced. But there is a bad news: it is not until April 2021, about a year after the first peak, that the epidemic will be brought under control.

Since this is obviously too long a period, the government has no other choice than imposing work from home whenever possible as well as strong social-distancing measures. Among other things, the latter involve forcing people to wear a protective mask when going to public places, including workplaces, shops, parks, etc. This is scenario 11 in which public health policies are at their maximum levels while only part of the economy is re-opened. This scenario is truly encouraging in the sense that there is no rebound and, moreover, the control of the epidemic is now advanced (compared to scenario 8) by no less than eight months: we indeed expect that by September 2020, the epidemic will be controlled. In addition, the number of dead as a result of covid-19 is expected to be significantly reduced. Like in all the previous scenarios, a huge proportion of the deceased should belong to the age classes above 60 years.


Figure 5: S11



4.2.2 Scenarios with strong social-distancing measures and varying levels of testing

The next question to ask is whether the government could still obtain satisfactory results if it is not in a position to practice testing on the ambitious scale assumed in the three aforementioned scenarios. The relevant scenarios are numbered 9 and 10: while there is no testing under scenario 9, $\tau = 0.01$ under scenario 10. Note that we have already looked at the case where $\tau = 0.05$, which corresponds to scenario 11).

Figure 6: S10



Covid Economics 23, 28 May 2020: 1-57

What do we see? If there is no testing but stringent social distancing is enforced, the effect is not as bad as could have been expected: there is no rebound but the epidemic will not be brought under control until early April 2021, as compared to early September 2020 under scenario 11. Moreover, the death-toll is expected to be only slightly higher than under the latter, best scenario. The explanation is that, if social-distancing measures are strong and well enforced, contamination is rather limited since infectious asymptomatic individuals have rather few opportunities to propagate the virus. When testing is improved to reach a scale of 1%, the situation is not fundamentally different even though the control of the epidemic now occurs a few months earlier (early January 2021). This confirms the result that once strong social distancing is in place, the effectiveness of stronger testing measures is seriously dampened. Does this imply that testing should be a low priority for

a government? The answer is negative for two main reasons. First, there is an intrinsic value for a country in acquiring sound knowledge about the scope and the dynamic of the epidemic, and this requires at the minimum that a strongly representative sample of the population be tested for both the presence of the virus (RT-PCR testing) and for the possible presence of antibodies (see Dewatripont et al., 2020, and Gilbert et al., 2020, for a defense of the case of joint testing). Second, social-distancing measures are very hard to enforce and, therefore, scenarios with less than strong social distancing are probably more realistic. This is the task to which we now turn.

4.2.3 Scenarios with moderate social-distancing measures and varying levels of testing


Here, we assume that the government is unwilling or unable to impose strong social distancing including the systematic use of protection masks in public spaces, that is, we assume that it can only impose moderate social-distancing measures. The scenarios concerned are numbered 6, 7, and 8 in which the value of τ is gradually raised from zero to 1%, and then to 5%. Note that we have already commented on scenario 8 (see subsection 4.2.1).

Under scenario 6, the rebound is not avoided, and the second peak number of infections (occurring only in early September 2020) is roughly as high as the peak reached under the lockdown. Furthermore, the death-toll should be very large as the number of cases is expected to remain high for a relatively long period. The considerable benefit of stringent distancing measures is thus manifest when we compare this scenario to scenario 9 in which such stringent measures are imposed (and τ is kept at zero) and the rebound is avoided. Under scenario 7, the rebound continues to occur but the number of infections at the new peak is reduced by almost 50% relative to its number under the lockdown. However, the death-toll is likely to remain intolerably large, testifying to the importance of having an effective testing programme when the government is unwilling or unable to impose strong distancing. The reference here is to scenario 8 in which no rebound exists. As for the number of dead, it remains high at a level expected to be about twice the number observed in the best scenario (where stringent social-distancing measures are added to strong testing).

The conclusion is that when strong social-distancing measures cannot be enforced, ambitious

testing programmes need to be put in place to (partly) substitute for the missing distancing discipline. It is a fortunate feature of our model that the marginal impact of more intense testing on infection (and, consequently, mortality rates) is stronger when social-distancing measures are weaker. In the presence of moderate (or poorly enforced) social-distancing measures, an increase in the scope of testing yields larger effects than in the presence of strong social-distancing measures.

Figure 7: S7

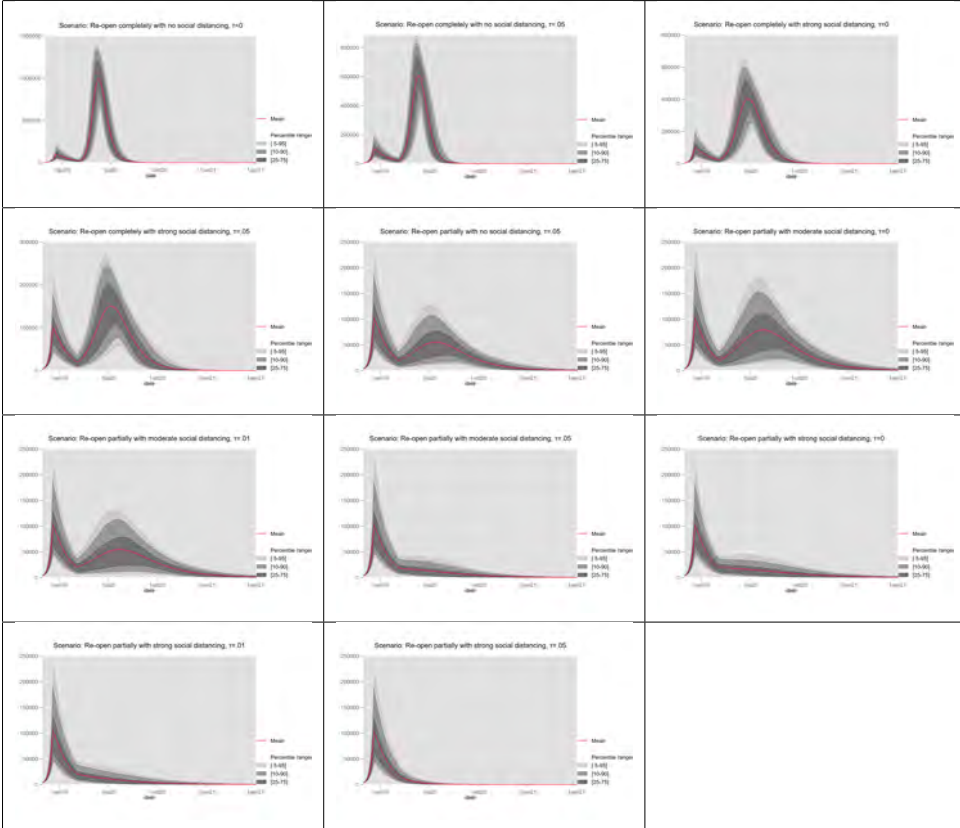


4.3 Sensitivity analysis

Given the high degree of uncertainty regarding the values of the parameters chosen, it is important to do a sensitivity analysis and thereby check that our results are not driven by a

specific combination of values. Once again, our objective is not to make predictions but to compare stereotyped scenarios. To do so, for each of the scenarios, we run 1,000 simulations randomly generating parameters from some specific distributions. In graphing the results, we shade in gray, by decreasing intensity, the areas defined by the i) 25-75 percentile range, ii) 10-90 percentile range and iii) 5-95 percentile range for the simulated values for each day. Parameters are generated from: $d_{IA} \sim U[9, 11]$, $d_{IS} \sim U[2, 4]$, $d_L \sim U[1, 3]$, $\rho \sim U[0.7, 0.9]$ and $\alpha \sim U[0.125, 0.375]$, where U stands for an uniform distribution. Furthermore, we generate the first element of vector z from a $U[0.2, 0.8]$, meaning that we allow β_i to increase substantially for the young while it is concomitantly decreased for the elderly so as to guarantee that the average β remains unchanged. We thus assume that positive cases for the young are heavily under-reported. Below, we present the results for total infections.

Table 3: Sensitivity analysis



Covid Economics 23, 28 May 2020: 1-57

As can be seen in our series of graphs in Table 3, the results do not seem to be significantly influenced by the combination of parameters used. In other words, other parameter combinations tend to yield trends in infections similar to those presented in Section 4.1 and 4.2. This is true even though the value of simulated infections may be different, which is an expected outcome since we have substantially raised the value of β_i for the young.

5 Results II: The role of culture


We are now ready to address the second and most important issue that motivated this paper: how does culture impact the Covid-19 epidemic? In line with our preliminary work under Section 4, we explore how different cultural habits change the lockdown exit paths under the different scenarios considered. Toward that purpose, we compare the benchmark country case of Belgium with two pseudo-cases in which the contact matrix of two other countries, Italy and Germany, are substituted for the Belgian matrix. In other words, we carry out a mind exercise whereby we inquire into the following question: how would the effects reported in Section 4 be affected if, when exiting the lockdown, Belgium was endowed with other cultural habits than those it actually inherited from the past? In a first exercise, the substitute contact matrix is borrowed from Italy, giving rise to a pseudo-Italian case, whereas in the second one, it is borrowed from Germany, giving rise to a pseudo-German case.¹⁹ Everything else is kept unchanged and, in particular, the infectious rates, β_i , remains those of real Belgium. Because of the presence of many confounding factors, this hypothetical approach seems more reasonable than the alternative approach consisting of comparing Belgium with Italy and Germany as simulated on the basis of their own premises.

That cultural differences may account for observed contrasts in the incidence of covid-19 has been noticed repeatedly in newspaper articles. Thus, the Japanese habit of keeping reasonable distances between interacting people strikingly contrasts with the Western European habit, especially in southern Europe, of kissing and hugging friends, relatives, and acquaintances. Moreover, in some countries like South Korea, China, and Japan again, people are accustomed to wearing face-masks as a way to protect themselves against air pollution, an attitude which is an oddity in Europe. In Vietnam, too, social comfort with wearing face-masks and acceptance of being isolated away from home have played a significant role (Economist, 2020, 9-15 May : 41). Because these East Asian attitudes are conducive to effective protection against contamination, they are a big help under the extraordinary circumstances of a pandemic.

¹⁹Note that matrices are not symmetric but need to satisfy reciprocity. See Mossong et al. (2008), Prem et al. (2017) and Towers and Feng (2012) for further details.

There is yet another important sense in which cultural variations do matter, and they relate to the frequencies of contacts between people. For example, the Italian society is strongly centered on the family with the consequence that relatives pay frequent visits to each other, implying in particular that children and grandchildren often visit their grandparents. These are attitudes mutually expected by Italian (and other Mediterranean) people, and they therefore constitute social norms. They contrast with the norms prevailing in northern European countries, Germany and Scandinavia, for example, where interpersonal contacts are not only more distant but also less frequent. These variations in contact habits should be on display when comparing the social interaction matrices of different countries. In figure (8), we compare the matrices of three European countries: Germany, Belgium, and Italy.

Figure 8: Interaction matrices




It is immediately evident that the contact matrix of Belgium is less dense than that of Italy, yet more dense than that of Germany. This is true both on and off the diagonal, meaning that not only contacts inside particular age groups, but also across them, differ significantly between the three countries, with Belgium occupying an intermediate position. The most striking feature is that, in accordance with the above-noted difference in the place of the family in society, children tend to pay more visits to their grand-parents in Italy than in Belgium, and even more than in Germany. How these differences get translated into the epidemiological effects of different exit scenarios is a matter of great interest since it raises the issue as to how a critical social variable that is obviously exogenous impinges upon the effectiveness of public policies in the circumstances of the Covid-19 crisis.

As has been already pointed out, our hypothetical approach is applied to the period succeeding the relaxation of the lockdown. We could also apply it to the period that also includes the lockdown time. Since the model has been initiated on the basis of the Belgian situation, which corresponds to the lockdown, the first option seems more sensible, and this is why we present the corresponding results below.

As in Section 4, our results are presented in the form of figures featuring the infection curves for each scenario. To facilitate the comparison, we have drawn the curves corresponding to the three country cases (one real and two pseudo) on a same graph. Yet, in order not to overburden the figures, we do not show decomposition by age group. The figures are now labeled Sci , where $i = 1, 11$. Like in Section 4, in the main text we only show the figures corresponding to scenarios 4, 7, 10, and 11 with the figures related to the seven other scenarios being displayed in Appendix 2.

The results are striking. Look at scenario 7 under which partial re-opening is combined with only moderate social-distancing and testing measures. (While in the case of real Belgium a rebound occurs and the second peak of infections is lower than the initial peak attained under lockdown, a severe rebound is observed in the pseudo-case of Italy where the second peak is considerably higher than the initial peak. In the case of pseudo-Germany, by contrast, there is no rebound and the infection curve quickly flattens. In other words, if Belgium was endowed with German contact habits, it could exit lockdown smoothly by contenting with moderate rather than strong public health measures in the event that the economy is only partially re-opened. Obviously, the big differences observed in infections imply huge differences in the respective death-tolls. While the number of dead remains very low in pseudo-Germany, it is intolerably high in pseudo-Italy and very high in (true) Belgium.

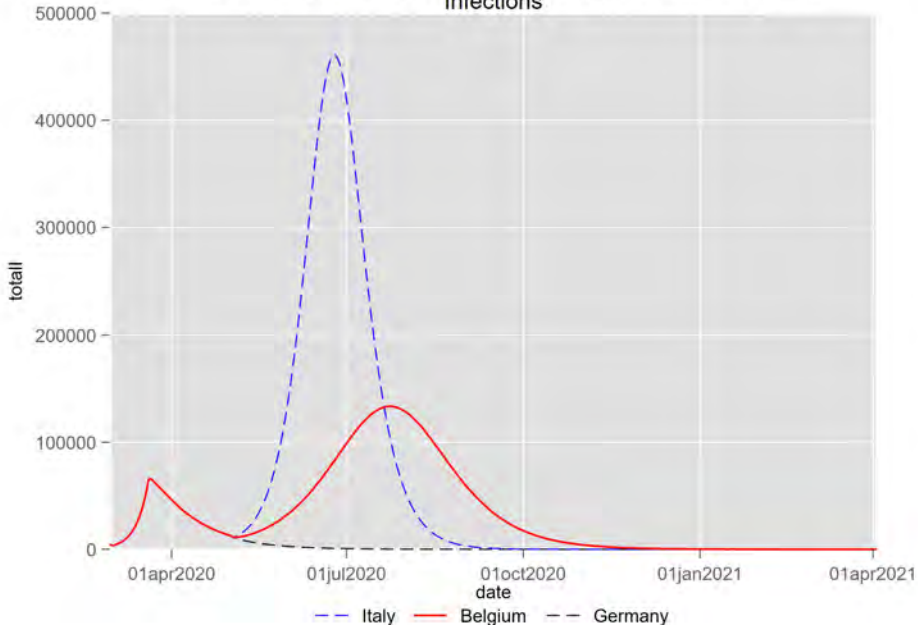
Figure 9: Sc7
Scenario: Re-open partially with moderate social distancing, $\tau=0.01$
Infections



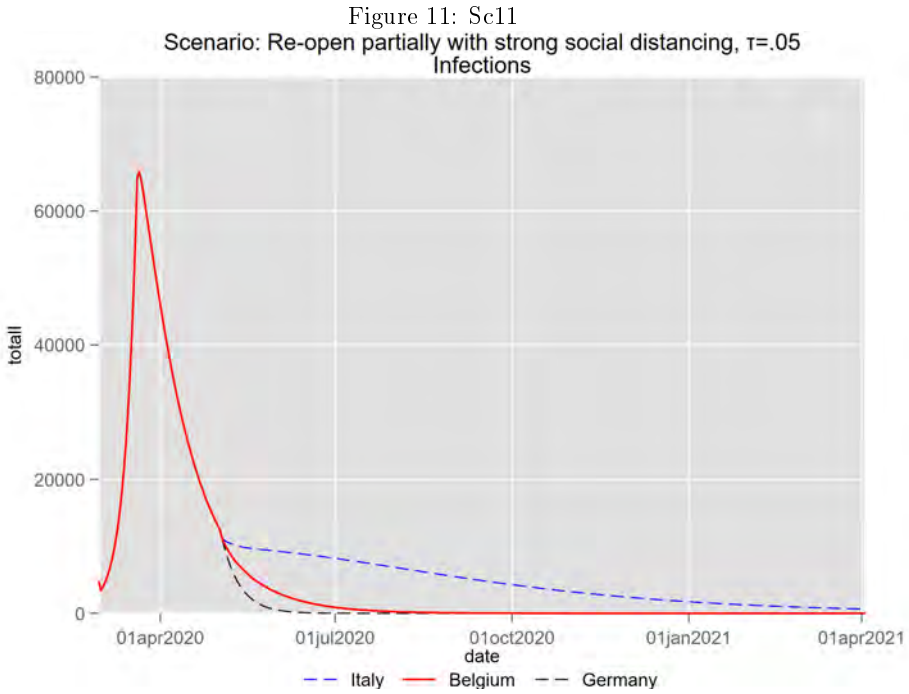
Covid Economics 23, 28 May 2020: 1-57

We have learned from Section 4 that (in our setup) all scenarios of complete re-opening are unrealistic in the case of real Belgium. It is now evident from figure 10, that pseudo-Germany could exit the lockdown by completely re-opening its economy without causing a rebound. This would nevertheless require that some public safety measures are imposed, social-distancing measures, in particular. With zero social distancing, pseudo-Germany would experience a rebound even with a moderate amount of testing, as can be seen from the simulations under scenario 1 and even scenario 2 where the rebound is quite flat, though (see Appendix 2). Clearly, when the rebound occurs in pseudo-Germany, its magnitude is small compared to the one which real Belgium, and pseudo-Italy even more so, would experience under the same conditions.

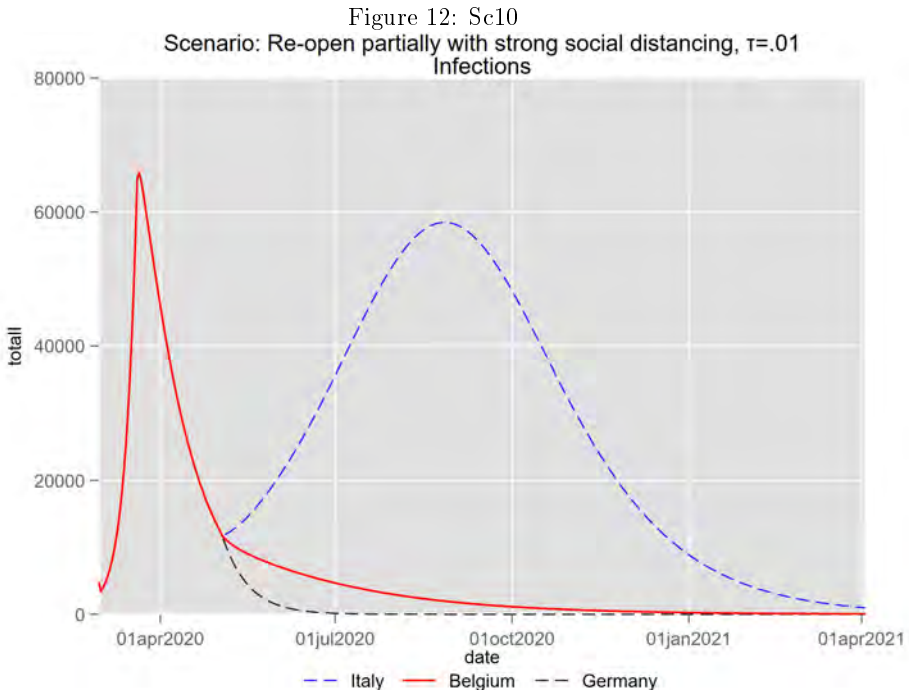
Figure 10: Sc4
Scenario: Re-open completely with strong social distancing, $\tau=.05$
Infections



If we now look at the most requiring scenario, scenario 11, in which the government combines strong social-distancing measures with ambitious testing (and only a partial re-opening of the economy), we find that pseudo-Germany is able to control the epidemic a few months earlier than Belgium (early June rather than early September 2020). The situation for pseudo-Italy is much less encouraging since even in these ideal conditions, it will not be able to avoid a long-lasting, slow decline in infections, and the death-toll will be much heavier than what we find for pseudo-Germany and real Belgium.



Finally, considering scenario 10 under which strong distancing measures are kept yet testing is less ambitious ($\tau = 0.01$), we see that pseudo-Italy goes through a second peak of infections almost as large as the initial one. By contrast, no rebound is observed for pseudo-Gemany and real Belgium. In terms of death-toll, pseudo-Italy would have to accept a considerable cost compared to the other two country cases. There is a noticeable difference between real Belgium and pseudo-Gemany in the sense that the time needed to get the epidemic under control will be much longer, by half a year, in the former than in the latter case.



The central lesson is that, from the standpoint of ability to battle against the covid-19 in a country like Belgium, the problem of public authorities would have been easier if the prevailing contact habits were those of the Germans rather than those of the Belgians themselves. But their problem would have been much more tricky to solve if the contact habits were those of the Italian people. Since cultural habits and social norms cannot be changed by government fiat, there is no escaping the fact that countries with a comparatively dense social interaction structure must be especially careful in the measures that they impose on their society. If it appears that stringent distancing measures are harder to implement in these countries for cultural, political or other reasons, testing will have to be especially effective lest the health cost borne by the population should be heavy.

Another interesting lesson to draw from our exercise is the following: even differences be-

tween contact frequencies that appear rather moderate (compare Belgium and Germany in Figure 8) get considerably amplified when made to play in the mechanical game of a highly contagious virus (which SARS-CoV-2 definitely is).

6 Conclusion and discussion

There are two central messages to extract from our simulation exercises. First, a reasonable exit strategy must involve a limited re-opening of the economy and the society, strict measures of social distancing and an ambitious and effective testing programme. These are apparently the conditions under which our benchmark country, Belgium, can hope to avoid a rebound in infections and follow a relatively quick path toward getting the epidemic under control. They are quite requiring since severe social distancing involves stringent social measures such as the wearing of face-masks and strong testing supposes an effective administration of RT-PCR tests on a large scale. The latter cannot be done unless many conditions are fulfilled in terms of logistics, availability of reagents, and large capacity for sample analysis. As for the former, it necessitates the intervention of a strong government ready to enforce measures that may be unpopular among some age groups, especially the youngsters and young and middle-aged adults. If the government lacks such authority or does not want to deploy the necessary body of inspectors and enforcement agents, the relevant scenario will not be the scenario nominally corresponding to the announced policy but the one associated with a milder version of it. This is true unless, of course, citizens display a strong civic sense (they are law abiding) or they gradually learn that it is their own interest to follow government prescriptions (or strong recommendations).

What bears emphasis is that cultural norms and habits play a significant role. There is thus a major contrast between the resistance put up against the imposition of protective masks as described in the news for some US states, on the one hand, and the documented ostracization of non-compliant persons in South Korea, on the other hand. Whereas in the former instance people tend to oppose a state-imposed rule in the name of individual freedom, in the latter it is out of a strong civic sense that people feel hurt by the behaviour of those who refuse to abide by a rule meant for the public good. Many European countries

fall in between these two polar attitudes, with German-speaking, Scandinavian, and eastern European countries believed to be leaning more toward compliance and southern European countries assumed to be more sensitive to limits to privacy and individual rights. In the latter countries, harsh rules supported by effective detection of rule-violation and sanctions can substitute for lack of self-disciplining, as attested by the lockdown experiences of Spain, Italy, and France.

A good news emerging from our simulations is that the marginal impact of testing on the epidemic increases as social-distancing measures become weaker. This implies that, in the presence of moderate (or poorly enforced) social-distancing measures, an increase in the scope of testing yields stronger effects than in the presence of strong social-distancing measures. If the government cannot increase testing enough to make up for the lack of strong distancing measures, the public health consequences of a (partial) re-opening of the economy will be severe and a rebound of the epidemic will be unavoidable.

The second lesson concerns the role of contact habits and the related social norms of conduct, which are another facet of a country's culture. There is considerable debate about why there exist large differences between countries with respect to the incidence of the covid-19 epidemic. For example, why is it that Germany, Austria, and Norway are so successful (at least for the moment) compared to other Western European countries, or that Eastern European countries are less affected than their Western counterparts, or that southern Italy has come out much better than northern Italy, or German-speaking Switzerland than French-speaking Switzerland and French-speaking Switzerland than Italian-speaking Switzerland? As the debate is going, attention tends to be focused on the effectiveness of public interventions. Germany, in particular, is singled out for having been remarkably well prepared for the epidemic and for having an efficiently decentralized system in place to carry out a high number of tests relatively quickly.

Whether this is the whole explanation is a moot question. Leaving aside the sometimes huge discrepancies in the ways of detecting infections and counting deaths, we still do not know much about which part of the inter-regional differences in covid-19 infection and mortality rates is to be ascribed to genetic variations, which part to variations in social behaviour, and which part to variations in the effectiveness of public health interventions

and facilities. Recently, microbiologists from the university of Ghent in Belgium have argued that part of the differences in the intensity of the epidemic may be attributable to genetic variations. More precisely, some population groups exhibit a comparatively high frequency of the polymorphism D of ACE1 gene, probably leading to a modification of the activity of this enzyme that reduces the lethality risk (Delanghe et al., 2020). Interestingly, the more one moves toward the eastern parts of Europe, the higher the incidence of this favourable variant of the ACE1 gene. Not only Eastern European countries but also Austria-Germany, Scandinavia, and southern Italy (where the Norman conquest left its biological imprint) are included in the zone where the polymorphism is found. Spain, Northern Italy, France, Belgium, the Netherlands, and the United Kingdom are not.

Our own contribution lies elsewhere. Differences in the way people interact, and more specifically the frequencies of their contacts within and between age groups, seem to also account for variations in the incidence of the virus and performances in battling against it. If only Belgium could have had the contact matrix of Germany, it could achieve the objective of (partial) re-opening of the economy with more moderate policies than the ones it actually needs. And, conversely, if it had the habits and norms of Italy, it would have to take even more stringent measures lest the cost to bear as a result of economic re-opening should be (much) heavier. Whether these stringent measures will necessitate strict enforcement by the government or will be largely self-enforced by compliant citizens attentive to its recommendations is another question, albeit one that also involves cultural considerations.

This finding suggests that a country like Germany is probably cumulating all the advantages that work toward a successful lockdown exit: (1°) it possesses a strong public health infrastructure and has chosen sound public health policies that prepared the ground for an effective battling against the covid-19; (2°) its people probably evince genetic characteristics that make them less vulnerable to the virus; and (3°) the social norms that guide individual behaviour, including the habits regarding meetings and visits, help slow down an epidemic. If in ordinary circumstances, the comparatively weak role of the family in Germany may not necessarily be an advantage and the family-based model of Italy would perhaps seem preferable, the situation is modified in conditions of a raging epidemic when such a model is suddenly transformed from an asset into a liability.

Our finding might also help explain why the French-speaking part of Switzerland (Romandy) has epidemiological statistics close to France whereas its German-speaking part (Alemannic Switzerland) evinces strong similarity with Germany and Austria, and its Italian-speaking part (Tessin) strong similarity with northern Italy. This said, important variations, such as those observed between northern and southern France, remain unaccounted for. But there is a key lesson that we can definitely learn from our foray into the comparative effectiveness of lockdown exit strategies: there is no one-size-fits-all solution that could be uniformly applied to all countries and even to all regions inside a given country. It is perhaps not coincidental that the European Union has been unable or unwilling to suggest, let alone prescribe, a common lockdown exit strategy for all its members, leaving them free to make their own decisions in the matter. The diversity of peoples and cultures inside Europe is too large to allow for a general solution to the complex problems raised by the present pandemic. The same conclusion also applies to large federal political entities, India, Russia, and the United States, for example.

References

- Alfaro, L., E. Faia, N. Lamersdorf, and F. Saidi, 2020. “Social Interactions in Pandemics : Fear, Altruism, and Reciprocity”, CEPR Discussion Paper, DP14716.
- Boast, A., A. Munro, and H. Goldstein, 2020. “An evidence summary of Paediatric COVID-19 literature”, *Don't Forget the Bubbles*. Available at: <http://doi.org/10.31440/DFTB.24063>
- Bunyavanich S., A. Do, and A. Vicencio, 2020. “Nasal Gene Expression of Angiotensin-Converting Enzyme 2 in Children and Adults”. *JAMA*. Published online May 20, 2020. doi:10.1001/jama.2020.8707
- Cori, A., N.M., Ferguson, C. Fraser and S. Cauchemez, 2013. “A new framework and software to estimate time-varying reproduction numbers during epidemic”, *American Journal of Epidemiology*, Novembre,178(9):1505-12.
- Deforche, Koen. (2020). An age-structured epidemiological model of the Belgian COVID-19 epidemic. pre-print, KULeuven, 10.1101/2020.04.23.20077115.

Delanghe, J.R., M.M. Speeckaert, and M.L. De Buyzere, 2020. “The host’s angiotensin-converting enzyme polymorphism may explain edipemiological findings in COVID-19 infections”, *Clinica Chimica Acta*, June, 505: 192-193.

Dewatripont, M., M. Goldman, E. Muraille, and J.P. Platteau, 2020. “Rapidly identifying workers who are immune to covid-19 and virus-free is a priority for restarting the economy”, *Vox EU*, 2-3 March 2020.

Dietz K., 1993. “The estimation of the basic reproduction number for infectious diseases”. *Stat Methods Med Res*, 2: 23–41.

Economist, April 2020. “Hard choices”, April 4-10, pp. 13-17.

Economist, 2020. “Infectious Doubt”, 2-8 May, pp. 16-17.

Economist, 2020. “Bargain Abatement”, 9-15 May, pp. 40-41.

Eichenbaum, M.S., S. Rebelo, and M. Trabandt, 2020. “The Macroeconomics of Epidemics”, NBER Working Paper, WP26882.

Garibaldi, P., E.R. Moen, and C.A. Pissarides, 2020. “Modelling Contacts and Transitions in the SIR Epidemics Model”, *Covid Economics*, Issue 5, 16 April 2020, pp. 1-19.

Gilbert, M., M. Dewatripont, E. Muraille, J.P. Platteau, and M. Goldman, 2020. “Preparing for a responsible lockdown exit strategy”, *Nature Medicine*, Comment, Published online, 14 April 2020.

Guirkinger, C., and J.P. Platteau, 2020. “The Dynamics of Family Systems: Lessons from Past and Present Times”. In Baland, J.M., F. Bourguignon, J.P. Platteau, and T. Verdier (Eds), *The Handbook of Economic Development and Institutions*, Princeton and Oxford: Princeton University Press, Chap. 13: 449-519.

Ichino, A., C.A. Favero, and A. Rustichini, 2020. “Restarting the Economy While Saving Lives Under Covid-19”, CEPR Discussion Paper, DP14664.

Long, QX., et al., 2020. “Antibody responses to SARS-Cov-2 in patients with covid-19”, *Nature Medicine*, Brief Communication, Published online, 29 April 2020.

Mizumoto K, Kagaya K, Zarebski A, and G. Chowell, 2020. “Estimating the asymptomatic

proportion of coronavirus disease 2019 (COVID-19) cases on board the Diamond Princess cruise ship, Yokohama, Japan”, *Euro Surveill.* 25(10):2000180.

Mossong J., N. Hens, M. Jit, P. Beutels, K. Auranen, R. Mikolajczyk, et al., 2008. “Social contacts and mixing patterns relevant to the spread of Infectious diseases”, *PLoS Med.*, Mar 25, 5(3).

Prem K, A.R. Cook, and M. Jit, 2017. “Projecting social contact matrices in 152 countries using contact surveys and demographic data”, *PLOS Comput Bio*, 13 (9): 1-21.

Prem K, Y. Liu, T. W. Russell, A. J. Kucharski, R. M. Eggo and N. Davies, 2020. “The effect of control strategies to reduce social mixing on outcomes of the COVID-19 epidemic in Wuhan, China: a modelling study?”, *Lancet Public Health*, Mar 25: S2468-266.

Todd, E., 2011, *L'origine des Systèmes Familiaux*, Tome 1 : l'Eurasie, Gallimard, Paris.

Towers, S. and Z. Feng, 2012. “Social contact patterns and control strategies for influenza in the elderly”, *Mathematical Biosciences*, 240 (2): 241-249.

Ferguson et al. 2020. “Report 9: impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand”, *Imperial College COVID-19 Response Team Report*.

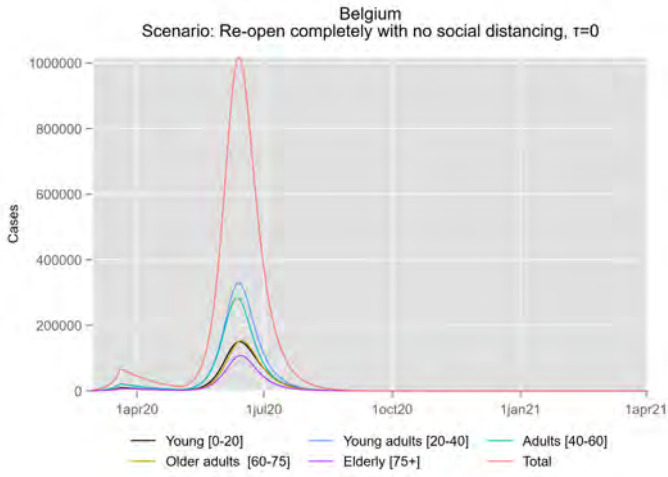
Wallinga J and P. Teunis P, 2004. “Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures”. *American Journal of Epidemiology*, 160 (6): 509-16.

Zhao S, Q. Lin, J. Ran, S.S. Musa, G. Yang, W. Wang, Y. Lou, D. Gao, L. Yang, D. He, M.H. 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.*, March; 92: 214-217.

Appendix

Appendix 1

Figure 13: S1



Covid Economics 23, 28 May 2020: 1-57

Figure 14: S2

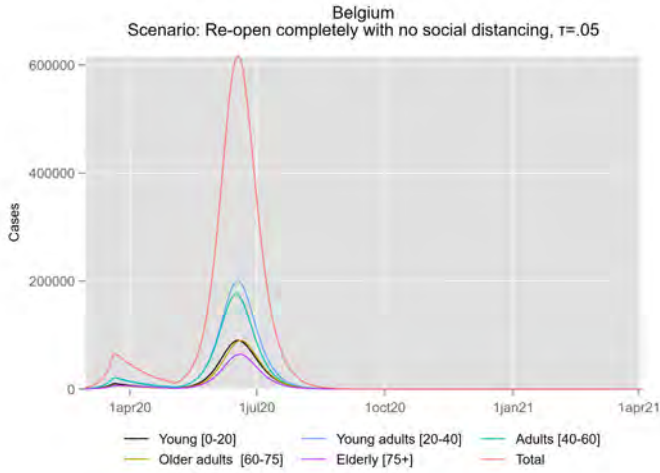


Figure 15: S3

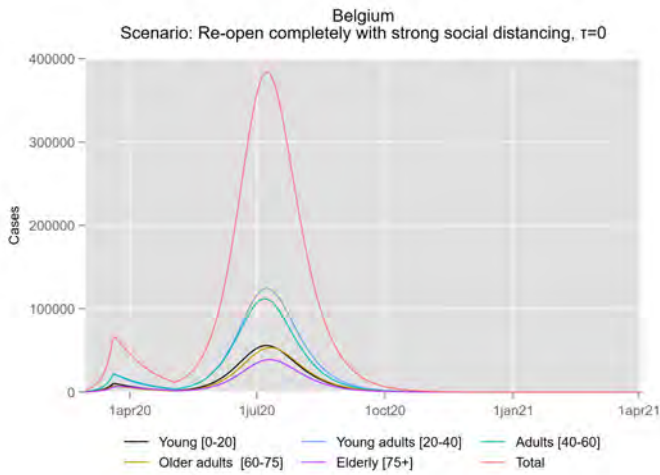


Figure 16: S5

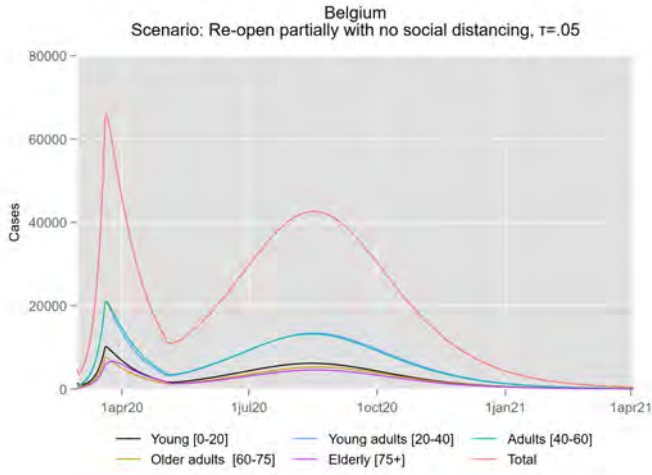


Figure 17: S6

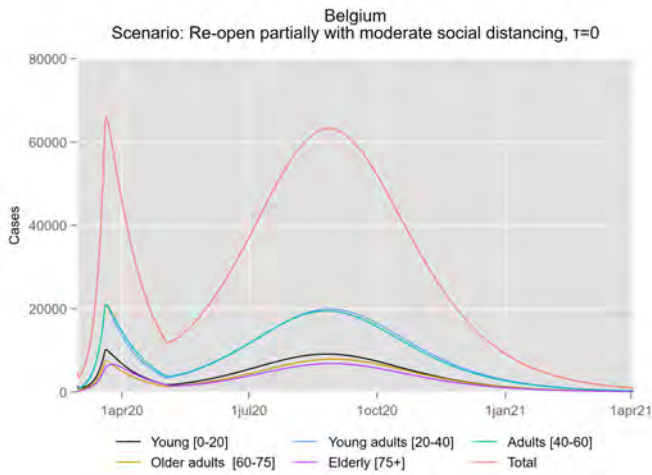


Figure 18: S8

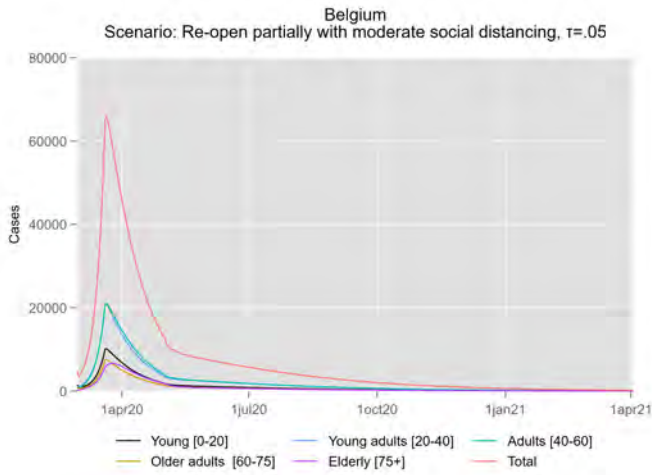
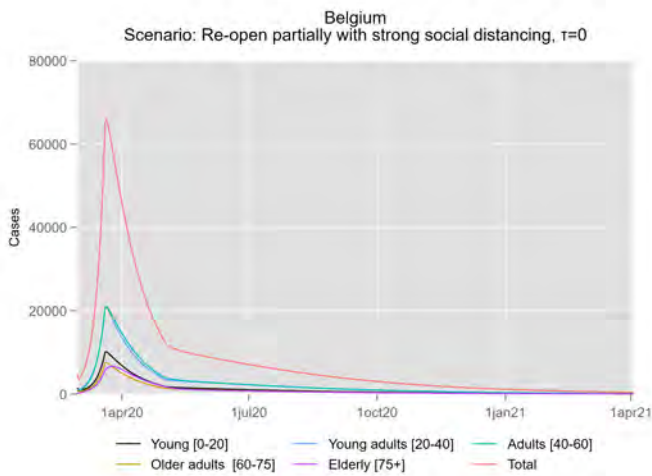
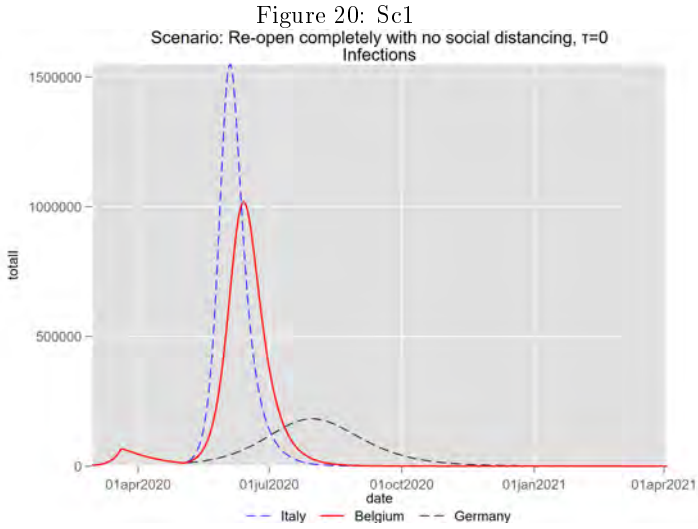


Figure 19: S9



Appendix 2



Covid Economics 23, 28 May 2020: 1-57

Figure 21: Sc2

Scenario: Re-open completely with no social distancing, $\tau=0.05$
Infections

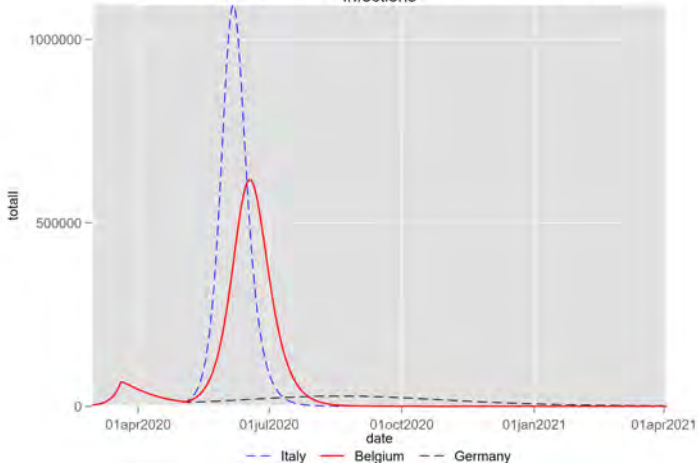


Figure 22: Sc3

Scenario: Re-open completely with strong social distancing, $\tau=0$
Infections




Figure 23: Sc5

Scenario: Re-open partially with no social distancing, $\tau=0.05$


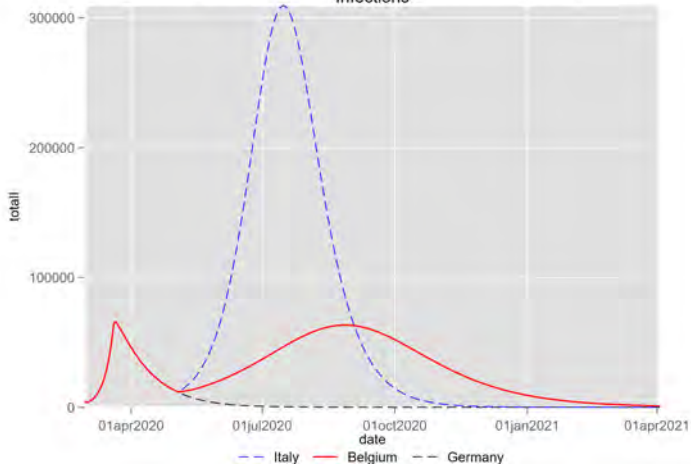
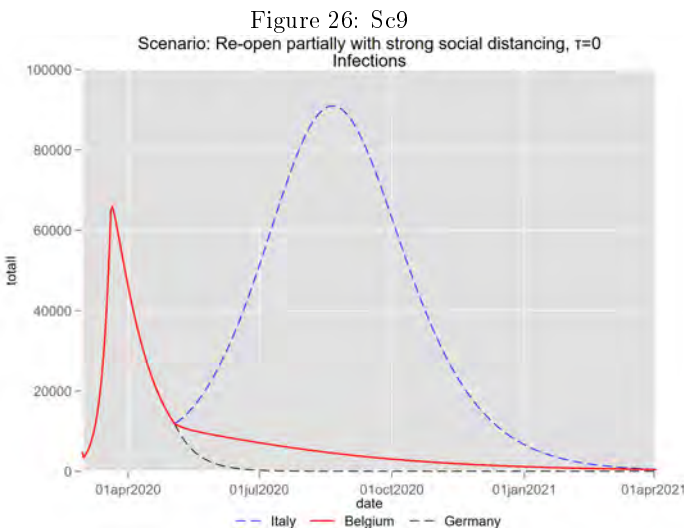
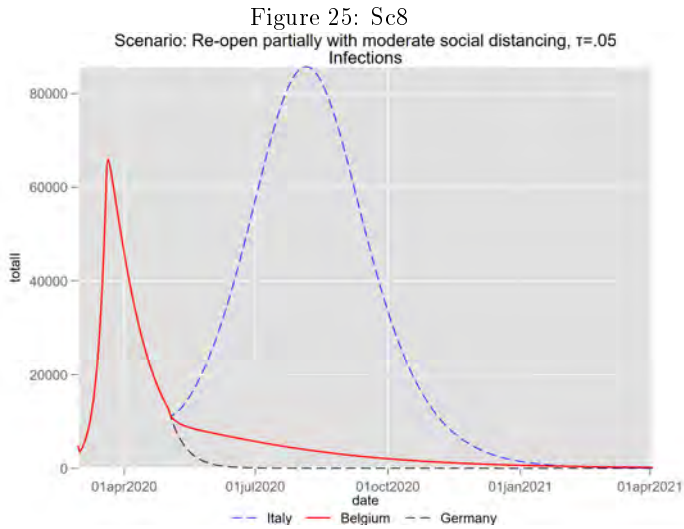


Figure 24: Sc6

Scenario: Re-open partially with moderate social distancing, $\tau=0$





Baby steps: The gender division of childcare during the Covid-19 pandemic

Almudena Sevilla¹ and Sarah Smith²

Date submitted: 22 May 2020; Date accepted: 22 May 2020

The Covid-19 pandemic has caused shocks to the demand for home childcare (with the closure of schools and nurseries) and the supply of home childcare (with many people not working). We collect real-time data on daily lives to document that UK families with young children have been doing the equivalent of a working week in childcare. Women have been doing the greater share, but overall, the gender childcare gap (the difference between the share of childcare done by women and the share done by men) for the additional, post-Covid-19 hours is smaller than that for the allocation of pre-Covid-19 childcare. However, the amount of additional childcare provided by men is very sensitive to their employment – the allocation has become more equal in households where men are working from home and where they have been furloughed/ lost their job. There are likely to be long-term implications from these changes – potentially negative for the careers of parents of young children; but also, more positively for some families, for sharing the burden of childcare more equally in the future.

¹ University College London.

² University of Bristol.

Copyright: Almudena Sevilla and Sarah Smith

1. Introduction

Since the COVID19 pandemic struck, governments around the world have introduced a range of non-pharmaceutical interventions in order to slow down the rate of transmission of the disease. In the UK, self-isolation measures for those with symptoms were imposed on March 12th, followed by social distancing measures encouraged for everyone on March 16th, school and nursery closures on March 20th and a general lockdown on March 23rd. These measures have brought about rapid and profound changes to people's everyday lives. For families with children, there have been major shocks both to the demand for home childcare and to the supply of home childcare. In the UK, following the example in many other countries, childcare providers (schools, colleges, nurseries and childminders) closed from Friday 20th March to all but the children of key workers and vulnerable children, requiring millions of children to stay, and be looked after, at home. Work life has also changed beyond recognition. Only a minority of people are working in their regular place of work. Wherever possible, people have been asked to work from home and to juggle work with childcare responsibilities if they have young children. With the required closure of most places of work, an estimated one-quarter of the UK workforce is on furlough (i.e. temporarily laid off and paid by the government at 80 per cent of their wages up to £2,500), while an estimated three per cent have lost their job altogether.¹ There are now many parents who are not working and who have more potential childcare time on their hands.

This paper provides new evidence on the combined effect of these demand and supply shocks on the gender allocation of childcare within couples. To that end we collected real-time data on the division of childcare within households pre- and post-lockdown. We document that in normal times (pre-COVID19), childcare is unequally distributed between men and women. Own calculations using evidence from the 2015 UK Time Use Survey (TUS) shows that mothers of young children (aged ≤ 12) spend about 2 hours per day in childcare during a normal weekday or weekend day, whereas fathers spend around 45 minutes on a weekday and about an hour and fifteen minutes on a weekend day. We find a similar gap in our survey; pre-COVID19, women in couples did 65 per cent of childcare, equating to a 30 percentage point gender childcare gap (i.e. 65 – 35).

We estimate the within-household change in the allocation of childcare post-COVID19. Estimating within-household changes is important as it allow us to control for unobserved heterogeneity that might be correlated with both post-COVID19 employment outcomes and childcare allocation. We confirm that there has been a dramatic increase in the total amount of childcare provided at home. A typical family with young children (aged ≤ 12) in our sample is now doing an average of 40 hours additional childcare each week that would previously have been provided by external providers. This

¹ Government figures indicate that 7.5million workers had been furloughed by 13th May 2020 out of a total workforce of 33 million (28 million employed and 5 million self-employed) as of January 2020. The number of people claiming unemployment benefit increased by 865,000 in April 2020.

is equivalent to an additional working week in childcare, and most of it is being done by women. On average, women do around ten more hours a week than men. Women have been more likely than men to lose employment as a result of the pandemic, but this does not explain all the gender gap in additional hours' childcare. Indeed, the amount of childcare provided by women is less sensitive to their own employment than it is for men.

On a more positive note, the allocation of additional childcare hours is more equal than the pre-COVID19 allocation of childcare. Estimates of the change in the difference in (female/ male) childcare shares indicates a small (c. 10 per cent) reduction in the gender childcare gap since lockdown. However, there is considerable variation by men's employment status – there is a small change to a more equal allocation when men work from home and bigger changes when men are furloughed/ not working. Hence any move towards a more equal allocation of childcare has been driven by the supply-side shock (more time not working) rather than the demand-side shock (the increase in childcare need); the additional burden of childcare is only shared more equally when men have more time on their hands.

This paper is related to recent studies that have discussed the gender impacts of the COVID19 pandemic. Most of these focus on the effects on male/female employment. Alon et al (2020) study pre-COVID19 employment and childcare in the US and make predictions about the likely impact of the pandemic. They predict that the negative employment effects of the pandemic for women are likely to be worse than those of a typical recession because of the impact of lockdown on retail and leisure industries, sectors that have a high female share. They also suggest that higher childcare needs will be a burden for women. Analysing data from the early days of the pandemic, Adams-Prassl et al (2020) find that women experienced a bigger drop in employment in the US, Germany and the UK; analysing time-use data, they also find that women do more of the additional childcare than men. Alon et al (2020) and Hapucheck and Petrongolo (2020) speculate that there may be some households where men do more childcare and that, following some of the evidence from paternity leave policies, the increase in man's childcare may have a positive effect in the longer term by changing social norms.²

Compared to these studies, our contribution is to show the actual (distribution of) changes in childcare allocation that have occurred since the pandemic started within a given household. Our paper is most closely related to preliminary work by Gonzalez and Ferre (2020). Using a self-selected sample of Spanish households, they show that there has been a shift to a more equal distribution of housework

² The quasi-experimental evidence on the effects of paternity leave on household specialization is not clear cut. Farré and González (2019) and Tamm (2019) show that paternity leave leads a persistent increase in fathers' involvement in childcare in the case of Spain and Germany respectively. However, Ekberg, Eriksson, and Friebe (2013) do not find an effect of "daddy months" in Sweden in father's likelihood to take medical leave to care for children.

(driven mainly by men taking responsibility for shopping) and childcare. Compared to this study, our contribution is to provide complementary analysis for a different country (UK not Spain) and to analyse a representative sample of households (Gonzales and Ferre analyse a self-selected sample). We also look at how the change in the allocation of childcare within households of childcare relates to both the demand- and supply-side shocks.

This paper also contributes to an earlier literature that studies the effect of unemployment shocks on childcare and housework. Aguiar et al (2013) study the effect of the recession following the financial crisis on time use. Using the American Time Use Survey (ATUS), they find that men and women increase their non-market work as the probability of unemployment increase. About 5 per cent of foregone market work is reallocated to childcare, and women tend to reallocate more of foregone market work to core production activities (e.g., cooking, cleaning, laundry), whereas in the sample of men foregone market work hours are relatively more likely to be reallocated to watching TV. This study cannot explicitly study the within-household allocation as the ATUS survey only asks one member of the household about their use of time. Gimenez et al (2009) use the Spanish Time Use survey (STUS) to show that unemployment increases own-time devoted by men and women to childcare and housework activities without affecting the time spend by the other partner to these activities. In the case of COVID19, there is simultaneously a demand-side shock as well as a supply-side shock, making it a particularly interesting setting to study the allocation of childcare within the household.

2. Sample and variables

2.1 Sample

The questions were asked by Ipsos MORI as part of their regular omnibus survey. Interviews were conducted online with 4,341 respondents aged 18-60 between the 5th and 11th May 2020. Quota controls were set upon the interviews achieved and the resultant survey data are weighted to the known offline population profile of this audience (18-60).

The total sample with non-missing gender is 4,250 individuals. In section 3 we analyse the employments effects of COVID19, testing to see whether the impact has been the same for men and women, following Adams-Prassl (2020). For this analysis, we focus on a sub-sample of 2,782 respondents who were employed prior to 23rd March and with non-missing information on employment characteristics. For the analysis of the gender childcare gap we focus on a smaller subsample of respondents who are in couples and who have children aged ≤ 12 .

2.2 Key variables

In order to capture changes pre and post-COVID19 we asked respondents about their work arrangements before and after the lockdown on March 23rd, the allocation of childcare within couples pre-COVID19 and about the number of additional hours, as well as their allocation, post-COVID19. As part of the omnibus survey, we have general demographics such as age and gender, as well as household socio-economic characteristics such as the number of children in the household below the age of 18, the age of children, respondent's educational attainment, and occupation categories.

We briefly summarise the questions used to elicit the key employment and childcare variables. Summary responses are in Table 1.

Employment status and working from home

Pre-COVID19

*“Which of the following best applied to your [and your partner’s] employment status **before** the coronavirus pandemic measures came into effect (prior to 23rd March 2020)?”* Responses are: In (full or part time) employment; Self-employed; Not in (full or part time) employment; Other.

*“Still thinking about your employment **before** the coronavirus pandemic measures came into effect (prior to the 23rd March 2020). In general, during the times you were working, how frequently, if at all, did you work from home?”* We group the responses (Working from home all the time; Working from home at least once a week; Working from home at least once a month and Working from home less often) to create a single, ever worked from home indicator.

Post-COVID19

“Which of the following, if any, best describes your [and your partner’s] employment status after measures against the coronavirus came into effect on the 23rd of March 2020?”

Responses are: In employment, working from home all of the time; In employment, working from home some of the time; In employment, working in a workplace elsewhere all of the time (not working from home); Furloughed (temporarily laid off with pay); Not employed; Other. We group the two working from home responses in our analysis.

Childcare

Pre-COVID19

*“Now thinking about childcare arrangements while education and childcare settings were **still open** (that is, during term times before the 20th March 2020) for your children aged 17 and under. Who took care of your children when they were **not** in an education or childcare*

setting?” Each respondent with a partner asked to give one option for each of self, partner and other out of: All or almost all of the time; Most of the time (about three quarters); About half of the time; Less than half of the time (about a quarter); None or almost none of the time; Don’t know. We use this information to calculate the allocation of childcare within couples (excluding Other) pre-COVID using values of 0.9 for all or almost all of the time, 0.75 for most of the time, 0.5 for about half of the time, 0.3 for less than half of the time and 0.1 for none or almost none of the time.

Post-COVID19

Thinking now about the education and/or childcare that education and childcare settings normally provided for your children (e.g. schools/colleges, nurseries, after school clubs, childminders, etc.):

How many extra hours, if any, are you (and partner) personally having to look after your children on each day during a typical week? Please give your answer to the nearest hour and if unsure, please give your best estimate. If not spending any extra time, please put ‘0’. We use the information on the additional number of hours done by self and partner (winsorized to 12) to calculate the allocation of additional childcare within couples and the number of additional hours done by men and women. This question is asked for each day of the week, including weekends.

3 Employment changes since COVID19

We first look at changes in employment. Focusing on the early phase of the pandemic, Adams-Prassl et al (2020) report that 15 per cent of their sample of UK workers lost their jobs (they do not differentiate furlough from no longer working). In our survey we distinguish between those who are on furlough (i.e., who are employed but temporarily not working and paid by the government up to 80 per cent of their wages up to £2,500), and those who report that they are not working and not on furlough.

Table 1 reports summary statistics on levels of employment among the whole sample (i.e. not conditioning on pre-COVID19 employment). The numbers highlight the reduction in employment that has occurred and the increase in the number of people who are not working, whether on furlough or out of work. Within those who are working, there has been a shift from working at work to working from home. Before the pandemic, around 40 per cent of people in work (employed and self-employed) said that they ever worked from home; Of those who were previously employed but had never worked from home, more than one-quarter (28 per cent) are now working from home. This shift to home-

working during the crisis is likely to have long-term implications for working arrangements in the future.

As in Adams-Prassl et al (2020) we find that women are more likely than men to have stopped working during lockdown. Of those who report that they were working or self-employed prior to 20th March, 22 per cent are still working at their usual place of work (25 per cent of men and 18 per cent of women), 44 per cent are working from home (men = 43 per cent, women = 45 per cent), 26 per cent report that they have been furloughed (24 per cent of men and 28 per cent of women), while 8 per cent are not working (8 per cent of men and 9 per cent of women).

Table 2 reports estimated marginal effects from a multinomial logit regression on four possible employment outcomes (1 = still at work, 2 = working from home, 3 = on furlough and 4 = not employed). This is estimated on the sample of people who report that they are working (employed or self-employed) pre-COVID19.

Column 1 shows the raw differences (without controls). Women are nearly 7 percentage points less likely than men to still be at work. They are 4 percentage points more likely to be furloughed [$p=0.020$] and 2 percentage points more likely to be not working [$p=0.011$]. The estimated magnitudes change little when we add controls (in column 2), including a full set of occupational dummies (column 3). One possible hypothesis for the gender gap is that, where there was an element of choice, women may have been more likely to stop working in order to meet the increased demand for home childcare. However, results in Table 2, panel B show that the gender gap is also present among women without kids, indicating that other (non-child related) factors account for at least some of the (unexplained) gender gap. Even so, we cannot rule out that employment changes are endogenous with respect to childcare arrangements. Results in Table 2, panel B, column 3 include as an additional control the pre-COVID19 childcare gap (ie the within-household difference between the share of childcare done by women and the share of childcare done by men) in the regression. We interact the childcare gap with gender. The results show that the allocation of childcare (pre-COVID19) is correlated with employment outcomes (post-COVID19). The opposite signs for men and women imply a similar relationship between individuals' share of childcare and their post-COVID19 employment – men and women who did a smaller share of childcare are less likely to be on furlough and more likely to be working from home. If men working from home are observed to do less childcare post-COVID19 than e.g. men who are not working, at least part of this may be explained by less childcare, pre-COVID. This endogeneity motivates our within-household, difference-in-differences estimates of changes in the gender childcare gap in Section 4, which allows us to control for unobserved household-level heterogeneity.

Before moving on, it is interesting to look at the impact of other characteristics on employment outcomes. Our findings broadly mirror those found by Adams-Prassl et al (2020). Table 2, panel A,

shows that having a degree is associated with working from home (rather than at work) and being furloughed rather than being out of work. Those who were self-employed were more likely not to be working and less likely to be furloughed, reflecting the fact that UK Government furlough support for self-employed workers did not come into effect until 13th May, after our data collection. Not surprisingly, those who had previously worked from home, were more likely to report working from home and were less likely to be furloughed and not working. Pre-COVID19, 40 per cent of people in work reported that they occasionally/ sometimes/ always worked from home (41 per cent of men and 39 per cent of women). But there are also many who are working from home for the first time. Of those who reported that they previously never worked from home, 28 per cent are now doing so. This move to home-working may lead to longer-term changes in working arrangements.

4 Childcare changes since COVID19

The closure of UK schools and other childcare providers to all but the children of key workers and vulnerable children from Friday 20th March left millions of children requiring home childcare.

Prior to COVID19, the allocation of childcare within households was uneven, with women bearing the greater share. Focusing only on the childcare done within the household (ie ignoring external childcare), the average share of childcare done by women was 65.3 per cent. This represents a gender childcare gap of 30.4 percentage points (=65.3 – 34.7).

The magnitude of the estimated gap from our survey (pre-COVID19) is similar that that observed in the UK 2014-15 Time Use Survey (UKTUS). The UKTUS collects diary information at the household level on 10-minute intervals for a 24-hour period during weekdays and weekend days for a representative sample of individuals in the UK. We calculate an equivalent share of childcare done by women in this sample by dividing the daily minutes in childcare by the women over the daily minutes in childcare by the man in a given household. Daily minutes of childcare for every respondent by adding up the minutes reported in childcare as the primary activity in a given day. Childcare includes physical care and supervision, feeding, teaching, reading, talking, and accompanying the child to do activities among others. Women spend an average of 126 minutes per weekday and 113 minutes per weekend day, whereas men spend an average of 46 minutes during a weekday and 72 minutes during a weekend day. These figures result in a share of childcare by women (men) between 60 (40) per cent during weekdays and 73 (27) per cent during the weekend.

Post COVID19, there have been several dramatic changes to home childcare that we summarize below.

There has been a **sizeable increase in the total amount of home childcare** provided. Couples with young children (aged 12 or under) self-report doing an average of 40 (median)/ 49.7 (mean) hours of

additional childcare per week. In other words, families are taking on the equivalent burden of a working week in additional childcare. This figure is more than double the time spent on childcare prior to COVID19. Evidence from the 2014-15 UKTUS reveals weekly time spent in childcare by households to be around 20 hours per week.³ Note that these are self-reported additional hours of childcare – there may be a concern that they are over-estimates. However, previous studies have shown that estimates on housework from stylized questions (such as how much housework you do per week) are meaningfully associated with actual housework measures derived from diaries (Borra, Browning, and Sevilla, forthcoming; Hill, 1985; Robinson, 1985). Similarly, for childcare Del Bono et. al., (2016) validate frequency in childcare activities from the Millennium Cohort Survey and show that these measures are meaningfully associated with actual maternal time in childcare in the 2014-15 UKTUS survey.

Women are doing the majority of the additional home childcare. Table 3 summarizes additional hours of childcare post-COVID19 by gender. On average, women have been doing 30 (median)/ 30.3 (mean) additional hours' childcare per week, compared to 15 (median)/ 19.4 (mean) done by men. It is important to emphasize that these figures suggest a substantial increase in childcare (in absolute number of hours) done by men. 2014-2015 UKTUS estimates indicate that, in “normal” (pre-COVID19) times, women do an average of 15 hours per week and men do average of 6 hours per week. But the gender childcare gap is also large in absolute number of hours. Comparing median hours, the gender childcare gap equates to an additional ten hours done by the “typical” mother compared to the “typical” father each week.

The amount of additional childcare that is done by men and women is correlated with their post-COVID19 employment. This is shown in Figure 1, panel A, which plots average total additional hours for men and women according to their own employment. The figure indicates that the amount of additional **childcare is more sensitive to own-employment in the case of fathers than it is in the case of mothers.** Total hours of childcare per week vary by employment status for both men and women but to a greater extent in the case of men.

Although women are more likely not to work than men, **lower levels of employment do not account for all of women's higher number of childcare hours.** Figure 1, panel A, shows that mothers are doing more childcare than fathers, irrespective of their employment. Indeed, women who are at work/working from home are doing as many additional hours of childcare as men who are furloughed. One possibility is that partner's employment status (which is correlated with own-employment status) might also drive variation in the amount of childcare by own-employment status. However, Figure 1,

³ As is standard in the literature, we calculate the weekly time in housework by adding up daily childcare for the couple. To that end, we multiply daily weekday numbers by five and daily weekend numbers by two for each individual in a couple.

panel B, suggests that there is less sensitivity to partner's employment status than to own-employment status, particularly in the case of women.

To explore the relationship between childcare hours and own- and partner-employment more systematically, we run an OLS regression of total additional hours on indicators for own- and partner's employment. We also allow the correlation with own-employment to differ for men and women. The results are reported in Table 4. Comparing the results in columns (1) and (2) shows that there is a gender gap in additional childcare, conditional on own-and partner's employment. Including employment controls reduces the gender gap (from 11.8 hours to 7.2 hours) but it remains statistically significant. The coefficients on partner's employment in column (2) are economically small and statistically significant, with the exception of not employed, confirming (similar to previous studies, eg Gimenez et al, 2009) that the number of hours of additional childcare is less sensitive to partner's employment than to own-employment. The results in column (3), including interactions between own-employment and gender, confirm that the number of additional hours' childcare is less sensitive to women's own-employment than it is to men's.

Although women are bearing more of the burden, **the average within-household gender childcare gap (i.e. the difference between the share done by women and men) is smaller**. Compared to an initial average within-household childcare gap of 30.4 percentage points, the post-COVID19 gap in additional childcare hours is 27.2 percentage points (see Table 3). In other words, the allocation of the additional burden of childcare is more equal than the pre-COVID19 allocation. However, the extent to which the gap is smaller depends heavily on men's employment.

We perform a "difference-in-differences" analysis of the effect of COVID19 on the within household childcare allocation. Specifically, we estimate the effect of COVID19 on the difference in the within-household (female/ male) shares before/ after lockdown. Our outcome variable (Gap_change_i) is the change before/ after COVID19 in the gender childcare gap for household i where the gap measures the percentage point difference in the share of childcare done by women and the share of childcare done by men, within the same household, i.e:

$$Gap_change_{i1} = Gap_{i1} - Gap_{i0}$$

$$\text{where } Gap_{it} = Share_{f_{it}} - Share_{m_{it}}, t = 0, 1$$

We calculate this Gap_change measure directly using our survey data, exploiting the fact that we observe the childcare allocation pre-COVID19 and the allocation of additional hours post-COVID19 for the same household. This allows us to control for unobserved heterogeneity that might be correlated with post-COVID19 employment outcomes when we look at the relationship between childcare allocation and employment. Averaging over all families, the gap has narrowed by 3.3

percentage points [$p=0.278$]. However, further analysis, presented in Table 5, shows that this average change masks considerable variation by (post-COVID19) employment status.

Table 5 reports results from regressing the household-level `gap_change` on indicators for men's and women's employment status and additionally (in column 3) the pre-COVID19 household childcare gap. The results in column 3 show that there has been a 12 percentage point narrowing of the childcare gap when men work from home [$p=0.089$]. There has been an even greater change when men have been furloughed or lost their jobs. The magnitudes of the changes in the gap in these cases when men are furloughed/ not employed are big enough to close the gap and move men to doing a majority share of the additional childcare. On the flipside, when women have lost their jobs, they have taken on an even greater share of the additional childcare than pre-COVID19 in these cases, the childcare gender gap has widened, moving women close to a 100 per cent share of the additional childcare.

5 Discussion

For families with young children, the months of lockdown have meant providing many additional hours of childcare – equivalent to a full-time, working week. In many cases, these hours have had to be provided in addition to working at work or from home.

Women have done more of this childcare than men (roughly ten hours a week more). Partly, this is because they are less likely to be working but that does not account for all the difference. The quantity of childcare done by women is much less sensitive to their employment than it is for men and, indeed, women have done more childcare than men, irrespective of their employment status; women working from home have done more childcare than men on furlough/ or who have lost their job.

The burden of additional childcare may have damaging long-term consequences for the career prospects of parents with young children – and particularly for women. Coviello et al., (2015) show evidence from the judiciary documenting how judges who juggle more trials at once instead of working sequentially on few of them at each unit of time take longer in closing a case. When working from home during the lockdown it is hard to be as productive as someone without children if you are juggling work with near full-time childcare. In fact, evidence from on-line job markets shows that women earn 20% less per hour on average, which can partly be explained by women, women with young children, having more fragmented work patterns which affects their ability to complete a task (Adams, 2020). Similarly, in academia there is anecdotal and some statistical evidence that the share of working papers being published and submissions to journals by women has fallen post-COVID19 (Shurchkov, Olga. 2020). Employers need to recognise – and perhaps take measures to compensate parents for – the lockdown childcare burden.

There has been a shift from working at work to working from home. Of those who had never worked from home pre-COVID19, 28 per cent are now working from home. This may bring about a permanent positive change in working arrangements.

There have also been baby steps towards a more equal allocation. For many families, the allocation of the additional hours of childcare is more equal than the previous allocation of childcare. The gender childcare gap (the gap between the share done by women and the share done by men) has narrowed (from 30.5 percentage points to 27.2 percentage points). However, this has not happened uniformly across all families, however, but has been driven to some extent where men are working from home and, to a far greater extent, where men are on furlough/ have lost their jobs. In that respect, the effect of lockdown is similar to that of other childcare supply shocks that occurred during previous recessions but on a grander scale because of the furlough scheme (Aguiar and Hurst, 2013 and Gimenez et al, 2009). It remains to be seen whether the change is a permanent one. Some evidence from paternity leave policies suggest that temporary changes can have longer-term effects on social norms, evidenced by increases in the time that fathers spend in household activities, including childcare (Ferre and Gonzalez 2019 and Patnaik, 2019). Two things are different about COVID19 lockdown. The first is the scale of the demand-side shock. The changes have been profound. The total amount of childcare being done at home completely dwarves usual amounts because of the closure of almost all formal childcare. The impact has also been across the board, affecting all families, meaning that almost all men have increased the quantity of childcare that they do. But the second difference is that this is not a deliberate policy to promote a more equal distribution of childcare, but an unintended consequence of measures to stop a virus spreading. The changes that have been brought about may need to be recognised and reinforced to have longer term effects.

References

- Adams, A. (2020) The Gender Wage Gap on an Online Labour Market: The Cost of Interruptions, CEPR Working Paper DP14294
- Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis (2013) "Time Use during the Great Recession." *American Economic Review*, 103 (5): 1664-96.
- Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey and Michèle Tertilt (2020) "The impact of Covid-19 on gender equality", *CEPR COVID Economics and Real Time Papers*, 4, 62-86.
- Decio Coviello, Andrea Ichino, Nicola Persico, (2015) "The Inefficiency of Worker Time Use", *Journal of the European Economic Association*, 13 (5) 906–947.
- Del Bono, Emilia, Marco Francesconi, Yvonne Kelly and Amanda Sacker (2016) "Early Maternal Time Investment and Early Child Outcomes", *Economic Journal*, 126 , F96–F135.
- Ekberg, Eriksson, and Friebel (2013) "Parental leave — A policy evaluation of the Swedish "Daddy-Month" reform", *Journal of Public Economics*, 97, 131-143.
- Farré, Lidia, and Libertad González. 2020. "¿Quién Se Encarga de Las Tareas Domésticas Durante El Confinamiento? Covid-19, Mercado de Trabajo Y Uso Del Tiempo En El Hogar." Nada Es Gratis. April 23, 2020. <https://nadaesgratis.es/admin/quien-se-encarga-de-las-tareas-domesticas>.
- Farré, Lidia, and Libertad González (2019) Does paternity leave reduce fertility?, *Journal of Public Economics*, 122, 52-66.
- Gimenez-Nadal, J. I., and Molina, J. A. (2014) "Regional unemployment, gender, and time allocation of the unemployed" *Review of Economics of the Household*, 12(1), 105–127.
- Hapucheck and Petrongolo (2020), "Work, care and gender during the Covid-19 crisis", CEP Center For Economic Performance COVID-19 Analysis, No. 002.
- Hill, M. 1985. "Patterns of Time Use," In: Juster, T., Stafford, J., (Eds.) *Time, goods, and well-being*. Ann Arbor, MI: Institute for Social Research, The University of Michigan.
- Patnaik, Ankita (2019) "Reserving Time for Daddy: The Consequences of Fathers' Quotas", *Journal of Labor Economics*, 37:4, 1009-1059.
- Robinson, J.P., 1985. "The validity and reliability of diaries versus alternative time use measures." In: Juster, T., Stafford, J., (Eds.) *Time, goods, and well-being*. Ann Arbor, MI: Institute for Social Research, The University of Michigan.
- Shurchkov, Olga. 2020. "Is COVID-19 Turning Back the Clock on Gender Equality in Academia?" Medium. April 23, 2020. <https://medium.com/@olga.shurchkov/is-covid-19-turning-back-the-clock-on-gender-equality-in-academia-70c00d6b8ba1>.

Table 1 Summary statistics

	Full sample		Couples with young children (aged <=12)	
	Men	Women	Men	Women
Age	41.6	37.5	37.0	36.3
Degree (0/1)	0.446	0.478	0.529	0.490
Pre COVID				
Working	0.639	0.592	0.806	0.502
Of whom, WFH	0.411	0.391	0.431	0.429
Self-employed	0.077	0.066	0.074	0.056
Not working	0.232	0.272	0.090	0.385
Other	0.053	0.080	0.031	0.056
Post COVID				
At work	0.197	0.135	0.265	0.115
WFH	0.330	0.315	0.441	0.269
Furloughed	0.186	0.197	0.155	0.141
Not employed	0.287	0.353	0.140	0.457
N	1983	2267	207	263

Notes to table: *Degree* includes NVQ4 / HNC / HND / Bachelor's degree or similar/ NVQ5 or post-graduate diploma. Current and retrospective employment status collected post-COVID19. For further information on the questions asked, see Section 2.

Table 2: Estimated marginal effects (multinomial logit), post-COVID employment status

A Sample of individuals working pre-COVID

	(1)		(2)		(3)	
	Marginal effect	p-value	Marginal effect	p-value	Marginal effect	p-value
Female						
At work	-0.066	0.000	-0.059	0.000	-0.052	0.001
Working from home	0.004	0.838	0.003	0.879	-0.001	0.951
Furloughed	0.039	0.020	0.028	0.086	0.032	0.058
Not employed	0.023	0.011	0.028	0.006	0.021	0.039
Kids (0/1)						
At work			0.026	0.110	0.023	0.143
Working from home			0.023	0.216	0.025	0.162
Furloughed			-0.021	0.250	-0.021	0.216
Not employed			-0.029	0.013	-0.026	0.020
Degree (0/1)						
At work			-0.104	0.000	-0.076	0.000
Working from home			0.217	0.000	0.128	0.000
Furloughed			-0.103	0.000	-0.051	0.003
Not employed			-0.009	0.313	0.000	0.982
Self-employed (0/1)						
At work			0.003	0.917	-0.008	0.767
Working from home			-0.049	0.151	-0.018	0.576
Furloughed			-0.121	0.001	-0.133	0.000
Not employed			0.167	0.000	0.161	0.000
Pre-WFH (0/1)						
At work			-0.107	0.000	-0.075	0.000
Working from home			0.241	0.000	0.189	0.000
Furloughed			-0.108	0.000	-0.081	0.000
Not employed			-0.025	0.016	-0.031	0.002
Age = 30s						
At work			0.010	0.618	0.020	0.308
Working from home			0.033	0.149	0.010	0.654
Furloughed			-0.028	0.212	-0.029	0.184
Not employed			-0.015	0.230	0.000	0.948
Age = 40s						
At work			0.030	0.163	0.033	0.116
Working from home			0.033	0.185	0.015	0.519
Furloughed			-0.043	0.062	-0.041	0.081
Not employed			-0.019	0.160	-0.008	0.564
Age = 50s						
At work			0.045	0.037	0.046	0.030
Working from home			-0.003	0.898	-0.015	0.540
Furloughed			-0.065	0.005	-0.066	0.004
Not employed			0.022	0.127	0.035	0.020
Occupation						
	No		No		Yes	
N	2,782		2,782		2,782	

Notes to table: Marginal effects estimated at mean values of co-variables. The sample includes only those respondents who reported that they were working (employed/ self-employed) pre-COVID19. *Degree* includes NVQ4 / HNC / HND / Bachelor's degree or similar/ NVQ5 or post-graduate diploma. Self-employed and pre-WFH (= ever worked from home) refer to pre-COVID19 status

Covid Economics 23, 28 May 2020: 58-78

B With/ without kids

	No kids		Kids <=12		Kids <=12	
	Marginal effect	p-value	Marginal effect	p-value	Marginal effect	p-value
Female						
At work	-0.069	0.000	-0.101	0.032	-0.043	0.693
Working from home	0.023	0.283	-0.010	0.846	-0.003	0.966
Furloughed	0.025	0.219	0.055	0.195	0.025	0.669
Not employed	0.021	0.113	0.056	0.044	0.021	0.648
Childcare gap (pre)						
At work					-0.059	0.390
Working from home					0.224	0.004
Furloughed					-0.130	0.080
Not employed					-0.034	0.561
Gap x Female						
At work					-0.043	0.693
Working from home					-0.235	0.052
Furloughed					0.168	0.116
Not employed					0.110	0.174
Other controls			Degree			
			Self-employed			
			WFH (pre)			
			Ageband			
N	1,845		312		312	

Notes to table: Marginal effects estimated at mean values of co-variables. The sample includes only those respondents who reported that they were working (employed/ self-employed) pre-COVID19. *Childcare gap (pre)* refers to the within-household difference between the share of childcare done by the woman and the share of childcare done by the man, prior to COVID-19 (retrospectively reported). A positive “gap” indicates that the woman does a greater share than the man.

Table 3: Allocation of childcare between women and men

	Women		Men	
Post-COVID19:				
Additional hours per person	Median	Mean	Median	Mean
Total per week (seven days)	30	30.3	15	19.4
Total, weekdays (five days)	28	25.4	10	15.2
Average per day (seven days)	4.3	4.5	2.4	2.9
Average per weekday (five days)	6	5.2	2.4	3.2
Within-household share of childcare:				
Mean, pre-COVID19	65.3%		34.7%	
Mean, post-COVID19	63.6%		36.4%	
N	290		290	

Notes to table: Additional hours of childcare refer to the (self-reported) additional hours done by men and women each day (compared to pre-COVID19). Respondents are asked to report the hours done by themselves and their partners each day.

These are aggregated to produce weekly totals. The shares are based on self-reported shares for respondents and their partners (pre-COVID) and self-reported total hours for respondents and their partners (post-COVID). For further information on questions asked, see Section 2.

Table 4: OLS regression results.**Outcome = total hours additional childcare per week (post-COVID19)**

	(1)	(2)	(3)
Female	11.834 (1.649)	7.248 (1.732)	14.074 (4.350)
WFH		1.337 (2.373)	3.247 (2.867)
Furloughed		8.539 (2.797)	13.038 (3.663)
NotEmployed		13.471 (2.663)	15.868 (3.882)
Partner_WFH		-1.595 (2.373)	-1.130 (2.383)
Partner_Furloughed		-1.355 (2.797)	-1.198 (2.796)
Partner_NotWorking		-8.705 (2.663)	-8.419 (2.665)
Female_WFH			-6.836 (4.953)
Female_Furloughed			-11.568 (5.802)
Female_Not employed			-7.408 (5.638)
Constant	18.938 (1.166)	19.003 (2.663)	16.756 (2.952)
N	580	580	580

Notes to table: Table reported estimated coefficients and standard errors. Additional hours of childcare refer to the (self-reported) additional hours done by men and women each day (compared to pre-COVID19). Respondents are asked to report the hours done by themselves and their partners each day. These are aggregated to produce weekly totals. For further information on questions asked, see Section 2.

Table 5: OLS regression results**Outcome = within household change in the gender childcare gap**

	(1)	(2)	(3)
Constant	-0.033 (0.030)	-0.023 (0.103)	0.083 (0.098)
Man_WFH		-0.076 (0.078)	-0.126 (0.074)
Men_Furloughed		-0.253 (0.097)	-0.317 (0.091)
Man_NotWorking		-0.221 (0.103)	-0.372 (0.099)
Woman_WFH		0.014 (0.111)	0.043 (0.103)
Women_Furloughed		0.255 (0.120)	0.290 (0.112)
Woman_NotWorking		0.136 (0.110)	0.338 (0.108)
Pre_gap			-0.476 (0.074)
N	290	290	290

Notes to table: The gap change refers to the within-household change before/after COVID19 in the gender childcare gap, where the gender childcare gap is defined as the difference between the share of childcare done by women and the share of childcare done by men. Pre-COVID refers to all childcare; Post-COVID19 refers to the additional hours of childcare. The Pre-COVID gender gap is 30.4. A negative number corresponds to a narrowing of the gender childcare gap. For further information on questions asked, see Section 2.


Figure 1: Additional hours' childcare (total per week), by post-COVID19 employment status



Notes to table: The figure shows average self-reported total hours additional childcare done by men and women post-COVID19. For further information on questions asked, see Section 2.

Covid Economics 23, 28 May 2020: 58-78

Figure 2: Changes in the within-household gender childcare gap



Notes to figure: The gap change refers to the within-household change before/after COVID19 in the gender childcare gap, where the gender childcare gap is defined as the difference between the share of childcare done by women and the share of childcare done by men. Pre-COVID refers to all childcare; Post-COVID19 refers to the additional hours of childcare. The Pre-COVID gender gap is 30.4. A negative number corresponds to a narrowing of the gender childcare gap. For further information on questions asked, see Section 2.

Production networks and epidemic spreading: How to restart the UK economy?¹

Anton Pichler,² Marco Pangallo,³ R. Maria del Rio-Chanona,⁴ François Lafond⁵ and J. Doyne Farmer⁶

Date submitted: 21 May 2020; Date accepted: 22 May 2020

We analyse the economics and epidemiology of different scenarios for a phased restart of the UK economy. Our economic model is designed to address the unique features of the COVID-19 pandemic. Social distancing measures affect both supply and demand, and input-output constraints play a key role in restricting economic output. Standard models for production functions are not adequate to model the short-term effects of lockdown. A survey of industry analysts conducted by IHS Markit allows us to evaluate which inputs for each industry are absolutely necessary for production over a two month period. Our model also includes inventory dynamics and feedback between unemployment and consumption. We demonstrate that economic outcomes are very sensitive to the choice of production function, show how supply constraints cause strong network effects, and find some counter-intuitive effects, such as that reopening only a few industries can actually lower aggregate output.

- 1 We would like to thank David Van Dijke, David Vines, Eric Beinhocker, Spencer Fox and John Muellbauer for many useful comments and discussions. We thank Baillie Gifford, IARPA, the Oxford Martin School and JSMF for the funding that made this possible. We appreciate that IHS Markit provided us with a survey on critical vs. non-critical inputs. (Note that JDF is on their advisory board).
- 2 Institute for New Economic Thinking at the Oxford Martin School and Mathematical Institute, University of Oxford; Complexity Science Hub Vienna, Austria.
- 3 Institute of Economics and EMbeDS Department, Sant'Anna School of Advanced Studies.
- 4 Institute for New Economic Thinking at the Oxford Martin School and Mathematical Institute, University of Oxford.
- 5 Institute for New Economic Thinking at the Oxford Martin School and Mathematical Institute, University of Oxford.
- 6 Institute for New Economic Thinking at the Oxford Martin School and Mathematical Institute, University of Oxford; Santa Fe Institute.

Copyright: Anton Pichler, Marco Pangallo, R. Maria del Rio-Chanona, François Lafond and J. Doyne Farmer

Occupation-specific data and contact surveys allow us to estimate how different industries affect the transmission rate of the disease. We investigate six different re-opening scenarios, presenting our best estimates for the increase in RO and the increase in GDP. Our results suggest that there is a reasonable compromise that yields a relatively small increase in RO and delivers a substantial boost in economic output. This corresponds to a situation in which all non-consumer facing industries reopen, schools are open only for workers who need childcare, and everyone who can work from home continues to work from home.

1 Introduction

The social distancing measures imposed to combat the COVID-19 pandemic have created severe disruptions to economic output. Governments throughout the world are contemplating or implementing measures to ease social distancing and reopen the economy, which may involve a tradeoff between increasing economic output vs. increasing the expected number of deaths due to the pandemic. Here we investigate several scenarios for the phased reopening of the economy. At one extreme, we find that reopening only a very limited number of industries can create supply chain mis-coordination problems that in some cases might actually decrease aggregate output. In contrast, reopening all industries would most likely increase R_0 above 1. We find a good scenario in-between these extremes: reopening a large part of the upstream industries, while consumer-facing industries stay closed, limits supply chain mis-coordination while providing a large boost to output and a relatively small increase in infection rates.

The shocks to the economy caused by social distancing are highly industry specific. Some industries are nearly entirely shut down by lack of demand, others are restricted by lack of labor, and many are largely unaffected. Feedback effects amplify the initial shocks. The lack of demand for final goods such as restaurants or transportation propagates upstream, reducing demand for the intermediate goods that supply these industries. Supply constraints due to lack of labor under social distancing propagate downstream, by creating input scarcity that can limit production even in cases where the availability of labor and demand would not have been an issue. The resulting supply and demand constraints interact to create bottlenecks in production. The resulting decreases in production may lead to unemployment, decreasing consumption and causing additional amplification of shocks that further decrease final demand.

Understanding these effects requires a model at the level of individual industries. Most of the economic analysis of the COVID-19 pandemic uses relatively aggregate macro models (Eichenbaum et al. 2020, Bodenstein et al. 2020), with only a few studies predicting the economic effects using input-output (IO) models. IO models are particularly relevant to evaluate the consequences of crises such as COVID-19, where different sectors are affected differently, and the propagation of shocks through supply chains is likely to amplify the initial effects. Table 1 summarizes the main features of several IO models that have been put forward recently to evaluate the macroeconomic effects of the COVID-19 crisis. Our paper differs in a number of important ways from the literature. On one hand, we provide comprehensive scenarios, an estimation of the epidemic spreading, non-equilibrium dynamics, and explicit demand shocks together with a sophisticated consumption response. On the other hand, we do not model prices, as we argue that price changes during the lockdown are relatively small.

The most important conceptual difference that distinguishes our model is our treatment of the production function, which dictates most of the behavior of the models listed in Table 1. Essentially, the literature can be ordered by the degree to which the production function allows substitutions between inputs. At one extreme, the Leontief production function assumes a fixed recipe for production, allowing no substitutions and restricting production based on the limiting input (Inoue & Todo 2020). Under the Leontief production function, if a single input is severely reduced, overall production will be reduced proportionately, even if that input is ordinarily relatively small. This can lead to unrealistic behaviours. For example, the steel industry has restaurants as an input, presumably because steel companies have a workplace canteen and sometimes entertain their clients and employees. A literal application of the Leontief production function would predict that a sharp drop in the output of the restaurant industry will dramatically reduce steel output. This is unrealistic, particularly in the short run.

Feature	Inoue & Todo (2020)	Barrot et al. (2020)	Mandel & Veetil (2020)	Fadinger & Schymik (2020)	Bonadio et al. (2020)	Baqae & Farhi (2020)	This paper
Date	31 Mar	2 April	29 Apr	April	April	7 May	May 26, 2020
Geographical scope	Japan	France + some EU countries	World	Germany	64 countries	US	UK
Data/ aggreg. level	1.67M firms	INSEE 38 + WIOD 54 sectors	WIOD 55 sectors	Eurostat 62 sectors	OECD 33 sectors	BEA 66 sectors	WIOD 55 sectors
Production function	Leontief	CES	CES	CD	CD nesting CES	nested CES	Leontief for "critical inputs" + Lack of non-critical inputs does not harm output
Utility and/or consumption function	Fixed demand	CES	CES	CD	2 nests CES	CRS non parametric	Linear
Supply shocks	Essential activities	WFH + Essential industries + School closures	Plausibility considerations from Dorn et al. (2020)	WFH	WFH + lockdown stringency + Essential for Health only	WFH + some share "not removed"	WFH + Essential industries
Final demand shocks	No	No	No	No	No	Theory only	Yes
Equilibrium	No	Yes	Transient	Yes	Yes	Yes	No
Prices	No	Yes	Yes, sticky	Yes	Yes	Yes	No
Inventories	Inputs inventories	No	Output inventories	No	No	No	Inputs inventories
Fiscal/ monetary policy	No	No	No	No	No	No	Yes
Reopening policy	Yes	Individual industries	Yes (lockdown lifted)	Individual industries	Yes, symmetric lockdown to	No	Yes, scenarios
Epidemic spreading	No	No	No	No	No	No	Yes
Other			Frictional input shares adjustment; hysteresis in final demand		Counterfactual for re-nationalizing supply chains	Productivity shocks (theory)	V-shape or instantaneous recovery consumption function

Table 1: Summary of features in IO papers modelling the effect of COVID-19 on the economy. WFH: Work From Home. CD: Cobb Douglas. CES: Constant Elasticity of Substitution.

The alternatives used in the literature are the Cobb-Douglas production function (Fadinger & Schymik 2020), which has an elasticity of substitution of 1, and the CES production function, where typically calibration for short term analysis uses an elasticity of substitution less than 1 (Barrot et al. 2020, Mandel & Veetil 2020, Bonadio et al. 2020). Some papers (Baqae & Farhi 2020) consider a nested CES production function, which can accommodate a wide range of technologies. In principle, nests could allow for substitution between some inputs and forbid it between others, in different ways for different industries. However, it is hard to calibrate all these elasticities, so that in practice many models end up using very limited nesting structure or assuming uniform substitutability. Consider again our example of the steel industry. With common calibrations of the (nested) CES production function, firms could substitute iron for energy, while still producing the same output. To the extent that certain production processes are encoded in fixed technological "recipes", this is clearly unrealistic¹.

We argue that modeling production during the COVID-19 crisis requires a new approach to production functions, that is different from both standard Leontief and CES production functions. In this paper, we mostly keep the basic Leontief assumption that firms cannot substitute one input for another. However, we depart from the Leontief assumption in that we

¹Papers using CES production functions also assume equilibrium mechanisms for substitution and prices that are not completely realistic for the short run dynamics in the current context. Consider again our restaurant example. In an equilibrium model with a CES production function, if the output of restaurants is constrained, the relative price of restaurants will immediately increase, and firms will choose a relatively lower level of restaurant inputs, and a relatively higher level of other inputs. In the current context the closure of work canteens in steel factories is not driven by increased prices but policy and it is not clear that it is affecting steel output, or that it is substituted by other inputs.

allow firms to keep producing as long as they have the inputs that are absolutely necessary, which we call “critical inputs”. The steel industry cannot produce steel without iron and energy, but it can operate for a considerable period of time without restaurants or logistics consultants. Specifically, we make the assumption that if restaurants cannot supply the steel industry, the steel industry simply keeps producing at the same rate. This is of course only an approximation. To keep the same example, by not using restaurants, the costs of the steel industry are reduced and, *ceteris paribus*, its profits increase. In reality, non-critical inputs may have an impact on steel output that could be modeled as a shock to productivity. However, we think that during the short time-scales of the pandemic, these problems are second-order effects, and our production function provides a better assumption than Leontief or CES production functions.

In order to determine which inputs are critical and which are not, we use a survey that IHS Markit performed at our request. In this survey they asked “Can production continue in industry X if input Y is not available for two months?”. The list of possible industries X and Y was drawn from the 55 industries in the World Input-Output Database. This question was presented to 30 different industry analysts who were experts in industry X. Each of them was asked to rate the importance of each of its inputs Y. They assigned a score of 1 if they believed input Y is critical, 0 if it is not critical, and 0.5 if it is in-between, with the possibility of a rating of NA if they could not make a judgement. We then apply the Leontief function to the list of critical inputs, ignoring non-critical inputs. We experimented with several possible treatments for industries with ratings of 0.5 and found that we get somewhat better empirical results by treating them as non-critical (though at present we do not have sufficient evidence to resolve this question unambiguously).

Besides the bespoke production function discussed above, we also introduce a COVID-19-specific treatment of consumption. Most models do not incorporate the demand shocks that are caused by changes in consumer preferences in order to minimize risk of infection. The vast majority of the literature has focused on the ability to work from home, and some studies incorporate lists of essential vs. inessential industries, but almost no papers have also explicitly added shocks to consumer preferences. (Baqae & Farhi (2020) is an exception, but the treatment is only theoretical). Here we use the estimates from del Rio-Chanona et al. (2020), which are taken from a prospective study by the Congressional Budget Office (2006). These estimates are crude, but we are not aware of estimates that are any better. As we write, data on actual consumption starts to become available; what we have seen so far is qualitatively consistent with the shocks predicted by the CBO, with massive shocks to the hospitality industry, travel and recreation, milder (but large) shocks elsewhere, and increased expenses in groceries/food retail (Andersen et al. 2020, Carvalho et al. 2020, Chen et al. 2020). Besides the initial shock, we also attempt to introduce realistic dynamics for recovery and for savings. The shocks to on-site consumption industries are more long lasting, and savings from the lack of consumption of specific goods and services during lockdown are only partially reallocated to other expenses.

The first step in the analysis of our model is empirical validation: We compare model predictions to the economic data that statistical agencies have started to disseminate. To compare to UK data, we start the lockdown in our model on March 23rd, and keep it for two months. For 2020Q1, we find a 1.7% reduction in GDP compared to 2019Q4, which is close to the 2% early estimate recently released by the Office for National Statistics. For 2020Q2, we forecast that GDP would be 21.5% lower than in 2019Q4, which is in the range of forecasts provided by economic institutions and consulting firms. We also compare model predictions to sectoral unemployment data, finding good agreement.

As a second step, we investigate some theoretical properties of the model. Our analysis makes it clear that bottlenecks in supply chains can strongly suppress aggregate economic

output. The extent to which this is true depends on the production function. These effects are extremely strong with the Leontief production function, are much weaker with a linear production function (which allows unrealistically strong substitutions) and have an intermediate effect with our modified Leontief function. Network effects can strongly inhibit recovery, and can cause counter-intuitive results, such as situations in which reopening a few industries can actually depress economic output.

Our third step, which is the key aim of this paper, is to find a good compromise between the economic benefit of reopening industries and negative health consequences of doing so. It is worth keeping in mind that many health outcomes depend on the state of the economy, so that keeping the economy closed also has negative health consequences.

The fundamental principles of epidemic spreading are relatively well understood, and it is clear that social distancing measures reduce the spreading of COVID-19 (Jarvis et al. 2020, Maier & Brockmann 2020, Arenas et al. 2020). The difficulty comes with obtaining good estimates of the key parameters that govern the fate of an epidemic, and in particular, the reproduction number R_0 , which gives the number of secondary cases for each primary case in a largely unaffected population. If R_0 is above one, the disease spreads to a given percentage of the population, otherwise the epidemic dies out. In this paper, we side-track the problem of developing a full-fledged epidemic spreading model, and focus on estimating R_0 . We decompose the reproduction number into the infections caused by contacts during work, during consumption, during public transport, and in other contexts, i.e. home and other social interactions. We use recent contact survey data from Sweden to estimate the share of infection due to each type of contact. For each industry we estimate its relative contribution to overall work and consumption infections. For instance, the Health sector is responsible for more work-related infections than the Forestry sector. This is because workers have more contacts, contacts are more risky, and there are more Health workers than Forestry workers. As another example, the Retail and Restaurant sectors contribute much more to consumption infections than the Mining sector, because there are virtually no direct consumption activities in the Mining industry.

We estimate the epidemiological consequences of scenarios for coming out of lockdown. Lifting the lockdown for a specific industry has several effects: workers of this industry contribute to increased work-related infection; consumers of this industry (if any) contribute to increased consumption-related infections; workers of this industry contribute to increased public transport infections; and finally, children of these workers go back to school if the workers cannot work from home, contributing to increased school-related infections. We assume in all the scenarios that workers who can work from home continue working remotely.

We present a summary of our re-opening scenarios results in the next section. We then present in detail our economic model and its calibration in Section 3. We show our model predictions for the UK economy in Section 4 and discuss production network effects and re-opening single industries in Section 5. We introduce the epidemic model and present effects of re-starting the economy on infectious contagion in Section 6. We conclude in Section 7.

2 A sweet spot for partially reopening the economy with only a minimal boost to the epidemic

Fig. 1 summarizes our bottom line results, presenting the trade-off between increasing economic production and mitigating the spread of the pandemic under five different scenarios. The bars on the left show estimates of R_0 and the bars on the right show GDP as a percentage of the pre-lockdown GDP. For comparison, pre-lockdown is shown on the left. The scenarios are: keeping lockdown; opening Manufacturing and Construction (which is short for fully re-opening

Agriculture, Mining, Manufacturing, Utilities and Construction); opening all industries except consumer-facing industries; same, but additionally opening schools; and opening all industries.

We find that a two-month lockdown has a strong impact on the economy, with gross output, value added, and consumption decreased by 26%, 25%, and 18%, respectively, compared to the UK pre-lockdown levels. Compared to the economy’s performance before the lockdown, an additional month of lockdown would decrease GDP from 75% to 74%, while re-opening only the Manufacturing and Construction sectors would increase GDP from 75% to 76% in a month. Re-opening all sectors except those that are consumer-facing would increase GDP to 82% of its pre-lockdown value, i.e. in a month it would increase GDP by 8 percentage points with respect to the lockdown scenario. Opening all industries only adds an additional 2% boost to GDP. Note that the scenario with all industries open has only 84% of pre-lockdown GDP. This is due to a combination of a persistent depression in demand for industries like restaurants (even if they are open) and the fact that consumer expectations take time to recover.




Figure 1: **How different policy scenarios affect R_0 and economic output.** The coloured bars show the expected reproduction number of the epidemic for each policy scenario. Different colours designate the activities that cause the epidemic to spread. The purple bars denote the percentage increase in value added relative to lockdown a month after the economy is opened under each scenario. Black lines are two standard deviation error bars. Note that we have normalized the infection rates for all scenarios so that they correspond to the *Jarvis et al. (2020)* study during lockdown. (Our estimate during lockdown is roughly $R_0 \approx 0.90$; their estimate is 0.62; these agree with the error bars). Note that Manufacturing and Construction also includes mining, agriculture, and a few others.

A comparison to our predictions for the increase in R_0 under each scenario shows that for the scenario where all industries except consumer facing industries are opened, the increase in R_0 relative to lockdown is small. In contrast, as soon as schools are open R_0 rises dramatically, and is very likely greater than one. If the economy is fully reopened, the predicted rise in R_0 is very likely substantially greater than one. Note that when the economy is fully reopened we find an R_0 still disturbingly greater than one, although much lower than the pre-lockdown value, as we assume that work from home continues and non-work related social distancing measures continue. Another reason for this is that we renormalize all our epidemiological results by the factor of 0.62/0.90, corresponding to the ratio of our original estimate to that of *Jarvis et al. (2020)* for the lockdown situation, thus ensuring that our estimate for the lockdown scenario corresponds to theirs. We did this because we feel that the relative values of our estimated R_0 across different scenarios are more reliable than the absolute values, and we defer to professional epidemiologists for estimating the absolute values. It is important to bear in mind that all these values are uncertain, and the uncertainties potentially make the difference in determining

Covid Economics 23, 28 May 2020: 79-151

whether there will be a second wave of the epidemic in the UK. We should also stress the uncertainties in the economic results – as we will show here, they depend rather sensitively on assumptions about the production function.

Thus, our results suggest that there is a “sweet spot”, corresponding to the scenario in which all except consumer facing industries reopen, with schools remaining closed for the children of parents who do not work or can work from home.² This scenario provides a good combination of a minimal predicted increase in R_0 and a substantial economic boost over remaining in lockdown.

The official UK government guidelines for COVID-19 recovery “Step one” (until June 1st)³ recommend that, in addition to sectors that were previously considered essential, manufacturing and construction should reopen, but that consumer-facing industries such as hospitality and non-essential retail should remain closed. Overall, this scenario corresponds to something in between our second and third scenario, depending on whether sectors such as business services fully reopen. By contrast, other countries (e.g. France) reopened personal services and non-essential retail soon after lockdown was lifted, which would correspond to something between our fourth and fifth scenario (depending on whether schools are open).

3 Economic Model

To analyse the economic benefits of staged re-opening we introduce a sectoral macroeconomic model that was inspired by the work of Battiston *et al.* (2007), Hallegatte (2008), Henriet *et al.* (2012) and Inoue & Todo (2019). We combine elements of these models and extend them to include new features. Our model incorporates production network effects that can amplify economic shocks both upstream and downstream.

In our model producers experience supply shocks caused by a nationwide lockdown. In the lockdown workers in non-essential industries who are unable to work from home become unproductive, resulting in lowered productive capacities of industries. At the same time demand-side shocks hit as consumers adjust their consumption preferences to avoid getting infected. We use the first-order supply and demand shocks predicted by del Rio-Chanona *et al.* (2020) to initialise our macro model.

Our model is open-source and can be downloaded together with all relevant data⁴. We also provide an interactive online interface for our model, allowing the user to explore alternative scenarios and parameter ranges⁵.

3.1 Timeline

A time step t in our economy corresponds to one day. There are N industries⁶, one representative firm for each industry, and one representative household that owns the industries. Every day:

1. Firms hire or fire workers depending on whether their workforce was insufficient or redundant to carry out production in the previous day.
2. The representative household decides its consumption demand and industries place orders for intermediate goods.

²This understandably creates childcare problems and will require creative solutions, and our model is not designed to account for the negative effects on health and future human capital that closed school can create.

³<https://www.gov.uk/government/publications/our-plan-to-rebuild-the-uk-governments-covid-19-recovery-strategy> (accessed: 2020/05/21).

⁴<https://doi.org/10.5281/zenodo.3834116>.

⁵https://anton-pichler.shinyapps.io/UK_COVID-19_economic_simulator/

⁶See Appendix G, Tables 12-13 for a comprehensive summary of notations used.

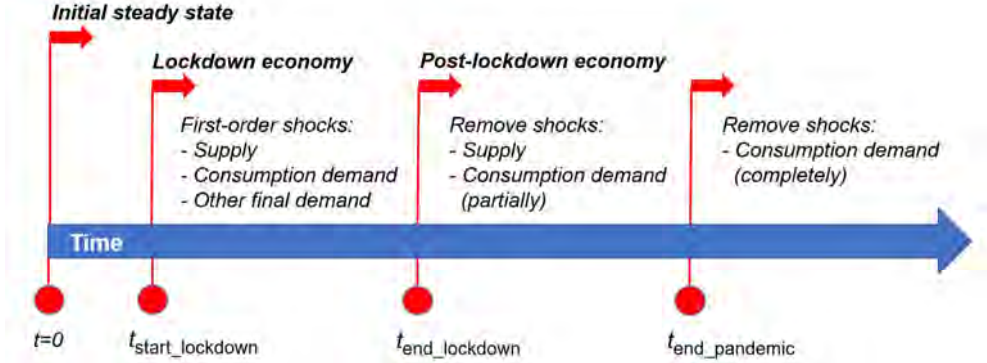


Figure 2: Schematic model chronology.

3. Industries produce as much as they can to satisfy demand, given that they could be limited by lack of critical inputs or lack of workers.
4. If industries do not produce enough, they distribute their production to final consumers and other industries on a pro rata basis, that is, proportionally to demand.
5. Industries update their inventory levels, and profits and labor compensation are distributed to households.

The model is initialized at time $t = 0$ in the steady state corresponding to input data. We apply the pandemic shocks on day $t_{\text{start_lockdown}}$ and keep the economy in lockdown until day $t_{\text{end_lockdown}}$. At that point we remove the supply-side restrictions corresponding to the scenario. Some consumption demand-side shocks remain in place until the pandemic is suppressed on day $t_{\text{end_pandemic}}$. In our scenarios we do not lift shocks of other final demand shocks (investment, international trade). Fig. 2 schematically displays the overall chronology of the model.

3.2 Accounting structure

Let $x_{i,t}$ denote total output of industry i at time t and $Z_{ji,t}$ the intermediate consumption by industry i of good j . Industry i is demand and j is supply. We adopt the standard convention that in the input-output matrix columns represent demand and rows represent supply. In an economy with no “excess” output, i.e. in which all produced output is used up, the output of i is equal to

$$x_{i,t} = \sum_{j=1}^N Z_{ij,t} + c_{i,t} + f_{i,t}, \tag{1}$$

where $c_{i,t}$ is household consumption of good i at time step t and $f_{i,t}$ is all other (exogenous) final demand, including government consumption and exports.

We let $l_{i,t}$ denote labor compensation to workers in industry i . This also indicates the number of workers employed in industry i , under the assumption that all workers employed in the same industry earn the same wage. Profits of industry i can then be written as

$$\pi_{i,t} = x_{i,t} - \sum_{j=1}^N Z_{ji,t} - l_{i,t} - e_{i,t}, \tag{2}$$

where $e_{i,t}$ represents all other expenses (taxes, imports, etc.). Note that we do not model physical capital explicitly, and we take prices as time-invariant.

For total output, total labor income, total profits and total household consumption we write

$$\tilde{x}_t \equiv \sum_{i=1}^N x_{i,t}, \tag{3}$$

$$\tilde{l}_t \equiv \sum_{i=1}^N l_{i,t}, \tag{4}$$

$$\tilde{\pi}_t \equiv \sum_{i=1}^N \pi_{i,t}, \tag{5}$$

$$\tilde{c}_t \equiv \sum_{i=1}^N c_{i,t}, \tag{6}$$

respectively. We focus on these four variables when discussing aggregate economic impacts of the pandemic in subsequent sections.

Our analysis is based on the UK economy. We use the latest release of the World Input-Output Database (WIOD) (Timmer et al. 2015) to determine the relevant values for gross output $x_{i,0}$, intermediate consumption $Z_{ij,0}$, household consumption $c_{i,0}$, other final demand $f_{i,0}$, labor compensation $l_{i,0}$, and all other expenses $e_{i,0}$ (2014 values). Overall, we can distinguish 55 separate industries.

3.3 Demand

It will become important to distinguish between demand, that is orders placed by customers to suppliers, and actual realized transactions. All the steps outlined above are realized transactions, which might or might not be equal to demand.

Industry demand. The total demand faced by industry i at time t , $d_{i,t}$, is the sum of the demand from all its customers,

$$d_{i,t} = \sum_{j=1}^N O_{ij,t} + c_{i,t}^d + f_{i,t}^d, \tag{7}$$

where $O_{ij,t}$ (for *orders*) denotes the demand from industry j , $c_{i,t}^d$ the demand from households and $f_{i,t}^d$ all other final demand.

Recipes. Industries produce output according to a production recipe encoded in the technical coefficient matrix A , where the element $A_{ij} = Z_{ij,0}/x_{j,0}$ is the expense in input i per unit of output j . We will relax the assumption of fixed production recipes, since not every input is critical for production in the short-run (see Appendix C). Industries always demand and aim to consume inputs according to their recipe, even if lacking non-critical inputs does not cause immediate effects on its output in the short time horizon considered here.

Inventories. Due to the dynamic nature of the model, production and demand are not immediate. Instead industries use an inventory of inputs in production. We let $S_{ij,t}$ denote the stock of material i held in j 's inventory. Each industry j aims to keep a target inventory $n_j Z_{ij,0}$ of

every required input i to ensure production for n_j further days⁷. We explain how we calibrate the parameters n_j in Appendix B.

Intermediate demand. Intermediate demand follows the dynamics originally introduced by [Henriet et al. \(2012\)](#) and adopted by [Inoue & Todo \(2019\)](#) in the context of firm-level production network models. To satisfy incoming demand (from $t - 1$) and to reduce the gap to its target inventory, an industry j makes orders to its suppliers at every time step t . More specifically, industry j demands from industry i

$$O_{ij,t} = A_{ij}d_{j,t-1} + \frac{1}{\tau}[n_j Z_{ij,0} - S_{ij,t}], \tag{8}$$

where τ indicates how quickly an industry adjusts its demand due to an inventory gap. Small τ corresponds to responsive industries that aim to close inventory gaps quickly. In contrast, if τ is large, intermediate demand adjusts slowly in response to inventory gaps. In the literature we find different choices for τ , ranging from 1 ([Henriet et al. 2012](#)) to 30 ([Hallegatte 2012](#)) time steps. In our simulations, we choose an intermediate value $\tau = 10$. We present sensitivity tests with respect to τ in Appendix D.4.

Consumption demand. We let consumption demand for good i be

$$c_{i,t}^d = \theta_{i,t} \tilde{c}_t^d, \tag{9}$$

where $\theta_{i,t}$ is a preference coefficient, giving the share of goods from industry i out of total consumption demand \tilde{c}_t^d . The coefficients $\theta_{i,t}$ evolve exogenously, following assumptions on how consumer preferences change during the various phases of the pandemic; see Section 3.5, Eq. (25).

Total consumption demand evolves following an adapted and simplified version of the consumption function in [Muellbauer \(2020\)](#). In particular, \tilde{c}_t^d evolves according to

$$\log \tilde{c}_t^d = \rho \log \tilde{c}_{t-1}^d + \frac{1-\rho}{2} \log (\tilde{m}_t^l) + \frac{1-\rho}{2} \log (\tilde{m}_t^p) + \tilde{\epsilon}_t, \tag{10}$$

where \tilde{l}_t is current labor income, \tilde{l}_t^p is an estimation of permanent income and m is the share of labor income that is used to consume final domestic goods, i.e. that is neither saved nor used for consumption of imported goods. From our data we find $m = 0.82$. Consumption demand during the pandemic is affected by a change of permanent income expectations and the exogenous shock term $\tilde{\epsilon}_t$; see Section 3.5, Eqs. (22) and (24). The parameter ρ indicates sluggish adjustment to new consumption levels. Assuming that a time step corresponds to a quarter, [Muellbauer \(2020\)](#) takes $\rho = 0.6$, implying that more than 70% of adjustment to new consumption levels occurs within two and a half quarters. We modify ρ to account for our daily timescale: By letting $\bar{\rho} = 0.6$, we take $\rho = 1 - (1 - \bar{\rho})/90$ to obtain the same time adjustment as in [Muellbauer \(2020\)](#)⁸. Note that, in the steady state, by definition permanent income corresponds to current income, i.e. $\tilde{l}_t^p = \tilde{l}_t$, and thus total consumption demand corresponds to $m\tilde{l}_t$.⁹

⁷Considering an input-specific target inventory would require generalizing n_j to a matrix with elements n_{ij} , which is easy in our computational framework but difficult to calibrate empirically.

⁸In an autoregressive process like the one in Eq. (10), about 70% of adjustment to new levels occurs in a time t related inversely to the persistency parameter ρ . Letting Q denote the quarterly timescale considered by [Muellbauer \(2020\)](#), time to adjustment t^Q is given by $t^Q = 1/(1 - \bar{\rho})$. Since we want to keep approximately the same time to adjustment considering a daily time scale, we fix $t^D = 90t^Q$. We then obtain the parameter ρ in the daily timescale such that it yields t^D as time to adjustment, namely $1/(1 - \rho) = t^D = 90t^Q = 90/(1 - \bar{\rho})$. Rearranging gives the formula that relates ρ and $\bar{\rho}$.

⁹To see this, note that in the steady state $\tilde{c}_t^d = \tilde{c}_{t-1}^d$. Moving the consumption terms on the left hand side and dividing by $1 - \rho$ throughout yields $\log \tilde{c}_t^d = \log (m\tilde{l}_t) + \tilde{\epsilon}_t$. With no exogenous shock, we find $\tilde{c}_t^d = m\tilde{l}_t$.

To test robustness, we present model results for alternative consumption functions in Appendix D.4. We find that our simulations are highly robust against alternative consumption models.

Other components of final demand. In addition, an industry i also faces demand $f_{i,t}^d$ from sources that we do not model as endogenous variables in our framework, such as government or industries in foreign countries. $f_{i,t}^d$ is not affected by the dynamics of the model. We discuss the composition and calibration of $f_{i,t}^d$ in detail in Section 3.5.

3.4 Supply

Every industry aims to satisfy incoming demand by producing the required amount of output. Production is subject to the following two economic constraints:

Productive capacity. First, an industry has finite production capacity $x_{i,t}^{\text{cap}}$, depending on the amount of available labor input. Initially every industry employs $l_{i,0}$ of labor and produces at full capacity $x_{i,0}^{\text{cap}} = x_{i,0}$. We assume that productive capacity depends linearly on labor inputs,

$$x_{i,t}^{\text{cap}} = \frac{l_{i,t}}{l_{i,0}} x_{i,0}^{\text{cap}}. \quad (11)$$

Input bottlenecks. Second, the production of an industry might be constrained due to an insufficient supply of critical inputs. This can be caused by production network disruptions. While the empirical intermediate consumption at the initial time step is embodied in the technical coefficient matrix A , not every input is necessarily critical for production. Modeling the severeness of intermediate input constraints realistically requires an understanding of how critical inputs are in the production of a given industry (Barrot & Sauvagnat 2016).

We use the ratings of IHS Markit analysts to differentiate three types of inputs: *critical*, *important* and *non-critical* inputs (Appendix C). If an industry runs out of critical inputs, economic production halts immediately. Conversely, if an industry runs out of non-critical inputs, we assume that economic production is not affected. We also have ratings on *important but not critical* inputs. As a baseline we treat important inputs as non-critical. In Section 5.1 we investigate in detail how alternative assumptions on the input-production relationship affect economic impacts.

For a given set of inputs if there are no limits on production capacities, industry i can produce

$$x_{i,t}^{\text{inp}} = \min_{j \in \mathcal{V}_i} \left\{ \frac{S_{j,i,t}}{A_{ji}} \right\}, \quad (12)$$

where \mathcal{V}_i is the set of *critical* suppliers to industry i . If every input is critical, every input is binding, and this reduces to a Leontief production function.

Output level choice and input usage. Since an industry aims to satisfy incoming demand within its production constraints, realized production at time step t is

$$x_{i,t} = \min\{x_{i,t}^{\text{cap}}, x_{i,t}^{\text{inp}}, d_{i,t}\}. \quad (13)$$

Thus production level of an industry is constrained by the smallest of three values: labor-constrained production capacity $x_{i,t}^{\text{cap}}$, intermediate input-constrained production capacity $x_{i,t}^{\text{inp}}$, or total demand $d_{i,t}$.

The level of output then determines the actual use of inputs according to the production recipe. Industry i uses an amount $A_{ji}x_{i,t}$ of input j , unless j is not critical and the amount of j in i 's inventory is less than $A_{ji}x_{i,t}$. In this case, the quantity consumed of input j by industry i is equal to the remaining inventory stock of j -inputs $S_{ji,t} < A_{ji}x_{i,t}$.

Rationing. Without any adverse shocks, industries are always able to meet total demand, i.e. $x_t = d_t$. But in case of production capacity or/and input bottlenecks, industries may not be able to meet total demand, $x_{i,t} < d_{i,t}$, so they need to ration their output. We assume simple proportional rationing, although alternative rationing mechanisms could be considered (e.g. [Inoue & Todo \(2019\)](#)).

The final delivery from industry i to industry j then is the share of orders received

$$Z_{ij,t} = O_{ij,t} \frac{x_{i,t}}{d_{i,t}}. \tag{14}$$

Households receive a share of their demand

$$c_{i,t} = c_{i,t}^d \frac{x_{i,t}}{d_{i,t}}, \tag{15}$$

and the realized final consumption of agents with exogenous final demand is

$$f_{i,t} = f_{i,t}^d \frac{x_{i,t}}{d_{i,t}}. \tag{16}$$

Inventory updating. The inventory of j for every input i is updated according to

$$S_{ij,t+1} = \min \{ S_{ij,t} + Z_{ij,t} - A_{ij}x_{j,t}, 0 \}. \tag{17}$$

In a Leontief production function, where every input is critical, the minimum operator would not be needed since production could never continue once inventories are run down. It is necessary here, since when inventories of non-critical inputs i are depleted, industry j produces output using less goods i than $A_{ij}x_{j,t}$.

Hiring and firing. Firms adjust their labor force depending on which production constraints in Eq. (13) are binding. If the capacity constraint $x_{i,t}^{\text{cap}}$ is binding, industry i decides to hire as many workers as necessary to make the capacity constraint no longer binding. Conversely, if either input constraints $x_{i,t}^{\text{inp}}$ or demand constraints $d_{i,t}$ are binding, industry i lays off workers until capacity constraints become binding. More formally, at time t labor demand by industry i is given by $l_{i,t}^d = l_{i,t-1} + \Delta l_{i,t}$, with

$$\Delta l_{i,t} = \frac{l_{i,0}}{x_{i,0}} \left[\min \{ x_{i,t}^{\text{inp}}, d_{i,t} \} - x_{i,t}^{\text{cap}} \right]. \tag{18}$$

Note that the term $l_{i,0}/x_{i,0}$ reflects the assumption that the labor share in production is constant over the considered period. We assume frictions in the labor market in a sense that it takes time for firms to adjust their labor inputs. Specifically, we assume that industries can increase their labor force only by a fraction γ_H in direction of their target. Similarly, industries can decrease their labor force only by a fraction γ_F in the direction of their target. In the absence of additional policies we usually have $\gamma_F > \gamma_H$, indicating that it is easier for firms to lay off employed than hire new workers. Industry-specific employment evolves then according to

$$l_{i,t} = \begin{cases} l_{i,t-1} + \gamma_H \Delta l_{i,t} & \text{if } \Delta l_{i,t} \geq 0, \\ l_{i,t-1} + \gamma_F \Delta l_{i,t} & \text{if } \Delta l_{i,t} < 0. \end{cases} \tag{19}$$

As we discuss further in Section 3.6, γ_H and γ_F can be interpreted as policy variables. For example, the implementation of a furloughing scheme makes re-hiring of employees easier, corresponding to an increase in γ_H . In our baseline simulations we choose $\gamma_H = 1/30$ and $\gamma_F = 2\gamma_H$. Given our daily time scale, this is a rather rapid adjustment of the labor force. We present sensitivity tests for these parameters in Appendix D.4.

3.5 Pandemic shock

Timeline. The simulation starts in the steady state. For simplicity we let the pandemic shock hit at the same time as the lockdown starts, i.e. we do not take into account reduced demand beforehand. We let the lockdown last for two months (60 time units), and then lift it according to the specifications below.

Supply shocks. At every time step during the lockdown an industry i experiences an (exogenous) first-order labor supply shock $\epsilon_{i,t}^S \in [0, 1]$ that quantifies labor reductions. These reductions are caused by the lack of labor that was previously provided by workers in non-essential industries (del Rio-Chanona et al. 2020, Fana et al. 2020, Galasso 2020) who cannot work remotely (del Rio-Chanona et al. 2020, Dingel & Neiman 2020, Gottlieb et al. 2020, Koren & Petó 2020). For instance, if an industry is non-essential, and none of its employees can work from home, it faces a labor supply reduction of 100% during lockdown i.e., $\epsilon_{i,t}^S = 1$, $\forall t \in [t_{\text{start_lockdown}}, t_{\text{end_lockdown}}]$. Instead, if an industry is classified as fully essential, it faces no labor supply shock and $\epsilon_{i,t}^S = 0 \forall t$.

Letting $l_{i,0}$ be the initial labor supply before the lockdown, the maximum amount of labor available to industry i at time t is given as

$$l_{i,t}^{\max} = (1 - \epsilon_{i,t}^S)l_{i,0}. \quad (20)$$

If $\epsilon_{i,t}^S > 0$, the productive capacity of industry i will be smaller than in the initial state of the economy. We assume that the reduction of total output is proportional to the loss of labor. In that case the productive capacity of industry i at time t is

$$x_{i,t}^{\text{cap}} = \frac{l_{i,t}}{l_{i,0}} x_{i,0}^{\text{cap}} \leq (1 - \epsilon_{i,t}^S)x_{i,0}. \quad (21)$$

Recall from Section 3.4 that firms can hire and fire to adjust their productive capacity to demand and supply constraints. Thus, productive capacity can be lower than the initial supply shock. However, during lockdown they can never hire more than $l_{i,t}^{\max}$ workers. If the lockdown is unbound for an industry i , first-order supply shocks are removed, i.e. we set $\epsilon_{i,t}^S = 0$, for $t \geq t_{\text{end_lockdown}}$.

Supply shock calibration. To initialise the economic model with first-order supply shocks from the pandemic, we use the shock predictions of the recent study by del Rio-Chanona et al. (2020). In del Rio-Chanona et al. (2020) supply shocks of the pandemic are derived by quantifying which work activities of different occupations can be performed from home (Remote Labor Index) and by using the occupational compositions of industries. Moreover, the predictions also take into account whether an industry is essential in the sense that it needs to continue operating during a lockdown. The predictions of first-order shocks are based on the US economy using a different industrial classification system. These predictions therefore need to be adopted for the UK economy and the WIOD industry classification as we outline in detail in Appendix A.

For the UK we estimate that 67% of the work force has an essential job. However, much of this essential work can be done remotely (e.g. government and financial services). In total we estimate that 44% of workers can work remotely and that 37% of workers are currently going to work, assuming that people work from home whenever possible.

Consumption demand shocks. A first shock to consumption demand occurs through reductions in current income and expectations for permanent income. Expectations for permanent income depend on whether households expect a V-shaped vs. L-shaped recovery, that is, whether they expect that the economy will quickly bounce back to normal or there will be a prolonged recession. Let expectations for permanent income \tilde{l}_t^p be specified by

$$\tilde{l}_t^p = \xi_t \tilde{l}_0 \tag{22}$$

In this equation, the parameter ξ_t captures the fraction of pre-pandemic labor income \tilde{l}_0 that households expect to retain in the long run. We first give a formula for ξ_t and then explain the various cases.

$$\xi_t = \begin{cases} 1, & t < t_{\text{start_lockdown}}, \\ \xi^L = 1 - \frac{1}{2} \frac{\tilde{l}_0 - \tilde{l}_{t_{\text{start_lockdown}}}}{\tilde{l}_0}, & t \in [t_{\text{start_lockdown}}, t_{\text{end_lockdown}}], \\ 1 - \rho + \rho \xi_{t-1} + \nu_{t-1}, & t > t_{\text{end_lockdown}}. \end{cases} \tag{23}$$

Before lockdown, we let $\xi_t \equiv 1$. During lockdown, following [Muellbauer \(2020\)](#) we assume that ξ_t is equal to one minus half the relative reduction in labor income that households experience due to the direct labor supply shock, and denote that value by ξ^L . (For example, given a relative reduction in labor income of 16%, $\xi^L = 0.92$.)¹⁰ After lockdown, we assume that 50% of households believe in a V-shaped recovery, while 50% believe in an L-shaped recovery. We model these expectations by letting ξ_t evolve according to an autoregressive process of order one, where the shock term ν_t is a permanent shock that reflects beliefs in an L-shaped recovery. With 50% of households believing in such a recovery pattern, it is $\nu_t \equiv -(1 - \rho)(1 - \xi^L)/2$.¹¹

In addition to the income effect, during a pandemic consumption/saving decisions and consumer preferences over the consumption basket are changing, leading to first-order demand shocks ([Congressional Budget Office 2006](#), [del Rio-Chanona et al. 2020](#)). For example, consumers are likely to demand less services from the hospitality industry, even if it is able to supply these services. Transport is very likely to face substantial demand reductions, despite being classified as an essential industry in many countries. A key question is whether reductions in demand for “risky” goods and services is compensated by an increase in demand for other goods and services, or if lower demand for risky goods translates into higher savings.

We consider a demand shock vector ϵ_t , whose components $\epsilon_{i,t}$ are the relative changes in demand for goods of industry i . These components evolve in the various phases of the pandemic,

¹⁰During lockdown, labor income may be further reduced due to firing. For simplicity, we choose not to model the effect of these further firings on permanent income.

¹¹The specification in Eq. (23) reflects the following assumptions: (i) time to adjustment is the same as for consumption demand, Eq. (10); (ii) absent permanent shocks, $\nu_t = 0$ after some t , ξ_t returns to one, i.e. permanent income matches current income; (iii) with 50% households believing in an L-shaped recovery, ξ_t reaches a steady state given by $1 - (1 - \xi^L)/2$: with $\xi^L = 0.92$ as in the example above, ξ_t reaches a steady state at 0.96, so that permanent income remains stuck four percentage points below pre-lockdown current income.

as defined in the following equations:

$$\epsilon_{i,t} = \begin{cases} 0, & \text{if } t < t_{\text{start_lockdown}}, \\ \epsilon_i^D, & \text{if } t_{\text{start_lockdown}} \leq t < t_{\text{end_lockdown}}, \\ 0, & \text{if no on-site consumption of } i \text{ \& } \\ & t \geq t_{\text{end_lockdown}}, \\ \frac{\epsilon_i^D}{\log 100} \log \left(100 - \frac{99t}{t_{\text{end_pandemic}}} \right), & \text{if on-site consumption of } i \text{ \& } \\ & t_{\text{end_lockdown}} \leq t < t_{\text{end_pandemic}}, \\ 0, & \text{if } t \geq t_{\text{end_pandemic}}. \end{cases} \quad (24)$$

We use the estimates by the [Congressional Budget Office \(2006\)](#), [del Rio-Chanona et al. \(2020\)](#), which we denote by ϵ_i^D , to parameterize $\epsilon_{i,t}$ during lockdown. Roughly speaking, these shocks are massive for restaurants and transport, mild for manufacturing, null for utilities, and positive for health (see [Appendix A](#)).

When the lockdown is lifted, demand shocks for industries that do not involve on-site consumption are removed; in contrast, demand for industries that involve on-site consumption (restaurants, theatres, etc.)¹² goes back to normal more slowly, and in a non-linear way. The specification in [Eq. \(24\)](#) captures the idea that demand for on-site consumption industries is likely to resume very slowly after lockdown and to accelerate towards its pre-pandemic level as the pandemic approaches an end (or at least is perceived to come to a conclusion).¹³ Recent evidence from transaction data in China ([Chen et al. 2020](#)) backs the assumption that demand for these industries resumes more slowly than for industries that do not face on-site consumption. An illustration for three industries that either do not experience any demand shock, experience a demand shock only during lockdown or experience a demand shock throughout the pandemic is given in [Fig. 3](#).

We now explain how the demand shock vector affects consumption demand. Recall from [Eq. \(9\)](#), $c_{i,t}^d = \theta_{i,t} \tilde{c}_t^d$, that consumption demand is the product of the total consumption scalar \tilde{c}_t^d and the preference vector θ_t , whose components $\theta_{i,t}$ represent the share of total demand for good i . We initialize the preference vector by considering the initial consumption shares, that is $\theta_{i,0} = c_{i,0} / \sum_j c_{j,0}$. By definition, the initial preference vector θ_0 sums to one, and we keep this normalization at all following time steps. To do so, we consider an auxiliary preference vector $\bar{\theta}_t$, whose components $\bar{\theta}_{i,t}$ are obtained by applying the shock vector $\epsilon_{i,t}$. That is, we define $\bar{\theta}_{i,t} = \theta_{i,0}(1 - \epsilon_{i,t})$ and define $\theta_{i,t}$ as

$$\theta_{i,t} = \frac{\bar{\theta}_{i,t}}{\sum_j \bar{\theta}_{j,t}} = \frac{(1 - \epsilon_{i,t})\theta_{i,0}}{\sum_j (1 - \epsilon_{j,t})\theta_{j,0}}. \quad (25)$$

The difference $1 - \sum_i \bar{\theta}_{i,t}$ is the aggregate reduction in consumption demand due to the demand shock, which would lead to an equivalent increase in the saving rate. However, households may not want to save all the money that they are not spending. For example, they most likely want to spend on food the money that they are saving on restaurants. Therefore, we define the

¹²For deciding whether an industry faces on-site consumption we use the same list that we compiled for the epidemic model, supplementing it with industries that are not very infectious collectively, but that individually could be perceived as risky. For example, infections while buying a car are a negligible share of all infections, but visiting a car seller might be perceived as risky. Specifically, we classify as industries involving on-site consumption those with the following codes: G45, G47, H49, H50, H51, H52, H53, I, L68, M69_M70, O84, P85, R,S, T.

¹³Note that the specification in [Eq. \(24\)](#) also allows for a small bump in consumption demand at the time the lockdown is lifted.




Figure 3: **Illustration of demand shocks.** Electricity (blue line) does not experience any demand shock; paper manufacturing (green line) only experiences a moderate demand shock during lockdown; restaurants (red line) experience a strong demand shock during lockdown, and then the demand shock remains persistent in the initial phase of reopening, disappearing only as the end of the pandemic approaches.

aggregate demand shock $\tilde{\epsilon}_t$ in Eq. (10) as

$$\tilde{\epsilon}_t = \Delta s \left(1 - \sum_{i=1}^N \bar{\theta}_{i,t} \right) (1 - \rho), \tag{26}$$

where Δs is the change in the savings rate. When $\Delta s = 1$, households save all the money that they are not planning to spend on industries affected by demand shocks; when $\Delta s = 0$, they spend all that money on goods and services from industries that are affected less. For our simulations, we take an intermediate value $\Delta s = 0.5$. Finally, the term $(1 - \rho)$ is needed to account for the autoregressive process in Eq. (10).¹⁴

Demand shock calibration. Note that WIOD distinguishes five types of final demand: (I) *Final consumption expenditure by households*, (II) *Final consumption expenditure by non-profit organisations serving households*, (III) *Final consumption expenditure by government* (IV) *Gross fixed capital formation* and (V) *Changes in inventories and valuables*. Additionally, all final demand variables are available for every country. The endogenous consumption variable $c_{i,0}$ corresponds to (I), but only for domestic consumption. All other final demand categories, including all types of exports, are absorbed into $f_{i,0}$.

We apply different initial shocks to the different demand categories presented above. For domestic final demand variables we assume the following initial shocks: We use the estimates from Congressional Budget Office (2006) and del Rio-Chanona et al. (2020) to calibrate the consumption demand shock variable ϵ_i^D which we apply to the final consumption variables (I) and (II). We assume that investment (IV) is reduced by 5.6%, in line with the US Bureau of Economic Analysis (BEA) estimates for the reduction in investment in the US from 2019Q4 to 2020Q1. We do not apply any exogenous shocks to categories (III) *Final consumption expenditure by government* and (V) *Changes in inventories and valuables*.

¹⁴If $\tilde{\epsilon}_t$ was constant, in the steady state $\log \tilde{c}_t^d$ would be reduced by $\Delta s \left(1 - \sum_{i=1}^N \bar{\theta}_{i,t} \right)$.

To initialise the model with foreign demand shocks, we use the recent estimates on trade by the World Trade Organisation. In their recent forecast international trade is predicted to decline between 12-33% for European countries (Bekkers et al. 2020). We follow the pessimistic scenario of the WTO and assume a drop of 33% in foreign intermediate and final demand.

A summary of all shocks is provided in Appendix A, Table 5. There is considerable uncertainty in our estimates of first-order demand shocks, which we aim to reduce in the future by collecting additional data. However, sensitivity tests shown in Appendix D.1 suggest that our model predictions are fairly robust against uncertainties in the shock estimates.

3.6 Policy intervention

An exogenous policymaker – the government – can influence economic outcomes in three possible ways. First, the key policy which we are considering is the implementation and withdrawal of a lockdown. While the implementation of a lockdown affects all industries simultaneously according to the exogenous first-order supply and demand shocks, the lockdown can be unwound for different sets of industries. We experiment with different re-opening scenarios which we also evaluate with respect to their impact on infectious contagion (Section 6).

Second, the government can also pay out additional social benefits to workers to compensate income losses. During the pandemic only a fraction of the initial labor force is employed, due both to direct shocks and subsequent firing/furloughing, resulting in lower labor compensation, i.e. $\tilde{l}_t < \tilde{l}_0$, for $t \geq t_{\text{start.lockdown}}$. The government can reimburse a fraction b of the income loss $\tilde{l}_t - \tilde{l}_0$ as social benefits, increasing disposable income of households to

$$\tilde{l}_t^* = \tilde{l}_t + b(\tilde{l}_0 - \tilde{l}_t). \quad (27)$$

Following the current UK policy on furloughing, we set $b = 0.8$ in our default simulations.

As a third policy dimension we consider labor force adjustment parameters γ_H and γ_F . Recall from Eq. (19) that the larger these parameters, the quicker firms can adjust their labor inputs. Hiring and firing of employees can be costly without further support by the government. We assume that a furloughing policy scheme increases the flexibility of adding and removing labor inputs to an industry.

We explore in a somewhat stylized way the effect of furloughing on the economy by varying parameters b , γ_H and γ_L . Setting these parameters to larger values represents a regime where furloughing is encouraged by the government, whereas smaller values indicate the business-as-usual scenario without furloughing.

4 Economic impact of COVID-19 on the UK economy

We now show results of the economic model and compare model predictions to data. We focus on the baseline calibration discussed above. For convenience all model parameters are reported again in Table 2.

We let the model start in the steady state at the beginning of 2020. The economy rests in steady state until March 23rd, at which point we apply the pandemic shock. For this simulation we assume that lockdown lasts two months, until May 23rd, at which point all supply-side restrictions are unwound. We show this specific scenario for illustration purposes, while we consider other reopening scenarios in Sections 5.3 and 6.5. We let the model run for another month and a half, until the end of June, to analyse its recovery path. We do not run the model further in the future, both because of the great uncertainties involved and because our assumptions on non-critical inputs are only valid for a limited time span.

Table 2: Parameters of the economic model for our baseline simulations.

Name	Symbol	Value
Consumption function		Eq. (10)
Production function		Eqs. (11)-(13)
Inventory adjustment	τ	10
Upward labor adjustment	γ_H	1/30
Downward labor adjustment	γ_F	1/15
Consumption adjustment	ρ	0.987
Government benefits	b	0.8
Change in savings rate	Δs	0.5

Fig. 4 shows model results for production (gross output); results for other important variables, such as profits, consumption and labor compensation (net of government benefits) are similar. When the lockdown starts, there is a sudden drop in economic activity, shown by a sharp decrease in production. A second smaller drop in production occurs at the beginning of April, due to some service sectors further reducing production. Throughout the simulation, however, service sectors tend to perform better than manufacturing, trade, transport and accommodation sectors. The main reason is that most service sectors face both lower supply and demand shocks, as a high share of workers can effectively work from home, and there is no on-site consumption for most business and professional services. In fact, consumption even increases for several industries (consumption of health is an example).

When the lockdown is lifted, the economy starts approaching its previous level, but this is not all achieved by the end of June. While some sectors quickly return close to pre-lockdown levels, recovery for other sectors is much slower. For example, Restaurants and Transport (green lines) recover very slowly, due to the assumption that consumers are cautious towards industries that involve on-site consumption (see Section 3.5 and Fig. 3). The aggregate level of consumption also does not return to pre-lockdown levels, due to a reduction in expectations of permanent income associated with beliefs in an L-shaped recovery (Section 3.3), and due to the fact that we do not remove shocks to investment and exports (see Section 3.5).

Considering both the lockdown period and the post-lockdown partial recovery, our forecast for GDP in the second quarter of 2020 compared to the last quarter of 2019 is -21.5%. This estimate is more pessimistic than the majority of forecasts for the UK economy done by economic institutions and consulting firms,¹⁵ which, on average, are around -15%. However, it is more optimistic than the estimate by the Bank of England, which predicts a -25% reduction in aggregate GDP.¹⁶

To test how realistic the results of our model are, we compare as many model predictions to data as possible. For aggregate data, we focus on the UK, also considering Spain, France, and Italy.

The UK recently released early estimates of national accounts in the first quarter of 2020.¹⁷ Because lockdown started in the UK only on March 23, the impact on the UK economy over all the first quarter was modest. Indeed, GDP reduced by about 2%, and consumption decreased by 1.7%. Running our model until the end of March, we find a GDP reduction of 1.7% and

¹⁵<https://www.gov.uk/government/collections/data-forecasts>

¹⁶See the monetary policy report of May 2020 (<https://www.bankofengland.co.uk/-/media/boe/files/monetary-policy-report/2020/may/monetary-policy-report-may-2020>), section 2 “Current economic conditions”.

¹⁷<https://www.ons.gov.uk/economy/grossdomesticproductgdp/bulletins/gdpfirstquarterlyestimateuk/januarytomarch2020>




Figure 4: **Economic production as a function of time.** We plot production (gross output) as a function of time for each of the 55 industries. Aggregate production is a thick black line and each sector is colored. Agricultural and industrial sectors are colored red; trade, transport, and restaurants are colored green; service sectors are colored blue. All sectoral productions are normalized to their pre-lockdown levels, and each line size is proportional to the steady-state gross output of the corresponding sector.

a consumption reduction of 1.3%. Labor compensation actually increased in the UK by 0.8%, while in our model it decreased by 0.3%. While data for labor compensation are probably not affected by only a week of lockdown, it is interesting that both in data and in our model consumption decreased more than income, a peculiar feature of this pandemic-induced recession (Muellbauer 2020).

Spain, France, and Italy have reported a larger effect. Remarkably, recently released data from statistical offices reveals that all these countries expect a decline in GDP of around 5% in 2020-Q1, which is substantially larger than the 2% reduction in the UK. Since this is mostly due to these countries starting widespread lockdowns between one and two weeks earlier than the UK, we rerun our model starting lockdown on March 15. In this case, we find that quarterly GDP decreases by 3.3%. This is somewhat off the 5% mark, but part of the error probably comes from the fact that we calibrate our model on UK data (e.g. IO tables), and we assume no reduction in GDP at all before March 15 (in contrast to the evidence of disruptions in supply chains, reduced international travel and early reaction by some consumers prior to this date).

So far, we considered testing data aggregated at the national level, and only pertaining at most to the first 15 days of lockdown. However, we want to test the predictions of our model at a detailed sectoral level, and also explore how our model fares deeper into the lockdown period. The best sectoral data that we could find were released by the States of Washington and Texas, which released weekly unemployment claims data at a high level of industrial disaggregation. In particular, Washington released data up to 6-digit NAICS, while Texas released data for 17 broad industries.¹⁸ To compare with the predictions of our model, we make the hypothesis that

¹⁸We had to clean the data imputing some missing information and to do a crosswalk from NAICS to WIOD




Figure 5: **Comparison of model predictions to data.** In both panels, on both axes, we show the ratio between employment 42 (Washington) or 35 (Texas) days after the shock and employment before the shock. Dots are scaled according to employment shares. The dots labels refer to industry codes, see Appendix for the corresponding names. The identity line is shown for reference.

relative reductions in employment across different sectors are similar in the UK as in Washington and Texas; for example, restaurant workers are more likely to be fired or furloughed than workers in food manufacturing in all these places. We consider unemployment claims from March 14 through April 25 for Washington, and through April 18 for Texas. While March 14 was not the official start of lockdown for either State, unemployment claims started to spike during that week, making it the ideal starting point to compare the predictions of our model to data. To run the model for the same time span as the data, we run it for 42 days after imposing lockdown to compare to Washington data, and for 35 days for Texas data.

Figure 5 shows the ratio of employment levels on April 25 to employment levels on March 14, both in the model and in the Washington data, across all sectors. The Pearson correlation between the model's predictions and the data is 0.44, and the correlation weighted by the employment share of each sector is 0.66. This indicates that predictions for the largest sectors are more accurate. In most cases the model somewhat overestimates the reduction in employment. However, in a few cases employment reduction is actually underestimated, for example in health (Q, large dot on the right).

Comparison with Texas data yields similar results, except that in this case the model vastly overestimates the number of firings. However, correlations are higher in this case, as the Pearson correlation coefficient is 0.68 and the weighted correlation coefficient is 0.72.

We perform the same comparison between model predictions and empirical data for alternative specifications in Appendix D.5. Considering all sensitivity cases studied in Appendix D, the correlation coefficients listed above are very robust. The only exception is the case in which we consider a Leontief production function, where the correlation between model predictions and data even becomes negative. The correlation is also low if we consider important inputs as critical or half-critical.

Overall, we take both the aggregate and sectoral results as an indication that the outcomes of the model are in good qualitative agreement with the reality. There remains some significant quantitative differences, but also a substantial margin for improvement as we have not fitted many parameters, demand shocks can be improved, and we have to compare a UK model to state-level US data. We also take the result as a clear indication that the Leontief production function produces predictions at odds with empirical data, supporting our modeling choice of considering non-critical inputs.

sector. Details are available upon request.

5 Supply, demand and network effects

5.1 Production and input criticality

A key innovation of our model is that it uses a production function that distinguishes between critical and non-critical inputs, based on the IHS Markit analyst ratings. To better understand how different assumptions on input criticality influence outcomes, we implemented various specifications.

The most rigid case is the Leontief production function, in which every positive entry in the technical coefficient matrix A is a binding input to an industry, reducing “input bottlenecks”, Eq. (12), to

$$x_{i,t}^{\text{inp}} = \min_{\{j: A_{ji} > 0\}} \left\{ \frac{S_{ji,t}}{A_{ji}} \right\}, \quad (28)$$

In this case an industry would halt production immediately if inventories of any input are run down, even for small and potentially negligible inputs .

At the other extreme, we implement a linear production function for which all inputs are perfectly substitutable. Here, production in an industry can still continue if inputs cannot be provided, as long as there is sufficient supply of alternative inputs. Thus, there is no input bottlenecks for individual inputs, however, production can be constrained if the input inventories are insufficient. In this case we have

$$x_{i,t}^{\text{inp}} = \frac{\sum_j S_{ji,t}}{\sum_j A_{ji}}. \quad (29)$$

Note that intermediate inputs are perfectly substitutable in this case, but a lack of labor supply cannot be compensated by other inputs.

We also implemented three different production function specifications based on the criticality ratings of IHS Markit analysts. Recall that we define *critical* inputs as those that were rated 1, *non-critical* those that were rated 0, and *important* those that were rated 0.5. First, we use the baseline production function of the main text where we set all input ratings of 0.5 equal to 0 (this makes them non-critical inputs). Second, we set all 0.5 input ratings equal to 1 to make them critical inputs. This moves us closer to the Leontief production function. As a third case, we implement the specification where an industry’s production scales proportionally to the 0.5 rating of important inputs. Thus, we have

$$x_{i,t}^{\text{inp}} = \min_{\{j \in \mathcal{V}_i, \forall k \in \mathcal{U}_i\}} \left\{ \frac{S_{ji,t}}{A_{ji}}, \frac{1}{2} \left(\frac{S_{ki,t}}{A_{ki}} + x_{i,0}^{\text{cap}} \right) \right\}, \quad (30)$$

where \mathcal{V}_i is the set of *critical* inputs and \mathcal{U}_i is the set of *important* inputs to industry i . This means that if an *important* input goes down by 50% compared to initial levels, production of the industry would decrease by 25%. In case the stock of this input is fully depleted, production drops to 50% of initial levels.

In Fig. 6 we show simulation results on total production \tilde{x}_t , labor compensation \tilde{l}_t , profits $\tilde{\pi}_t$ and household consumption \tilde{c}_t ; Eqs. (3)-(6). Note that value added is the sum of profits and labor compensation. In the simulation, $t_{\text{start_lockdown}} = 2$,¹⁹ and there is no re-opening in the following six months (180 days).

As expected we find the largest drop for all economic variables for the Leontief production function, where every input can potentially become binding (black line). For the Leontief

¹⁹Choosing a time step for applying the pandemic shock is arbitrary, as the model rests in steady state beforehand. In Section 4 we chose March 23rd as the starting date of lockdown to easily compare to UK quarterly data, in the following we apply the pandemic shock at $t = 2$, with no loss in generality.

economy our model predicts roughly a 70% drop in gross output, consumption and labor compensation over the next six months if the lockdown continues.

Again in line with intuition, we obtain the mildest economic impacts for the linear production function (blue). There is still a substantial drop in all variables due to first-order shocks, but little further adjustment resulting from network effects.

The production function specifications using the results of the IHS Markit survey lie somewhere in between these two extremes. Treating all important inputs as critical also yields a severe drop in economic production and all other variables (green). The red line indicates the scenario where all important inputs are considered to be non-critical, which is the baseline specification of our model. The results for this case are more similar to the linear production function, although slightly more severe due to the higher risk of lacking critical inputs. Treating important inputs as half-critical such that output scales with the inputs by the factor of 1/2 is the ‘median’ scenario in these simulations, as indicated by the orange line.

In Appendix D.3 we show the recovery paths after re-opening the entire economy. Although the production functions can yield substantially different economic impacts in lockdown, they converge in the long-run after the economy is re-opened.

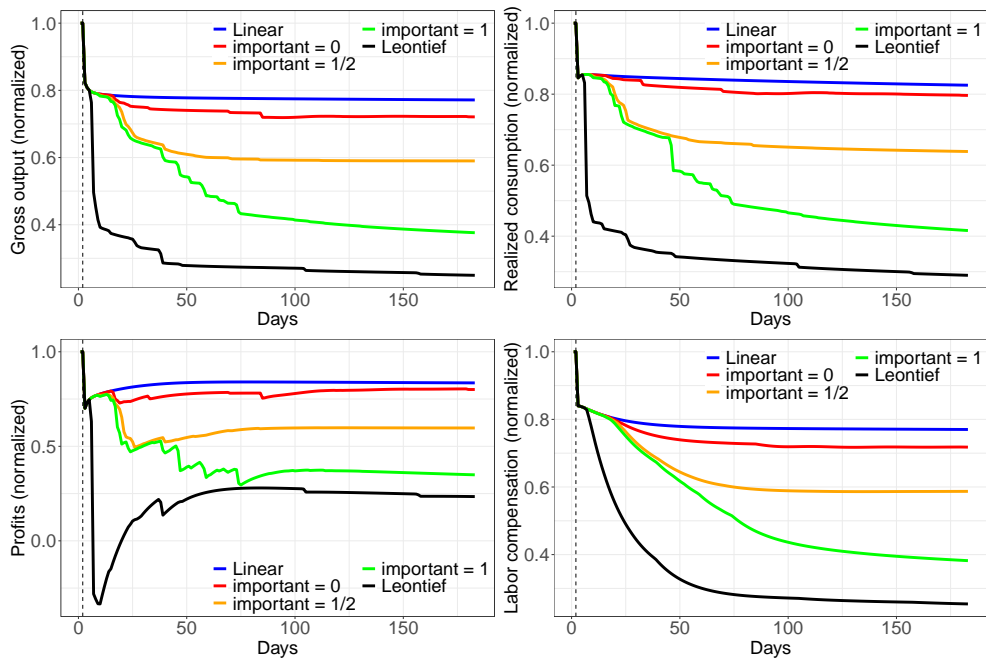


Figure 6: **Comparison of production functions for indefinite lockdown.** Total output, realized consumption, profits and labor compensation for five production function specifications after the lockdown is imposed at $t = 2$ (dashed vertical line). All values are normalized to the initial no-lockdown steady state and there is no re-opening of industries. Linear (blue) denotes a linear production. The red line represents the baseline production function with all important (0.5) inputs being set to be non-critical. The orange line is the case where production scales with a factor of 1/2 with important inputs. The green line (important=1) is the case where all important inputs are treated as critical. The black line the Leontief production function where all inputs are critical.

5.2 First-order shocks, shock propagation and total impact

Our model is initialised with the first-order supply and demand shocks discussed in Section 3.5 and simulates how the economic system translates these shocks into overall economic impacts. Recall that exogenous supply shocks lead to an immediate reduction in gross output. We let the direct output shock be

$$OS_i^{\text{direct}} \equiv \frac{x_{i,0}^{\text{shocked}} - x_{i,0}}{x_{i,0}} = \frac{(1 - \epsilon_{i,0}^S)x_{i,0} - x_{i,0}}{x_{i,0}} = -\epsilon_{i,0}^S, \tag{31}$$

which is equivalent to first-order supply shocks. Similarly, exogenous demand shocks instantaneously decrease final consumption. More formally, we define direct final consumption shocks as

$$CS_i^{\text{direct}} \equiv \frac{c_{i,0}^{\text{shocked}} + f_{i,0}^{\text{shocked}} - c_{i,0} - f_{i,0}}{c_{i,0} + f_{i,0}}. \tag{32}$$

The upper panel of Fig. 7 visualises for each sector of the UK economy the reduction in gross output OS_i^{direct} and final consumption CS_i^{direct} as a result of first-order shocks. While some industries such as Forestry and Logging (A02) face larger immediate reductions in gross output than final consumption, transport industries (H49-51) experience much larger negative shocks in final consumption. Reductions in both final consumption and gross output are enormous for Accommodation and Food Service Activities (I).

We also quantify higher-order impacts on gross output (supply-side) and final consumption (demand-side). These indirect effects are time-dependent since overall economic performance changes in time as can be seen from the simulations above. Note that higher-order impacts in gross output do not necessarily need to be caused by supply-side shocks but could also result from a lack of demand. Conversely, final consumption reductions can stem from lowered production levels. We let the total output shock at any time denote

$$OS_{i,t}^{\text{total}} = x_{i,t}/x_{i,0} - 1.$$

The indirect output shock is then computed as the residual of the total output shock $OS_{i,t}^{\text{total}}$ and direct output shock OS_i^{direct} . Thus, for the indirect output shock we have

$$OS_{i,t}^{\text{indirect}} = OS_{i,t}^{\text{total}} - OS_i^{\text{direct}}. \tag{33}$$

Similarly, we compute indirect final consumption shocks as

$$CS_{i,t}^{\text{indirect}} = CS_{i,t}^{\text{total}} - CS_i^{\text{direct}}, \tag{34}$$

where $CS_{i,t}^{\text{total}} = (c_{i,t} + f_{i,t})/(c_{i,0} + f_{i,0}) - 1$ is the overall final consumption shock.

The center panel of Fig. 7 shows a scatter plot where the x-axis denotes direct output shocks OS_i^{direct} and the y-axis indirect output shocks $OS_{i,t}^{\text{indirect}}$ for $t = 60$, i.e. two months into lockdown. Points scatter in an inverted L-shape indicating that industries that experience large direct output shocks do not reduce production much more in the course of the lockdown. In contrast, many industries that experience little or no direct shocks to their productive capacities downsize economic production substantially after two months.

The bottom panel of Fig. 7 is the same but for final consumption instead of output. Although similar patterns can be observed, there are also a few larger differences. For the majority of industries there are less extreme direct and indirect effects on final consumption. These industries thus lie closer to the identity line. While all higher-order supply-side effects are non-positive, almost half of the industries face positive higher-order consumption effects, although they tend

to be very small. Note that the vast majority (92%) still face negative total consumption effects. Positive values on the y-axis indicate that higher-order impacts on final consumption are slightly mitigating initial shocks. For example, Transport and Warehousing industries (H49-52) are substantially hit by direct demand shocks. The total impact after two months of lockdown is somewhat below these levels.

We also used our first-order shocks to calibrate simpler traditional input-output models (see [Miller & Blair \(2009\)](#) for an excellent overview). We show in [Appendix D.6](#) that under much simplified model parametrizations we can recover the classic Leontief model and get similar sectoral predictions as the Gosh model in steady state. The problem with these models is that they are not able to take supply and demand shocks into account at the same time. While alternative IO models such as the mixed endogenous/exogenous model ([Dietzenbacher & Miller 2015](#)) can be used with simultaneous supply and demand constraints, the model does not necessarily yield feasible solutions corresponding to positive output and consumption values. This is indeed what we find for the UK economy. Calibrating the mixed endogenous/exogenous IO model to our first-order shocks results in infeasible economic allocations.




Figure 7: *Upper panel: Sectoral first-order shocks in the UK economy on the supply and demand sides.* The axes give %-reductions in gross output (y-axis) and in final consumption (x-axis). A blue disk indicates that the monetary shock in absolute terms is larger on the supply side than on the demand side (blue disks can thus lie above the identity line and vice versa). Disk size corresponds to initial gross output of industries. Details on the shocks and industry labels can be found in Appendix A, Table 5. *Center panel: Comparison of direct and indirect output shocks.* Direct supply shocks shown as reduction in sectoral output (x-axis) plotted against indirect impacts on sectoral production (y-axis). Industries that face a large initial supply shock tend to experience smaller higher-order impacts, while higher-order effects can be large for industries that experienced little or no initial supply shock. *Bottom panel: Comparison of direct and indirect final consumption shocks.* Direct final demand shocks (x-axis) plotted against indirect impacts on final demand (y-axis). Disk size corresponds to initial level of final demand satisfied per industry.

We now demonstrate how the first-order shocks shown in Fig. 7 are amplified through the production network and how they affect overall economic impact. To evaluate the relative contributions of supply and demand shocks to overall outcomes, we run the following simulations: First, we run the baseline model setup where both initial supply and demand shocks are present as discussed above. Second, we run the model considering only supply shocks, i.e. we set all consumer demand shocks to zero, $\epsilon_t^D = 0$, as well as remove all shocks to other final demand categories, i.e. $f_{i,t}^d = f_{i,0}^d$. As a third simulation scenario we run the model with all initial supply shocks switched off and include only initial demand shocks.

Fig. 8 shows four key macroeconomic variables (total output \tilde{x}_t , consumption \tilde{c}_t , profits $\tilde{\pi}_t$ and labor compensation \tilde{l}_t) for all three simulation scenarios, both when lockdown continues (solid lines) and when the lockdown is lifted for all industries after 60 days (dashed lines). In Appendix D.2, we run the same simulations for alternative production functions. It becomes clear that demand shocks lead to much smaller economic impacts than supply shocks (blue vs. red solid lines). On the other hand, we find that the economy recovers much quicker after the lockdown if there are no demand shocks, as indicated by the large positive slope of the red dashed line. When there are only demand shocks, recovery is slow. In this case unwinding the economy from lockdown brings limited positive effects due to persistence in exports and investment shocks, sluggish consumption adaptation and a portion of consumers believing in an L-shaped recovery.

Strikingly, we observe that overall negative economic impact in lockdown is larger if the model economy faces only supply shocks instead of being exposed to supply and demand shocks simultaneously (red vs. black solid line). If demand shocks are absent, total output lies roughly 5% below the baseline scenario where both types of shocks are present, except for the first eleven days.

Why is it that turning off demand shocks leads to larger adverse overall impacts? The reason is that in case of large supply constraints and no reduction in final demand, there is higher competition for relatively few goods. If producers cannot satisfy aggregate demand, they need to ration their output to customers (recall that we use proportional rationing). In case of large aggregate demand every customer receives only a relatively small share, which could be even less for some industries compared to the scenario where demand shocks are turned on. If these goods are critical inputs, production in concerned industries will come down once inventories of these inputs are run down. Thus, removing demand shocks can increase the risk of input bottlenecks in production. Put simply, decreasing final demand of some key intermediate goods ensures continued supply of these intermediate goods to other intermediate industries.

In the particular case considered here, it turns out that several large industries have to reduce production as a consequence of fierce competition for critical inputs, as can be seen in Fig. 9 a). Without demand shocks industries such as Health (Q), Education (P85), Real Estate (L68) or IT (J62-3) produce up to 25% less when compared to the baseline case two weeks after lockdown. Although several other industries such as Manufacturing Chemicals (C20) and Pharmaceuticals (C21) produce substantially more without demand shocks, this does not offset the overall adverse effect on the economy as a whole. This can be thought of as a coordination failure.

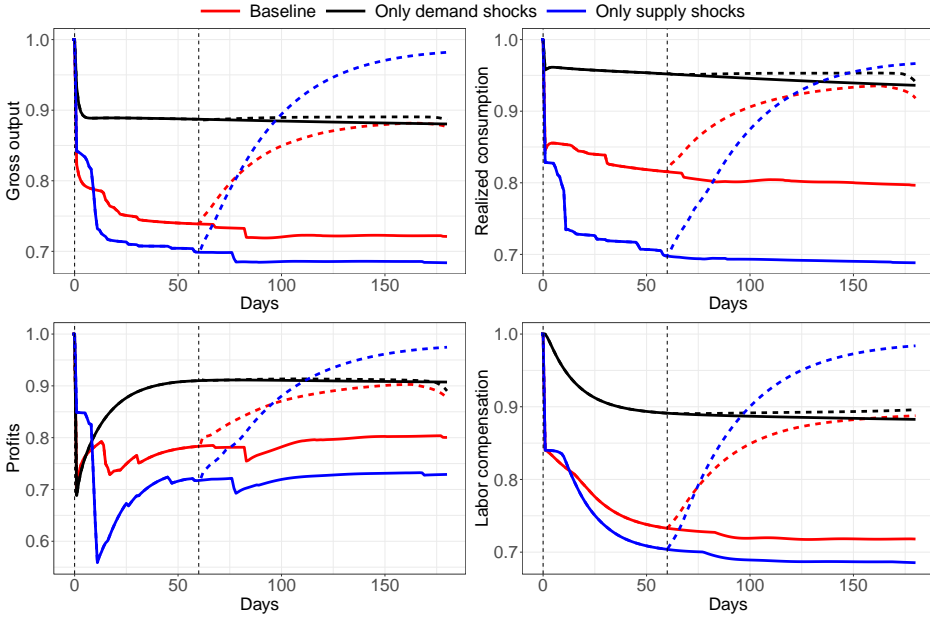


Figure 8: **Dynamic effect of supply shocks vs. demand shocks.** Normalized values of gross output, realized consumption, profits and labor compensation for different shock scenarios. Baseline (red) denotes the model default setup where both supply and demand shocks are used. The black/red line shows the case where only demand/supply shocks are switched on. The lockdown starts at $t = 0$ and ends for all industries after two months at $t = 60$ (vertical dashed lines).

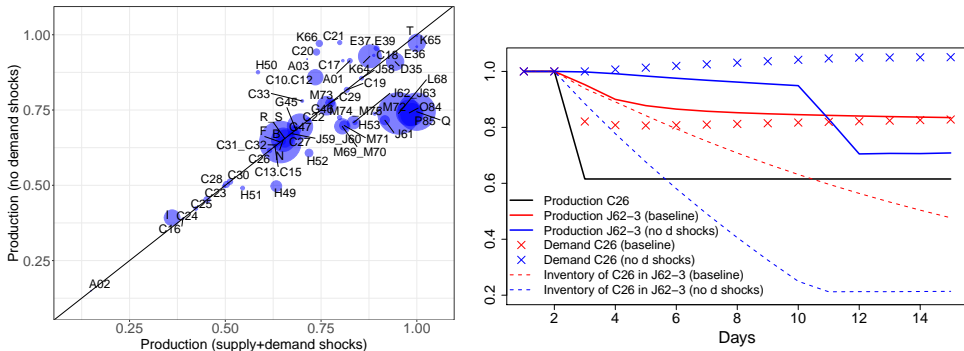


Figure 9: **How can production decrease when demand shocks are removed?** *Left panel:* Comparing normalized production at the sectoral level two weeks after lockdown when supply and demand shocks are turned on (x-axis) and with demand shocks turned off (y-axis). A few large sectors achieve substantially less production if there are no initial demand shocks in the economy. *Right panel:* A sectoral example of downstream shock propagation of the “Baseline” and “Only supply shocks” scenarios in Fig. 8. Sector IT (ISIC J62-3) produces less after 11 time steps if there are no demand shocks in the economy (blue solid line) compared to both supply and demand shocks being present (red solid), since it quickly runs out of critical input C26, Manuf. Electronic, inventories (blue vs. red dashed line). C26 produces the same in both cases (black line) due to binding capacity constraints. If there are no initial demand shocks, C26 faces higher aggregate demand (blue vs. red crosses). Due to higher demand for C26 goods and lower production of C26 goods, IT receives less C26 if there are no demand shocks in the economy.

To better understand how this happens we zoom into the first 15 days of Fig. 8 a). As

Covid Economics 23, 28 May 2020: 79-151

an illustrative example in Fig. 9 b), we show how production constraints in Manufacturing Electronic (C26) lead to larger input bottlenecks in IT (J62-3) in the absence of initial demand shocks. There is no difference in Manufacturing Electronic production (black line) for both scenarios due to binding capacity constraints, but it faces larger aggregate demand if there are no demand shocks (blue crosses) compared to both shocks being present (red crosses). Thus, sector IT for which Manufacturing Electronic goods are critical inputs has to run down its input inventories quicker if there are no demand shocks (blue dashed line), since sector Manufacturing Electronic can deliver less goods in this case. This is reflected in total output of IT. If there are no demand shocks present, IT production is higher for the first few time steps (blue solid line), but drops below the baseline production of both supply and demand shocks (red solid line). We observe similar dynamics for other industries as well.

The case of Manufacturing Electronic production and the coupled output of the IT sector exemplifies the complexity of shock propagation through production networks. We highlight several striking features of this analysis which tend to be entirely neglected in most macroeconomic studies, even if they incorporate industry-specific effects.

First, the specification of production functions and input criticality plays an important role. Most economic analyses use some form of CES production functions with non-zero substitution coefficients. Under this approach, while it is in principle possible to construct elaborated CES nests where different degrees of substitutability are allowed between different inputs of a given sector (Baqae & Farhi 2020), in practice it is so hard to calibrate the parameters that only one (Barrot & Sauvagnat 2016) or a few (Baqae & Farhi 2020, Bonadio et al. 2020) parameters are specified. None of the recent IO papers have considered different degrees of substitution between groups of intermediate inputs. The survey considered here (Appendix C) instead introduces a distinction between critical and non-critical inputs, for each separate industry, allowing us to keep the Leontief assumption of a strong lack of substitutability for critical inputs, which is arguably a key feature of short-run dynamics after large shocks, while at the same time not allowing some non-critical inputs to prevent production. This is a step toward more realism, as in exceptional circumstances like a pandemic, we believe that it is likely that firms can still operate even if several inputs that they usually use are not available. Of course, this is admittedly imperfect and could be improved, and we have made the strong assumption that the lack of use a non-critical input simply does not decrease production and translate into higher profits. Assuming a drop in productivity in this case would change the quantitative results, but would not, however, fundamentally change the dynamics.

Second, the size of inventories held by industries is crucial. Similar to equity buffers in financial distress models, inventories act as buffers against production shocks originating upstream and propagating downstream. Inventory effects are not present in most macroeconomic studies (Favero et al. 2020, Bodenstein et al. 2020, Eichenbaum et al. 2020, McKibbin & Fernando 2020) and IO models (see Table 1), and only appear in a very stylized manner in other empirical work (Mandel & Veetil 2020, Inoue & Todo 2020). Detailed information on input-specific inventories on industry and firm levels, as well as on behaviour and inventory management rules, could vastly improve our understanding of shock propagation in production networks.

Third, dynamics really matter over the short time horizons relevant for the pandemic lockdown and its immediate aftermath. It can take days to weeks for shocks to cascade through several layers of a large production network and our simulations suggest that it can take months until the dynamics reach a steady state. Moreover, the presence of input bottlenecks due to the lack of critical inputs can amplify initial shocks in highly nonlinear ways (see Appendix D.1). The propagation of shocks is path-dependent. This is due to the fact that different industries can have very different customers, resulting in heterogeneous contagion dynamics that depend on “who gets hit first”. General equilibrium models and most input-output models implicitly

assume zero adjustment time and compare pre- and post-shock equilibrium states of the economy to quantify overall impacts. Our analysis suggests that the shock propagation dynamics play an important role in the short time horizons of pandemic lockdowns. In other words, equilibrium comparative static is warranted only when adjustment is faster than the arrival of new shock (Ando et al. 1963). This is not the case currently, where the lifting of the lockdown happens before the system has had a chance to reach the lockdown steady-state.

5.3 Re-opening a network economy

We next investigate how unwinding social distancing measures in certain sectors affects overall economic output. We consider more realistic re-opening scenarios in Section 6.5 together with their impact on infection. Here, we focus on more stylized, theoretically interesting, examples under different production function assumptions to better understand the driving forces behind overall impacts of staged re-opening.

We study the following simulations: As before we represent the economy in lockdown by initialising the model as usual with first-order shocks. We then consider two cases. First, the re-opening scenario where lockdown is relieved after two months for a given set of industries, i.e. for those industries we set $\epsilon_{i,t}^S = 0$ and demand adjusts as discussed in Section 3.5. Second, the lockdown scenario, where the lockdown continues and no shocks are removed. We then compare the two scenarios to quantify the boost in economic activity of re-opening a given industry compared to the lockdown.

Fig. 10 summarises our findings. Each panel shows total production normalized by pre-shock output on the y-axes for both scenarios (re-open sectors in red, continued lockdown in black). The x-axes shows the number of days where day zero is when the lockdown is lifted in the re-opening scenario. Thus, the economy was already two months in lockdown before day zero which is not shown since production is identical for both scenarios during that period. Panel columns represent simulation results for different production function specifications. Panel rows indicate the industries which are re-opened if the lockdown is relaxed.

The economic boost of re-opening varies largely between different sectors and also depends strongly on the production mode assumed. Let us first consider the Leontief production function (left panels). Here, we find a huge increase in economic activity if the highly upstream primary sectors (Agriculture and Mining) are re-opened. Note that primary sectors only account for 2% of UK's total economic output. Opening primary sectors has much smaller effects when using the baseline production function, where inputs are only partially critical, and the linear production function, where inputs are not critical at all.

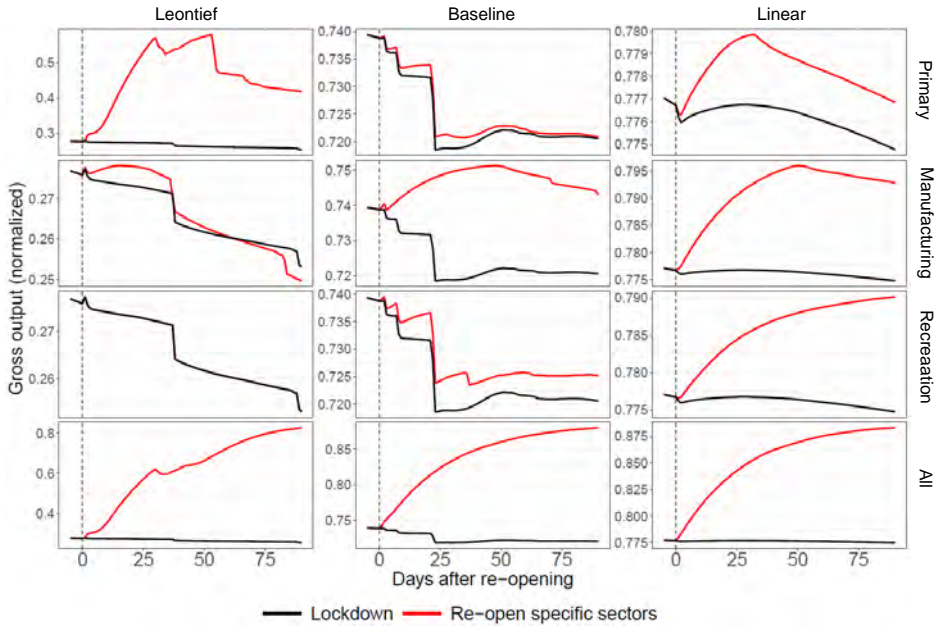


Figure 10: **Economic performance of indefinite lockdown and re-opening single sectors for different production functions.** Black lines indicate total output if the lockdown continues and red lines if only one given sector is re-opened at $t = 0$. Columns represent different production functions, rows denote the sectors which are re-opened. The left column show the effect of re-opening for different sectors in a Leontief economy. The center and right columns show the same for the baseline and linear production functions, respectively. Row 1 shows economic effects of re-opening only primary sectors (ISIC A and B). Row 2 shows the same for opening Manufacturing (ISIC C), row 3 for “Recreation” (ISIC I, R, S) and row 4 for opening all sectors simultaneously. Note that the scale of the y-axes varies between panels.

When re-opening the much larger manufacturing sectors (15% of total output), we obtain a completely different impact on economic output. Strikingly, we find for a Leontief production function that economic output can be lower if manufacturing sectors are opened. The reason is similar to why smaller aggregate shocks can lead to larger overall impacts as discussed in Fig 8. Manufacturing is a large sector relying on many inputs which are critical inputs for other sectors too. Production constraints in other industries might render it impossible to provide larger amounts of those inputs. If manufacturing sectors are re-opened, competition for those scarce inputs increases, resulting in less intermediate consumption for non-manufacturing industries which might face input bottlenecks as a consequence. We do not observe this for the alternative production function setups which relax the strong Leontief assumption. Here, economic output increases by 2-3 percentage points.

We find again very different results when reopening Other Services and Food and Accommodation, here for brevity called recreation. These sectors are large (6% of total output) and heavily affected by the lockdown since they include theaters, hotels, restaurants and other social activities. It is interesting that opening these industries has no impact on overall economic production when assuming a Leontief production function. This is because these are highly downstream industries and their economic output is of little significance for the intermediate consumption of other industries. Thus, opening recreational sectors has mostly demand-side

effects, and given the capacity constraints of upstream sectors this extra demand cannot be satisfied, resulting in no change in overall production. We find positive effects of opening recreational sectors under the alternative production function assumptions where total production increases slightly above lockdown levels.

The bottom row of panels compares the restart of the economy when the lockdown is lifted for all industries simultaneously. The largest economic boost is given in the Leontief economy which re-starts at very low levels after lockdown. The recovery paths are similar of our baseline and the linear production function. Note that recovery is not instantaneous, but takes a considerable amount of time. A month after re-opening, the economy still operates well below initial production levels.

6 The effect of reopening on the reproduction number

We attempt to quantify the effects of the re-opening scenarios on R_0 , the reproduction number of the epidemic that ultimately determines the overall number of deaths.

Rather than coupling a complete epidemiological model with the economic model, we focus on determining R_0 by modelling only the infection rate as it varies across economic scenarios. It is illustrative to consider a simple SIR model,

$$\begin{aligned}\dot{S} &= -\beta SI/M, \\ \dot{I} &= \beta SI/M - \gamma I, \\ \dot{R} &= \gamma I,\end{aligned}\tag{35}$$

where dots over a variable denote its time derivative, S is the number of people who are susceptible, I is the number who are infected, and R is the number who have recovered, and $M = S + I + R$ the total population, which we assume constant.

The model has two parameters, the transmission rate β and the recovery rate γ . We focus on the early stage of the epidemic, that is, when the number of recovered individuals is small with respect to the rest of the population and there is no herd immunity²⁰. The rate of exponential growth in the early stages of the epidemic is $R_0 = \beta/\gamma$, and is the key parameter determining the outcome of the outbreak. When $R_0 < 1$, the outbreak is minimal, but when $R_0 \geq 1$, the outbreak reaches a finite proportion of the population, and larger R_0 implies a larger final number of individuals infected.

While γ is largely unaffected by public health measures absent any treatments for the disease, β depends strongly on public policies and individual behavior. Since we are interested in R_0 , and γ is relatively constant, we focus on modelling β . To compute R_0 , we use

$$R_0(t) = \beta_t/\gamma \propto \beta_t.\tag{36}$$

We now discuss how we compute β_t , and come back to computing R_0 at the end of the section.

6.1 Decomposing infection across activities

The parameter β encompasses two factors: the number of contacts and the risk of infection during a contact. Usually epidemiologists try to distinguish between contacts at home, school, work, and other places (Mossong et al. 2008, Mikolajczyk & Kretzschmar 2008, McCreesh et al. 2019, Strömgren et al. 2017, Ferguson et al. 2020, Du et al. 2020). While we also want to distinguish home-based and work-based contacts, in the context of reopening industries a key question

²⁰To give an upper bound, Vollmer et al. (2020) recently reported an attack rate of around 13% for Lombardy, the most affected region of Italy.

is that of consumption-based infections (Eichenbaum et al. 2020). For instance, reopening cinemas would pose a threat to people working in cinemas, but the number of consumers attending cinemas is vastly larger than the number of workers. Thus, data on the number of contacts at work would fail to capture this risk.

We decompose β into five sources of infection: infections at the workplace (β_w), infections at school (β_s), infections during consumption activities (β_c), infections during commuting (β_T) and infections that are not influenced by whether or not an industry is open (for short, “home” infections β_h). We write

$$\beta(t) = \beta^* \left(\beta_w(t) + \beta_s(t) + \beta_c(t) + \beta_T(t) + \beta_h(t) \right), \tag{37}$$

where β^* is a disease specific parameter, and the components sum up to one when the economy is fully open ($t = 0$),

$$\beta_w(0) + \beta_s(0) + \beta_c(0) + \beta_T(0) + \beta_h(0) = 1. \tag{38}$$

In Appendix E we show how we can derive this equation and the functional form for each $\beta_x(t)$ where $x = w, s, c, T, h$. With this formulation we can measure β^* using data on the speed of diffusion of the virus in a fully open economy. Since there are estimates of R_0 before the lockdown, and considering Eq. (36), to evaluate R_0 during the lockdown and for various scenarios we only need to evaluate the changes to each of the terms within the parenthesis in Eq. (37). To do this we rely on social contact surveys that estimate the intensity-weighted number of contacts of an average person across each activity (see Table 3 and Appendix F for details).

Source of infection	Symbol	Share of intensity-weighted contacts
Work	β_w	0.29
School	β_s	0.28
Consume	β_c	0.16
Transport	β_T	0.06
Home-related	β_h	0.21
Sum		1

Table 3: Calibration of values for Eqs. (37)-(38), using our intensity-weighted share of contacts per activity derived from Strömgen et al. (2017), see Appendix F.

6.2 Decomposing work and consumption infection across industries

Our goal is to focus on what happens when schools and/or a group of industries are opened for work and/or for consumption. To do this we split the population into three categories: a fraction η^s are students and pupils, a fraction η^u are non-working adults, and the rest are workers, which we further split into N industries, each containing a share η_i of the population, so that we have

$$\eta^s + \eta^u + \sum_{i=1}^N \eta_i = 1. \tag{39}$$

The adult non-working population (including the inactive and unemployed population) do not produce output and cannot get infected through the work or study channel. Students and pupils do not produce output, but they do interact with others and can get infected at school. Workers produce economic output and can get infected through the work channel. The economic output and risk of infection of a worker is determined by the industry they work in.

In what follows we consider policy variables that affect each of the β terms in Eq. (37). For instance, one of the key policy variables is

$$\delta_{i,w} \equiv \text{Share of workers of industry } i \text{ that go to work physically.}$$

We defer to the next subsection the details of how we compute this share in each scenario. We now show how each term of Eq. (37) (work, schools, consumption, transport and home) is computed, and in particular how they depend on $\delta_{i,w}$ and other policy variables (see Appendix E and F for more details).

Home-related. We assume that $\beta_h(t)$ is unaffected by whether or not a given industry opens. However, we account for reduction in household interactions during lockdown due to social distancing guidelines. We assume that

$$\beta_h(t) = \beta_h(0) \left((1 - \delta_h) \kappa + \delta_h \right),$$

where δ_h is a policy variable that takes the value 0 if social distancing for family/friends contacts is in place and 1 otherwise, and κ is the share of social/family/friends contacts that are not avoidable by social distancing. We calibrate this by assuming that κ is the share of household/home contacts, and $(1 - \kappa)$ is the share of contacts due to visiting friends and relatives, time in family cars, and contacts in public urban spaces. As discussed in Appendix F, we find $\kappa = 0.76$.

Work-related. Ideally, we would want to know the share of work-related infections that are due to workers of industry i . We are unable to obtain this, but from O*NET data we can estimate an index of exposure to infection. To incorporate this information, we assume

$$\beta_w(t) = \beta_w(0) \sum_{i=1}^N \delta_{i,w} \frac{\eta_i b_{i,w}}{\sum_{k=1}^N \eta_k b_{k,w}}, \tag{40}$$

where $\delta_{i,w}$ is a policy variable that is equal to 1 in the pre-lockdown period (more details below), $b_{i,w}$ is an indicator of intensity weighted number of contacts in industry i , and η_i is the share of population in industry i . To calibrate $b_{i,w}$, we take O*NET occupation-level data on the exposure to infection and on physical proximity. We construct exposure to infection and physical proximity indexes at the industry level by using the share of occupations in each industry, and then construct $b_{i,w}$ as the average of the industry-level exposure to infection and physical proximity.

Schools. We model students and pupils separately (workers in Education face a risk under the “Work-related” category). We assume that the school closure implies that all children above 14 (a share $1 - g$) are not allowed to school, and those at 14 or below (a share g) are allowed if their parents work but cannot work from home. This excludes from school the children of the adult non-working population, and we assume that the students and pupils do not work. Let δ_s be one if schools are open as normal, and zero if they are partially closed. Then, the fraction of the students and pupils population that are attending school is

$$\mu^s = \left(\delta_s + (1 - \delta_s) \left(g \sum_{i=1}^N \delta_{i,w} \eta_i \right) \right). \tag{41}$$

For simplicity, we assume that the school infection rate scales linearly with the fraction of students attending schools as follows,

$$\beta_s(t) = \beta_s(0) \mu^s. \tag{42}$$

Consumption-related. We proceed as for work-related infections, and write

$$\beta_c(t) = \beta_c(0) \sum_{i=1}^N \delta_{i,c} b_{i,c}, \tag{43}$$

where $\delta_{i,c}$ is a policy variable that is equal to 1 in the pre-lockdown period (more details below), and the $b_{i,c}$ are such that $\sum_{i=1}^N b_{i,c} = 1$. To calibrate $b_{i,c}$, we derive from Strömberg et al. (2017) a breakdown of consumption-related contacts into those related to Retail, those related to Restaurants, and those related to Sports Venues. We then map each of these three categories into a single separate WIOD classification.

Public transport. Formally transport is just an industry where consumers risk catching the virus. However, in other consumption-related industries, the number of consumers depends only whether this particular industry is open. Transport is different because the number of people taking transport depends on how many *other* industries are open - if all industries are open, trains are packed and there are more contacts/infections. Therefore we treat Transport separately, and we assume that all transport-related infections are between commuters. If we assume that the number of contacts of one commuter is proportional to the number of other commuters, infections are proportional to the square of the proportion of usual commuters that do commute in a given scenario. Thus, we have

$$\beta_T(t) = \beta_T(0) \left(\frac{\mu^s \eta^s + \sum_{i=1}^N \delta_{i,w} \eta_i}{\eta^s + \sum_{i=1}^N \eta_i} \right)^2 \tag{44}$$

where the left term inside the parenthesis corresponds to student commuters and the right term to work commuters. The term in the denominator is a normalizing factor that guarantees consistency at time $t = 0$. (See Appendix E for details on the derivation).

6.3 Policy scenarios: reopening selected industries

A policy is a set $\Lambda \equiv \left\{ \{\delta_{i,c}\}_{i=1\dots N}, \{\delta_{i,w}\}_{i=1\dots N}, \delta_s, \delta_h \right\}$. It is helpful to first note the values of Λ before and after the lockdown.

Pre-Lockdown. All industries are open for workers and consumers. All schools are open and there is no friends and family social distancing.

$$\begin{aligned} \delta_{i,w}(\text{Pre-Lockdown}) &= 1 \\ \delta_{i,c}(\text{Pre-Lockdown}) &= 1 \\ \delta_s(\text{Pre-Lockdown}) &= 1 \\ \delta_h(\text{Pre-Lockdown}) &= 1 \end{aligned}$$

In this case, $\beta(t) = \beta^*$.

Scenario I: Full lockdown. Workers go to work physically if and only if they are essential and they cannot work from home, which happens for a share of workers equal to $e_{iw}(1 - r_i)$, where e_{iw} is the degree to which the industry is essential, and r_i is the Remote Labor Index (see

Appendix A for details on how we estimate e_{iw} and r_i for each industry). Consumers consume physically only what they can consume physically.

$$\begin{aligned} \delta_{i,w}(\text{Lockdown}) &= e_{i,w}(1 - r_i) \\ \delta_{i,c}(\text{Lockdown}) &= e_{i,c} \\ \delta_s(\text{Lockdown}) &= 0 \\ \delta_h(\text{Lockdown}) &= 0 \end{aligned}$$

where $e_{i,c}$ is an “essential consumption index”. For each industry (in fact, retail is the only relevant one), it indicates how much of infection-related consumption is still open during lockdown. In practice, we assume $e_{i,c} = 0$ for all industries except retail, and $e_{i,c} = e_{i,w}$ for retail, that is, we assume that the share of the retail-based infections that continue during lockdown can be proxied by the share of retail workers who are essential.

We consider four degrees of reopening: roughly speaking, reopening only manufacturing and construction, reopening everything except consumer-facing industries (with or without fully reopening schools), and reopening everything. For all scenarios we assume that friends and family social distancing remains in place, $\delta_h = 0$, and that everyone that can work from home continues to work remotely. Table 4 shows the main scenarios and the key dimensions in which they differ.

Scenario	I Lockdown	II Manufacturing and Construction	III All except consumer- facing.	IV All except consumer- facing, Schools opened	V Open	VI Pre- lockdown
Work from home?	Yes	Yes	Yes	Yes	Yes	No
Schools open?	Partially	Partially	Partially	Fully	Fully	Fully
On-site consumption open?	No	No	No	No	Yes	Yes
Home-related distancing?	Yes	Yes	Yes	Yes	Yes	No

Table 4: Six scenarios for reopening the economy.

Scenario II: Manufacturing and Construction. We label this scenario “Manufacturing and Construction” for short, but we mean opening all A-F industries: Agriculture, Mining, Manufacturing, Utilities and Construction. Because all workers who can work from home do so, we have

$$\begin{aligned} \delta_{i,w}(\text{Manufacturing and Construction}) &= \begin{cases} (1 - r_i) & \forall i \in \text{A-F}, \\ e_{i,w}(1 - r_i) & \forall i \in \text{G-T}. \end{cases} \\ \delta_{i,c}(\text{Manufacturing and Construction}) &= e_{i,c} \quad \forall i. \end{aligned}$$

Scenario III: All except consumer-facing. This means opening all industries (A-T), except the three “consumer-facing” industries G47 (Retail), I (Accommodation and Food), and RS (Other Services, which includes recreation and personal services).

$$\begin{aligned} \delta_{i,w}(\text{All except consumer-facing.}) &= \begin{cases} e_{iw}(1 - r_i) & \forall i \in \{G47, I, RS\}, \\ (1 - r_i) & \text{otherwise.} \end{cases} \\ \delta_{i,c}(\text{All except consumer-facing.}) &= e_{i,c} \quad \forall i. \end{aligned}$$

Scenario IV: All except consumer-facing. + Schools. This is the same as Scenario III, but we assume that schools are fully open instead of receiving only children of workers who cannot work from home, that is $\delta_s = 1$

Scenario V: Open. We reopen all industries for work and consumption.

$$\begin{aligned}\delta_{i,w}(\text{All open}) &= (1 - r_i) \quad \forall i \\ \delta_{i,c}(\text{All open}) &= 1 \quad \forall i.\end{aligned}$$

6.4 Computing R_0

We assume that before the lockdown, $R_0(\text{pre-lockdown}) = 2.6$ (Jarvis et al. 2020). Then we compute how β is reduced by social distancing measures during the lockdown. This gives us an estimate of the lockdown R_0 which we denote \tilde{R}_0

$$\tilde{R}_0 = R_0(\text{pre-lockdown}) \times \frac{\beta(\text{lockdown})}{\beta(\text{pre-lockdown})} \quad (45)$$

We find $\tilde{R}_0 \approx 0.90$. However, the recent study by Jarvis et al. (2020) finds that in the UK during lockdown, $R_0 \approx 0.62$. It is not surprising that we overestimate the lockdown R_0 , as our model does not incorporate all the basic sanitary measures that would apply to the contacts that have not been reduced. For instance, we consider supermarket infections to stay the same because supermarkets are open, but social distancing applies in supermarkets and there are extra cleaning procedures in place. Another example would be essential workers. In our model, essential workers contribute to infections now just as they did before, while in reality it is likely that their conditions have been made at least a bit safer. A final example includes the effect of information campaigns on hand washing. To take this into account, we rescale all our estimates for the scenarios so that they start from a lockdown value at 0.62.

$$R_0(\Lambda) = \frac{0.62}{\tilde{R}_0} \times \frac{\beta(\Lambda)}{\beta(\text{pre-lockdown})} \quad (46)$$

This rescaling implies that fully re-opening the economy back to the pre-lockdown situation (that is, removing all limitations to work, consumption, school and social contact, while keeping the extra sanitary precautions) would bring R_0 to $(0.62/1.04) \times 2.6 = 1.55$.

We obtain standard errors for R_0 as follows. Jarvis et al. (2020) report a pre-lockdown mean R_0 of 2.6 with a standard error of 0.54, that is, a standard error of $0.54/2.6 = 21\%$ of the mean. Their post-lockdown estimate is 0.62 with a 95% confidence interval (0.37 - 0.89), that is, a standard error of $((0.62 - 0.37)/2)/0.62 = 20\%$ of the mean. In view of this, we assume that for all our estimates, one standard error always equal 20% of the mean R_0 estimates. We report confidence intervals as two standard errors around the mean.

6.5 Economic performance vs. infections: sector-specific re-opening

We now show simulations for the four economic scenarios outlined in Section 6.3: lockdown, open manufacturing and construction, open all industries except consumer-facing ones, open all industries (school opening does not affect economic scenario).




Figure 11: **Reopening scenarios.** We show production normalized to pre-lockdown levels, for four months since lockdown starts, assuming lockdown is lifted. We consider the four economic scenarios outlined in Section 6.3.

As can be seen in Figure 11, keeping the lockdown leads to a further reduction in production, due to firms exhausting their inventories. Re-opening only all A-F industries (Agriculture, Mining, Manufacturing, Utilities and Construction) leads to an increase in production, however the jump is much bigger when opening all industries except consumer-facing ones. Conversely, there is not a large benefit in terms of production from opening consumer-facing industries. To compute the increases in production in Figure 1, we consider the first 30 days after lockdown is lifted; as it is easy to see, the economy is far from completely recovering after 30 days, as discussed in Section 4.

There is a trade-off between an increase in production and mitigating the epidemic spread when opening industries. In the bar plot in Fig. 1 we illustrate this trade-off. The bars on the left show our estimates of R_0 , the higher the bar, the faster the epidemic spread. The bars on the right show the GDP (as a percentage of the pre-lockdown GDP) of each scenario. In Scenario II, where Manufacturing and Construction open, the effect on R_0 is negligible, but has a 3 percentage points higher GDP than the lockdown scenario. The negligible increase in R_0 is due to the low percentage of the labor force that resumes work. 15% of the employed people work in A–F industries, and only 10% cannot work from home. Furthermore, 6% were already working on-site due to the essential nature of their work. Thus, under scenario II, only 4% of the employed labor force resume on-site work.

Scenario III presents a slight increase in R_0 with respect to the lockdown scenario. The slight increase is mostly because non-consumer-facing industries from $G - R.S$ have a high remote labor index. Therefore, assuming that all the workers that can telework stay home, the number of people returning to work is small. Scenario III has a 8 percentage points higher GDP than the lockdown scenario. This is due to a direct effect as well as an indirect effect where business services now resume work in tandem with primary and secondary sector, lifting key bottlenecks in supply chains.

Scenario IV, where we include opening schools for all children (not only for the below 14 children of people working on-site), increases R_0 substantially. With 17% of the population being 14 years old or younger, it is not surprising that opening schools increases the speed of the epidemic spread. Since we do not consider productivity decrease due to childcare work,

the GDP of Scenario IV is equal to the GDP in Scenario III. Scenario V, where consumer-facing industries open, increases substantially epidemic spread due to the share of contacts that happen in restaurants, hotels, gyms, etc. However, there is only a 2 percentage point increase in GDP compared to the previous scenario. This is mostly due to a lack of recovery in demand for the hospitality and recreation industries.

7 Discussion

In this paper we have investigated how locking down and re-opening the economy as a policy response to the COVID-19 pandemic affects economic performance and contagion. We introduced a novel economic model specifically designed to address the unique features of the current pandemic. The model is industry-specific, incorporating the production network and inventory dynamics. We use survey results by industry experts to model how critical different inputs are in the production of a specific industry. We calibrate the model to the UK economy and find that two months after lockdown gross output and consumption are down by 27% when compared to pre-lockdown levels.

We find that industries are affected by direct demand and supply shocks in highly heterogeneous ways. While many manufacturing industries face large supply shocks, transport industries experience mostly demand-side shocks. Other industries including hotels and restaurants are substantially exposed to both shock types simultaneously. We find similar industrial heterogeneity for higher-order impacts.

We analysed how shocks propagate through the production network, resulting in non-trivial economic impacts. First, we have shown that input criticality plays an important role in the downstream amplification of shocks. Second, we found that inventory levels can act as buffers against production shocks and are crucial for understanding economic impacts – an aspect usually neglected in other studies. Third, it has become evident that time scales matter as shock propagation is not immediate but takes time. Overall, we find that first-order shocks can be translated into overall impacts in highly nonlinear ways. We even find cases where smaller aggregate shocks can lead to larger economic impacts as a result of unbalanced supply and demand dynamics. This “coordination failure” suggests that it could be dangerous to re-open single sectors of the economy by themselves without understanding how they are embedded in the production network. Our results suggest that the economic boost from opening an industry depends on the up-/downstream location of that industry as well as how severe the economy suffers from input bottlenecks. In case the economy faces serious productive constraints, re-opening a single sector can even have adverse effects on economic output.

There is a trade-off between re-opening the economy and facing an increase in epidemic spreading. To help understand this we develop an epidemic model where the infection rate is divided between different economic-related activities: work, schools, consumption, transport, and others. Within work and consumption, we consider the relative risk of infection between industries. We find that there is little variance between the risk of infection workers face in different industries. In contrast, the risk of infection due to consumption is concentrated in three consumer-facing industries: retail, restaurants and hotels, and other services (including gym and entertainment events). Our results show that keeping consumer-facing industries and schools closed, and having people who can telework work from home can significantly increase the economic output while having a relatively small increase in the spread of the epidemic.

8 Bibliography

- Andersen, A. L., Hansen, E. T., Johannesen, N. & Sheridan, A. (2020), ‘Consumer responses to the covid19 crisis: Evidence from bank account transaction data’, *Covid Economics 7*, 20 April 2020: 88-114 .
- Ando, A., Fisher, F. M. & Simon, H. A. (1963), *Essays on the structure of social science models*, MIT Press.
- Arenas, A., Cota, W., Gomez-Gardenes, J., Gómez, S., Granell, C., Matamalas, J. T., Soriano-Panos, D. & Steinegger, B. (2020), ‘A mathematical model for the spatiotemporal epidemic spreading of covid19’, *medRxiv* .
- Arto, I., Andreoni, V. & Rueda Cantuche, J. M. (2015), ‘Global impacts of the automotive supply chain disruption following the japanese earthquake of 2011’, *Economic Systems Research 27*(3), 306–323.
- Baqaei, D. R. & Farhi, E. (2020), ‘Nonlinear production networks with an application to the covid-19 crisis’.
- Barrot, J.-N., Grassi, B. & Sauvagnat, J. (2020), ‘Sectoral effects of social distancing’, *Covid Economics 3*, 10 April 2020: 85-102 .
- Barrot, J.-N. & Sauvagnat, J. (2016), ‘Input specificity and the propagation of idiosyncratic shocks in production networks’, *The Quarterly Journal of Economics 131*(3), 1543–1592.
- Battiston, S., Gatti, D. D., Gallegati, M., Greenwald, B. & Stiglitz, J. E. (2007), ‘Credit chains and bankruptcy propagation in production networks’, *Journal of Economic Dynamics and Control 31*(6), 2061–2084.
- Bekkers, E., Keck, A., Koopman, R. & Nee, C. (2020), ‘Trade and covid-19: The wto’s 2020 and 2021 trade forecast’, *CEPR VoxEU column*, 24 April .
- Bodenstein, M., Corsetti, G. & Guerrieri, L. (2020), ‘Social distancing and supply disruptions in a pandemic’.
- Bonadio, B., Huo, Z., Levchenko, A. A. & Pandalai-Nayar, N. (2020), ‘Global supply chains in the pandemic’, *CEPR DP14766* .
- Borgonovo, E. & Plischke, E. (2016), ‘Sensitivity analysis: a review of recent advances’, *European Journal of Operational Research 248*(3), 869–887.
- Carvalho, V. M., Hansen, S., Ortiz, Á., Garcia, J. R., Rodrigo, T., Rodriguez Mora, S. & Ruiz de Aguirre, P. (2020), ‘Tracking the covid-19 crisis with high-resolution transaction data’.
- Chen, H., Qian, W. & Wen, Q. (2020), ‘The impact of the covid-19 pandemic on consumption: Learning from high frequency transaction data’, *Available at SSRN 3568574* .
- Congressional Budget Office (2006), ‘Potential influenza pandemic: Possible macroeconomic effects and policy issues’. <https://www.cbo.gov/sites/default/files/109th-congress-2005-2006/reports/12-08-birdflu.pdf>.
- del Rio-Chanona, R. M., Mealy, P., Pichler, A., Lafond, F. & Farmer, D. (2020), ‘Supply and demand shocks in the covid-19 pandemic: An industry and occupation perspective’, *Covid Economics 6*, 17 April 2020: 65-103 .
- Dietzenbacher, E. & Miller, R. E. (2015), ‘Reflections on the inoperability input–output model’, *Economic Systems Research 27*(4), 478–486.
- Dingel, J. & Neiman, B. (2020), ‘How many jobs can be done at home?’, *Covid Economics 1*, 3 April 2020: 16-24 .

- Dorn, F., Fuest, C., Götttert, M., Krolage, C., Lautenbacher, S., Link, S., Peichl, A., Reif, M., Sauer, S., Stöckli, M. et al. (2020), 'The economic costs of the coronavirus shutdown for selected european countries: A scenario calculation', *EconPol Policy Brief* 25. https://www.ifo.de/sites/default/files/2020-04/EconPol-Policy-Brief_21_Appendix-March2020.pdf.
- Du, Z., Wang, X., Pasco, R., Petty, M., Fox, S. J. & Meyers, L. A. (2020), 'Covid-19 healthcare demand projections: 22 texas cities'. https://sites.cns.utexas.edu/sites/default/files/cid/files/covid-19_analysis_for_22_texas_cities_033020.pdf?m=1585596798.
- Eichenbaum, M. S., Rebelo, S. & Trabandt, M. (2020), 'The macroeconomics of epidemics', *NBER WP* 26882.
- Fadinger, H. & Schymik, J. (2020), 'The effects of working from home on covid-19 infections and production a macroeconomic analysis for germany', *Covid Economics* 9, 24 April 2020: 107-139.
- Fana, M., Tolan, S., Perez, S. T., Brancati, M. C. U., Macias, E. F. et al. (2020), The covid confinement measures and eu labour markets, Technical report, Joint Research Centre (Seville site).
- Favero, C. A., Ichino, A. & Rustichini, A. (2020), 'Restarting the economy while saving lives under covid-19'.
- Ferguson, N., Laydon, D., Nedjati Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunuba Perez, Z., Cuomo-Dannenburg, G. et al. (2020), 'Report 9: Impact of non-pharmaceutical interventions (npis) to reduce covid19 mortality and healthcare demand'.
- Galasso, V. (2020), 'Covid: Not a great equaliser', *Covid Economics* 18, 15 May 2020: 157-169.
- Ghosh, A. (1958), 'Input-output approach in an allocation system', *Economica* 25(97), 58-64.
- Goldfarb, A. & Tucker, C. (2020), Which retail outlets generate the most physical interactions?, Technical report, National Bureau of Economic Research.
- Gottlieb, C., Grobovšek, J. & Poschke, M. (2020), 'Working from home across countries', *Covid Economics* 8, 22 April, 70-91.
- Hallegatte, S. (2008), 'An adaptive regional input-output model and its application to the assessment of the economic cost of katrina', *Risk Analysis: An International Journal* 28(3), 779-799.
- Hallegatte, S. (2012), 'Modeling the roles of heterogeneity, substitution, and inventories in the assessment of natural disaster economic costs', *The World Bank, Policy Research Working Paper* 6047.
- Henriet, F., Hallegatte, S. & Tabourier, L. (2012), 'Firm-network characteristics and economic robustness to natural disasters', *Journal of Economic Dynamics and Control* 36(1), 150-167.
- Inoue, H. & Todo, Y. (2019), 'Firm-level propagation of shocks through supply-chain networks', *Nature Sustainability* 2(9), 841-847.
- Inoue, H. & Todo, Y. (2020), 'The propagation of the economic impact through supply chains: The case of a mega-city lockdown against the spread of covid-19', *Covid Economics* 2, 8 April 2020: 43-59.
- Jarvis, C. I., Van Zandvoort, K., Gimma, A., Prem, K., Klepac, P., Rubin, G. J., Edmunds, W. J., working group, C. C.-. et al. (2020), 'Quantifying the impact of physical distance measures on the transmission of covid-19 in the uk', *medRxiv*.
- Koren, M. & Petó, R. (2020), 'Business disruptions from social distancing', *Covid Economics* 2, 8 April 2020: 13-31.
- Leontief, W. W. (1936), 'Quantitative input and output relations in the economic systems of the United States', *The Review of Economics and Statistics* pp. 105-125.

- Maier, B. F. & Brockmann, D. (2020), ‘Effective containment explains subexponential growth in recent confirmed covid-19 cases in china’, *Science* .
- Mandel, A. & Veetil, V. P. (2020), ‘The economic cost of covid lockdowns: An out-of-equilibrium analysis’, *Available at SSRN 3588421* .
- McCreech, N., Morrow, C., Middelkoop, K., Wood, R. & White, R. G. (2019), ‘Estimating age-mixing patterns relevant for the transmission of airborne infections’, *Epidemics* **28**, 100339.
- McKibbin, W. J. & Fernando, R. (2020), ‘The global macroeconomic impacts of covid-19: Seven scenarios’, *Covid Economics 10*, 27 April 2020: 116-156 .
- Mikolajczyk, R. T. & Kretzschmar, M. (2008), ‘Collecting social contact data in the context of disease transmission: prospective and retrospective study designs’, *Social Networks* **30**(2), 127–135.
- Miller, R. E. & Blair, P. D. (2009), *Input-output analysis: foundations and extensions*, Cambridge University Press.
- Mossong, J., Hens, N., Jit, M., Beutels, P., Auranen, K., Mikolajczyk, R., Massari, M., Salmaso, S., Tomba, G. S., Wallinga, J. et al. (2008), ‘Social contacts and mixing patterns relevant to the spread of infectious diseases’, *PLoS medicine* **5**(3).
- Muellbauer, J. (2020), ‘The coronavirus pandemic and U.S. consumption’, *VoxEU.org*, 11 April . <https://voxeu.org/article/coronavirus-pandemic-and-us-consumption>.
- Strömgren, M., Holm, E., Dahlström, Ö., Ekberg, J., Eriksson, H., Spreco, A. & Timpka, T. (2017), ‘Place-based social contact and mixing: a typology of generic meeting places of relevance for infectious disease transmission’, *Epidemiology & Infection* **145**(12), 2582–2593.
- Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R. & De Vries, G. J. (2015), ‘An illustrated user guide to the world input–output database: the case of global automotive production’, *Review of International Economics* **23**(3), 575–605.
- Vollmer, M. A. C., Mishra, S., Unwin, H. J. T., Gandy, A. et al. (2020), ‘Report 12: Using mobility to estimate the transmission intensity of covid-19 in italy: a subnational analysis with future scenarios’.

Appendix

A First-order economic shocks and work context industry variables

A.1 NAICS-WIOD mapping of shocks

Due to the COVID-19 pandemic industries experience supply-side reductions due to the closure of non-essential industries and workers not being able to perform their activities at home. Many industries also face substantial reductions in demand. [del Rio-Chanona et al. \(2020\)](#) provide quantitative predictions of these first-order supply and demand shocks for the US economy. To calculate supply-side predictions, [del Rio-Chanona et al. \(2020\)](#) classified industries as essential or non-essential and constructed a Remote Labor Index, which measures the ability of different occupations to work from home. Under the assumption that the distribution of occupations across industries and that the percentage of essential workers within an industry is the same for the US and the UK, we can map the supply-shocks estimated by [del Rio-Chanona et al. \(2020\)](#) into the UK economy as follows.

First, we build a crosswalk from the NAICS 4-digit industry classification to the classification system used in WIOD, which is a mix of ISIC 2-digit and 1-digit codes. We build this crosswalk using the NAICS to ISIC 2-digit crosswalk from the European Commission and then aggregating the 2-digit codes that are presented as 1-digit in the WIOD classification system. We then do an employment-weighted aggregation of the supply shocks from [del Rio-Chanona et al. \(2020\)](#) for the 277 industries at the NAICS 4-digit classification level to the 55 industries in the WIOD classification. Some of the 4-digit NAICS industries map into more than one WIOD industry classification. When this happens we assume employment is split uniformly among the WIOD industries the NAICS industry maps into. Finally, we make one modification to deal with imputed rents for the Real Estate Sector. Imputed rents account for 69% of the monetary value of the sector²¹. We assume that the supply shock does not affect imputed rents for the Real Estate Sector and thus consider that the supply shock only affects 31% of the sector. With this modification the final supply shock to the Real Estate Sector is 15%.

For calibrating consumption demand shocks, we use the same data as [del Rio-Chanona et al. \(2020\)](#) which are based on the [Congressional Budget Office \(2006\)](#) estimates. These estimates are available only on the more aggregate 2-digit NAICS level which are straightforward to map into WIOD ISIC categories. Table 5 gives an overview of all first-order shocks applied to WIOD industries.

A.2 Essential score, remote labor index, and industries' work context

Using the same methodology as before, i.e., doing a crosswalk from NAICS-4 digit to the classification system used in WIOD and using employment shares to aggregate, we map the essential score and remote labor index computed in [del Rio-Chanona et al. \(2020\)](#) into the WIOD list of industries. We use these industry remote labor index and essential score at the WIOD industry classification level to estimate the number of people working in each industry for each scenario i.e., to estimate $\delta_{iw}(t)$.

O*NET provides different Work Context²² indices for occupations, including “Exposure to disease and infection” and “Physical proximity”, for brevity we refer to these indexes as exposure

²¹Calculation obtained from monetary values reported in <https://www.ons.gov.uk/economy/grossvalueaddedgva/datasets/nominalandrealregionalgrossvalueaddedbalancedbyindustry>

²²https://www.onetonline.org/find/descriptor/browse/Work_Context/




Figure 12: **Remote Labor Index of industries.** Remote labor index of the WIOD industry classification. See Table 5 for code-industry name.

to infection and physical proximity. Using the same methodology than [del Rio-Chanona et al. \(2020\)](#) we map these occupation indexes into the NAICS 4-digit industry classification. In particular, we use the data from the BLS, which indicates the occupational composition of each industry, and take the employment weighted average of the occupation’s work context employed in each industry. After computing the exposure to infection, physical proximity and outdoors work for the industries at the NAICS 4-digit industry classification we map them into the WIOD classification system with the above mentioned crosswalk methodology. As we explain in Appendix F we use the exposure to infection and physical proximity index of each industry to estimate the relative risk of contracting COVID-19 workers in each industry face.

Covid Economics 23, 28 May 2020: 79-151

ISIC	Sector	Supply				Cons. demand		Other demand	
		x	ϵ^S	RLI	ess.	c	ϵ^D	f	f shock
A01	Agriculture	0.8	0.0	14	100	0.9	-10	0.3	-20
A02	Forestry	0.0	-85.0	15	0	0.0	-10	0.0	-23
A03	Fishing	0.1	0.0	36	100	0.0	-10	0.1	-32
B	Mining	1.3	-35.3	31	51	0.1	-10	1.4	-33
C10-C12	Manuf. Food-Beverages	2.8	-0.6	22	99	2.5	-10	1.3	-32
C13-C15	Manuf. Textiles	0.4	-37.1	31	47	0.1	-10	0.5	-27
C16	Manuf. Wood	0.2	-61.1	27	18	0.1	-10	0.1	-16
C17	Manuf. Paper	0.4	-7.5	31	89	0.1	-10	0.2	-30
C18	Media print	0.3	-6.0	39	90	0.1	-10	0.1	-18
C19	Manuf. Coke-Petroleum	0.9	-18.3	36	71	1.5	-10	0.7	-33
C20	Manuf. Chemical	1.1	-2.6	37	96	0.3	-10	1.6	-32
C21	Manuf. Pharmaceutical	0.7	-1.1	40	98	0.3	-10	1.2	-31
C22	Manuf. Rubber-Plastics	0.7	-28.3	29	60	0.1	-10	0.6	-28
C23	Manuf. Minerals	0.5	-50.3	36	20	0.1	-10	0.2	-27
C24	Manuf. Metals-basic	0.6	-57.7	27	20	0.0	-10	1.7	-33
C25	Manuf. Metals-fabricated	1.1	-54.8	34	18	0.1	-10	0.7	-18
C26	Manuf. Electronic	0.8	-38.5	57	10	0.2	-10	1.4	-28
C27	Manuf. Electric	0.4	-33.3	37	46	0.1	-10	0.8	-30
C28	Manuf. Machinery	1.1	-49.7	38	20	0.2	-10	2.1	-30
C29	Manuf. Vehicles	1.6	-22.6	30	65	1.3	-10	2.6	-29
C30	Manuf. Transport-other	1.0	-48.8	40	17	0.1	-10	2.5	-31
C31_C32	Manuf. Furniture	0.6	-36.6	35	43	0.2	-10	0.8	-25
C33	Repair-Installation	0.4	-3.3	39	95	0.0	-10	0.0	-17
D35	Electricity-Gas	3.2	0.0	42	100	3.6	0	0.1	-20
E36	Water	0.2	0.0	33	100	0.6	0	0.0	-11
E37-E39	Sewage	0.8	0.0	30	100	0.6	0	1.1	-16
F	Construction	7.9	-35.6	28	51	0.3	-10	11.5	-6
G45	Vehicle trade	1.7	-31.6	45	41	2.1	-10	0.6	-23
G46	Wholesale	3.5	-23.6	50	51	3.3	-10	4.3	-30
G47	Retail	4.7	-30.5	50	37	16.4	-10	0.6	-28
H49	Land transport	2.0	-11.1	31	83	2.7	-67	0.2	-28
H50	Water transport	0.6	-12.4	35	81	0.6	-67	0.7	-32
H51	Air transport	0.6	-0.1	29	100	1.2	-67	0.4	-32
H52	Warehousing	1.4	-0.5	30	99	0.1	-67	0.4	-26
H53	Postal	0.7	0.0	36	100	0.1	0	0.1	-30
I	Accommodation-Food	2.9	-60.8	35	6	8.4	-80	0.7	-32
J58	Publishing	0.6	-14.4	70	48	0.5	0	0.6	-24
J59_J60	Video-Sound-Broadcasting	0.9	-32.8	56	34	1.1	0	1.2	-15
J61	Telecommunications	1.6	-0.9	55	99	1.9	0	0.8	-22
J62_J63	IT	2.3	-0.2	71	99	0.2	0	2.6	-13
K64	Finance	4.3	0.0	71	100	3.2	0	2.9	-32
K65	Insurance	3.2	0.0	71	100	6.3	0	1.5	-31
K66	Auxil. Finance-Insurance	1.1	0.0	72	100	0.2	0	2.0	-33
L68	Real estate	7.8	-4.8	49	6	25.3	0	1.0	-8
M69_M70	Legal	2.8	-2.0	64	94	0.0	0	1.3	-25
M71	Architecture-Engineering	1.7	0.0	54	100	0.1	0	1.4	-19
M72	R&D	0.5	0.0	59	100	0.0	0	1.1	-11
M73	Advertising	0.6	-22.5	60	36	0.0	0	0.3	-30
M74_M75	Other Science	0.7	-3.0	61	94	0.3	0	1.0	-30
N	Private Administration	4.4	-34.9	36	42	1.0	0	2.7	-29
O84	Public Administration	4.8	-1.1	45	97	0.6	0	11.9	-1
P85	Education	4.2	0.0	54	100	1.6	0	8.7	-2
Q	Health	7.0	-0.1	36	100	2.8	15	15.0	1
R_S	Other Service	3.2	-34.5	39	47	5.6	-5	2.5	-10
T	Household activities	0.2	0.0	0	0	0.8	0	0.0	-32

Table 5: **Industry-specific first-order shocks.** Column x denotes relative shares of gross output, ϵ^S the supply shock, RLI the Remote Labor Index and ess. the essential score of industries. Column c represents relative shares of consumer consumption and ϵ^D the demand shock to consumption. Column f denotes relative shares of other final consumption (exports, gross capital formation, inventory changes, government) and f shock the shock to other final demand. All values are in %.

B Inventory data and calibration

We could not find UK data to calibrate inventory target parameters n_j in Eq. (8). The only reliable data that we could find are those from the US National Income and Product Accounts (NIPA). We describe these data and how we processed them in this appendix. We obtain ratios of level of inventories to monthly sales across industries, with a high level of disaggregation for manufacturing industries, and more uniform information for other industries. Among other things, we show that these ratios are remarkably stable over time. We take this evidence as supporting the idea that inventory to sales ratios are fundamental concepts that have to do with the nature of production rather than with specificities of the US economy. Therefore, we consider these ratios as proxies of n_j , after multiplying them by 30 to take into account that we are considering a daily rather than a monthly time scale.

B.1 Data sources

All inventory data have been sourced from the Bureau of Economic Analysis (BEA)'s National Income and Product Accounts (NIPA). In particular, we used data from the following tables:

- Table 5.8.6B. Real Private Inventories and Real Domestic Final Sales by Industry, Chained Dollars - LastRevised: March 26, 2020 [BEA/NIPA-T50806B]
- Table 208. Real Gross Output by industry, Chained Dollars - LastRevised: April 6, 2020 [BEA/GDPbyIndustry-208Q]
- Table 1BU. Real Manufacturing and Trade Inventories, Seasonally Adjusted, End of Period [Chained 2012 dollars, 1997 forward, NAICS] - LastRevised: March 26, 2020 [BEA/NIUnderlyingDetail-U001B]
- Table 2BU. Real Manufacturing and Trade Sales, Seasonally Adjusted at Monthly Rate [Chained 2012 dollars, 1997 forward, NAICS] - LastRevised: March 26, 2020 [BEA/NIUnderlyingDetail-U002BU]
- Table 4BU1. Real Manufacturing Inventories, by Stage of Fabrication (Materials and supplies), Seasonally Adjusted, End of Period [Chained 2012 dollars, 1997 forward, NAICS] - LastRevised: March 26, 2020 [BEA/NIUnderlyingDetail-U004B1]
- Table 4BU2. Real Manufacturing Inventories, by Stage of Fabrication (Work-in-process), Seasonally Adjusted, End of Period [Chained 2012 dollars, 1997 forward, NAICS] - LastRevised: March 26, 2020 [BEA/NIUnderlyingDetail-U004B2]
- Table 4BU3. Real Manufacturing Inventories, by Stage of Fabrication (Finished goods), Seasonally Adjusted, End of Period [Chained 2012 dollars, 1997 forward, NAICS] - LastRevised: March 26, 2020 [BEA/NIUnderlyingDetail-U004B3]

For inventory data, we started considering tables 4BU1, 4BU2 and 4BU3, focusing on all 3-digit NAICS manufacturing sectors. We then added information on inventories in trade sectors from table 1BU, focusing on the following industries: Merchant wholesale industries (NAICS 42 except 4251), motor vehicles, parts, and supplies wholesalers (4231), retail trade industries (44-45), motor vehicle and parts dealers (441). We finally added information on inventories in all other industries from table 5.8.6B.²³

²³We also complemented merchant wholesale industries by adding non-merchant wholesale (4251), so as to recover the wholesale sector as a whole.

We then used tables 208 and 2BU to extract information on gross output by industry. Information about gross output in wholesale and retail trade is not consistent in the two tables: in table 208 (and several other sources), yearly gross output for these sectors is around 1800 billion dollars; in table 2BU, it is around 6000 billion dollars. We use information from table 208, as, according to the BEA, “underlying detail” tables such as 2BU may be less accurate.²⁴

The latest available data (2019Q4) are reported in Table 7, where we also report monthly gross output and the ratio between the level of inventories and monthly gross output. For presentation purposes, we split inventories in certain aggregate sectors according to output shares of subsectors within these sectors — for example, we disambiguate between mining, utilities and construction, although these are given together in table 5.8.6B.

As shown in Fig. 13, the ratios between the level of inventories and monthly gross output are remarkably stable (ratios are normalized to their values in 2019Q4), varying by no more than 20% in the last 10 years. There appears to be an upward trend from 1997, where the average ratio was around 80% of the 2019Q4 value, and the only industries whose ratios increased substantially are apparel manufacturing and leather and allied products manufacturing. We view the temporal stability of these ratios as supporting the idea that they can be used for other countries.

B.2 Mapping to WIOD codes

We next map data from NAICS to the industrial classification used in the World Input Output Database (WIOD), which is an aggregation of 2-digit International Standard Industrial Classification (ISIC) sectors. Using official concordance tables, which are valid for 4-digit NAICS and ISIC codes, is not the best option, as our data are not available to that level of disaggregation. We resort instead to manual mapping between the NAICS sectors for which we have data and the WIOD sectors. In particular, we use the crosswalk available in Table 6. When a NAICS sector maps uniquely to a WIOD sector, we directly attribute inventory and gross output data. When multiple NAICS sectors map to (one or more) WIOD sectors, we aggregate data for all relevant NAICS sectors. When one or more NAICS sector map to multiple WIOD sectors, we attribute data using as weights sectoral gross outputs from the 2014 WIOD table for the U.S.

The results are shown in Table 8. They make sense. As an example, consider the ratio between level of inventories and monthly gross output in the NAICS sectors 313, 314, 315, 316 and in the WIOD sector C13-C15. It is clear that the latter is a weighted average of the former ratios, weighted by the size of the NAICS subsectors.

²⁴We do not find any other discrepancy, for example information provided in tables 4BU1-4BU2-4BU3 was consistent with more aggregate information in table 1BU, and that information was in turn consistent with data from table 5.8.6B.

NAICS	WIOD
321	C16
327	C23
331	C24
332	C25
333	C28
334	C26
335	C27
336	C29, C30
337, 339	C31_C32, C33
311, 312	C10-C12
313, 314, 315, 316	C13-C15
322	C17
323	C18
324	C19
325	C20, C21
326	C22
4231_441	G45
42-4231	G46
44_45-441	G47
11	A01, A02, A03
21_23	B, D35, E36, E37-E39, F
48_92	H49, H50, H51, H52, H53, I, J58, J59-J60, J61, J62-J63, K64, K65, K66, L68, M69_M70, M71, M72, M73, M74_M75, N, O84, P85, Q, R_S, T

Table 6: Crosswalk NAICS to WIOD




Figure 13

Covid Economics 23, 28 May 2020: 79-151

Table 7: Data are from 2019Q4. Everything in millions of 2012 chained dollars.

naics_code	naics_name	stock_mat_sup	work_in_process	finished_goods	total_inv	go_yearly	go_monthly	ratio_inv_go_monthly
321	Wood product manufacturing	4975	2072	4638	11685	93672	7806	1.50
327	Nonmetallic mineral product manufacturing	6497	1728	8456	16681	110484	9207	1.81
331	Primary metal manufacturing	15409	12215	15307	42931	249984	20832	2.06
332	Fabricated metal product manufacturing	25870	20066	25938	71874	368076	30673	2.34
333	Machinery manufacturing	29528	18621	20667	68816	353964	29497	2.33
334	Computer and electronic product manufacturing	24361	19743	16948	61052	335520	27960	2.18
335	Electrical equipment, appliance, and component manufacturing	10743	7907	7457	26107	132144	11012	2.37
336	Transportation equipment manufacturing	34067	98210	17703	149980	936336	78028	1.92
337	Furniture and related product manufacturing	5057	1883	3538	10478	71112	5926	1.77
339	Miscellaneous durable goods manufacturing	11069	6560	15186	32815	151440	12620	2.60
311	Food manufacturing	21590	9238	38713	69541	801432	66786	1.04
312	Beverage and tobacco product manufacturing	8981	5681	7141	21803	124788	10399	2.10
313	Textile mills	2056	1179	2721	5956	28332	2361	2.52
314	Textile product mills	1419	683	2296	4398	23880	1990	2.21
315	Apparel manufacturing	2127	1496	5193	8816	10692	891	9.89
316	Leather and allied product manufacturing	645	311	1288	2244	4572	381	5.89
322	Paper manufacturing	11145	2110	8814	22069	174264	14522	1.52
323	Printing and related support activities	2820	1163	3137	7120	76968	6414	1.11
324	Petroleum and coal product manufacturing	18917	20958	32313	72188	1033320	86110	0.84
325	Chemical manufacturing	41896	23577	52226	117699	734484	61207	1.92
326	Plastics and rubber product manufacturing	14955	3936	15273	34164	223452	18621	1.83
42	Wholesale industries	823900	0	0	823900	1861400	155117	5.31
4231	Motor vehicles, parts, and supplies wholesalers	63705	0	0	63705	138677	11556	5.51
44-45	Retail trade industries	674155	0	0	674155	1844000	153667	4.39
441	Motor vehicle and parts dealers	257911	0	0	257911	363100	30258	8.52
111-113	Farms	0	0	174416	174416	487200	40600	4.30
114-115	Forestry, fishing, and related activities	0	0	19368	19368	54100	4508	4.30
21	Mining	0	0	32662	32662	732800	61067	0.53
22	Utilities	0	0	20316	20316	455800	37983	0.53
23	Construction	60363	0	0	60363	1354300	112858	0.53
48-49	Transportation and warehousing	19737	0	0	19737	1163000	96917	0.11
51	Information	18463	0	0	18463	1999800	166650	0.11
52-53	Finance, insurance, real estate, rental, and leasing	55949	0	0	55949	6060000	505000	0.11
54-56	Professional and business services	37009	0	0	37009	4008500	334042	0.11
61-62	Educational services, health care, and social assistance	25173	0	0	25173	2726500	227208	0.11
71-72-81	Arts, entertainment, recreation, accommodation, and food services	12343	0	0	12343	1336900	111408	0.11
92	Government	33553	0	0	33553	3634200	302850	0.11

Covid Economics 23, 28 May 2020: 79-151

Table 8: Data are from 2019Q4. Everything in millions of 2012 chained dollars.

wiod_code	wiod_name	stock_mat_sup	work_in_process	finished_goods	go_yearly	tot_inv	go_monthly	ratio_inv_go_monthly
A01	Crop and animal production, hunting and related service activities	0	0	172754	482557	172754	40213	4.30
A02	Forestry and logging	0	0	13056	36470	13056	3039	4.30
A03	Fishing and aquaculture	0	0	7974	22273	7974	1856	4.30
B	Mining and quarrying	0	0	31740	712120	31740	59343	0.53
C10-C12	Manufacture of food products, beverages and tobacco products	30571	14919	45854	926220	91344	77185	1.18
C13-C15	Manufacture of textiles, wearing apparel and leather products	6247	3669	11498	67476	21414	5623	3.81
C16	Manufacture of wood and of products of wood and cork, except furniture; manufacture of articles of straw and plaiting materials	4975	2072	4638	93672	11685	7806	1.50
C17	Manufacture of paper and paper products	11145	2110	8814	174264	22069	14522	1.52
C18	Printing and reproduction of recorded media	2820	1163	3137	76968	7120	6414	1.11
C19	Manufacture of coke and refined petroleum products	18917	20958	32313	1033320	72188	86110	0.84
C20	Manufacture of chemicals and chemical products	30865	17369	38475	541095	86709	45091	1.92
C21	Manufacture of basic pharmaceutical products and pharmaceutical preparations	11031	6208	13751	193389	30990	16116	1.92
C22	Manufacture of rubber and plastic products	14955	3936	15273	223452	34164	18621	1.83
C23	Manufacture of other non-metallic mineral products	6497	1728	8456	110484	16681	9207	1.81
C24	Manufacture of basic metals	15409	12215	15307	249984	42931	20832	2.06
C25	Manufacture of fabricated metal products, except machinery and equipment	25870	20066	25938	368076	71874	30673	2.34
C26	Manufacture of computer, electronic and optical products	24361	19743	16948	335520	61052	27960	2.18
C27	Manufacture of electrical equipment	10743	7907	7457	132144	26107	11012	2.37
C28	Manufacture of machinery and equipment n.e.c.	29528	18621	20667	353964	68816	29497	2.33
C29	Manufacture of motor vehicles, trailers and semi-trailers	21551	62129	11199	592339	94879	49362	1.92
C30	Manufacture of other transport equipment	12516	36081	6504	343997	55101	28666	1.92
C31-C32	Manufacture of furniture; other manufacturing	13924	7290	16168	192166	37382	16014	2.33
C33	Repair and installation of machinery and equipment	2202	1153	2556	30386	5911	2532	2.33
D35	Electricity, gas, steam and air conditioning supply	0	0	19237	431602	19237	35967	0.53
E36	Water collection, treatment and supply	0	0	657	14746	657	1229	0.53
E37-E39	Sewerage; waste collection, treatment and disposal activities; materials recovery; remediation activities and other waste management services	0	0	4371	98070	4371	8173	0.53
F	Construction	57335	0	0	1286362	57335	107197	0.53
G45	Wholesale and retail trade and repair of motor vehicles and motorcycles	321616	0	0	501777	321616	14815	7.69
G46	Wholesale trade, except of motor vehicles and motorcycles	760195	0	0	1722723	760195	143560	5.30
G47	Retail trade, except of motor vehicles and motorcycles	416244	0	0	1480900	416244	123408	3.37
H49	Land transport and transport via pipelines	5299	0	0	573956	5299	4785	0.11
H50	Water transport	653	0	0	70739	653	5890	0.11
H51	Air transport	1957	0	0	211979	1957	17665	0.11
H52	Warehousing and support activities for transportation	2031	0	0	219937	2031	18328	0.11
H53	Postal and courier activities	1144	0	0	123877	1144	10323	0.11
I	Accommodation and food service activities	9256	0	0	1002493	9256	83541	0.11
J58	Publishing activities	3427	0	0	371174	3427	30931	0.11
J59-J60	Motion picture, video and television programme production, sound recording and music publishing activities; programming and broadcasting activities	3351	0	0	362983	3351	30249	0.11
J61	Telecommunications	6735	0	0	729434	6735	60786	0.11
J62-J63	Computer programming, consultancy and related activities; information service activities	5816	0	0	629926	5816	52494	0.11
K64	Financial service activities, except insurance and pension funding	7307	0	0	791410	7307	65951	0.11
K65	Insurance, reinsurance and pension funding, except compulsory social security	10347	0	0	1120663	10347	93389	0.11
K66	Activities auxiliary to financial services and insurance activities	4967	0	0	537998	4967	44833	0.11
L68	Real estate activities	28803	0	0	3119741	28803	259978	0.11
M69-M70	Legal and accounting activities; activities of head offices; management consultancy activities	11781	0	0	1276047	11781	106337	0.11
M71	Architectural and engineering activities; technical testing and analysis	4626	0	0	501018	4626	41752	0.11
M72	Scientific research and development	2492	0	0	269945	2492	22495	0.11
M73	Advertising and market research	2510	0	0	271811	2510	22651	0.11
M74-M75	Other professional, scientific and technical activities; veterinary activities	819	0	0	88716	819	7393	0.11
N	Administrative and support service activities	11268	0	0	1220478	11268	101706	0.11
O84	Public administration and defence; compulsory social security	35648	0	0	3861125	35648	321760	0.11
P85	Education	3317	0	0	359267	3317	29939	0.11
Q	Human health and social work activities	21640	0	0	2343861	21640	195322	0.11
R.S	Other service activities	7819	0	0	846949	7819	70579	0.11
T	Activities of households as employers; undifferentiated goods- and services-producing activities of households for own use	216	0	0	23374	216	1948	0.11
U	Activities of extraterritorial organizations and bodies							

Covid Economics 23, 28 May 2020: 79-151

C Critical vs. non-critical inputs

A survey was designed to address the question when production can continue during a lockdown. For each industry, IHS Markit analysts were asked to rate every input of a given industry. The exact formulation of the question was as follows: “For each industry in WIOD, please rate whether each of its inputs are essential. We will present you with an industry X and ask you to rate each input Y. The key question is: Can production continue in industry X if input Y is not available for two months?” Analysts could rate each input according to the following allowed answers:

- 0 – This input is *not* essential
- 1 – This input is essential
- 0.5 – This input is important but not essential
- NA – I have no idea

To avoid confusion with the unrelated definition of essential industries which we used to calibrate first-order supply shocks, we refer to inputs as *critical* and *non-critical* instead of *essential* and *non-essential*.

Analysts were provided with the share of each input in the expenses of the industry. It was also made explicit that the ratings assume no inventories such that a rating captures the effect on production if the input is not available.

Every industry was rated by one analyst, except for industries Mining and Quarrying (B) and Manufacture of Basic Metals (C24) which were rated by three analysts. To improve input ratings, we aim to increase the sample size of analyst ratings for every industry in the next few weeks. In case there are several ratings we took the average of the ratings and rounded it to 1 if the average was at least 2/3 and 0 if the average was at most 1/3. Average input ratings lying between these boundaries are assigned the value 0.5.

The ratings for each industry and input are depicted in Fig. 14. A column denotes an industry and the corresponding rows its inputs. Blue colors indicate *critical*, red *important, but not critical* and white *non-critical* inputs. Note that under the assumption of a Leontief production function every element would be considered to be critical, yielding a completely blue-colored matrix. The results shown here indicate that the majority of elements are non-critical inputs (2,338 ratings with score = 0), whereas only 477 industry-inputs are rates as critical. 365 inputs are rates as important, although not critical (score = 0.5) and NA was assigned eleven times.

The left panel of Fig. 15 shows for each industry how often it was rated as critical input to other industries (x-axis) and how many critical inputs this industry relies on in its own production (y-axis). Electricity and Gas (D35) are rated most frequently as critical inputs in the production of other industries (score=1 for almost 60% of industries). Also frequently rated as critical are Land Transport (H49) and Telecommunications (J61). On the other hand, many manufacturing industries (ISIC codes starting with C) stand out as relying on a large number of critical inputs. For example, around 27% of inputs to Manufacture of Coke and Refined Petroleum Products (C19) as well as to Manufacture of Chemicals (C20) are rated as critical.

The center panel of Fig. 15 shows the equivalent plot for 0.5 ratings (important, but not critical inputs). Financial Services (K64) are most frequently rated as important inputs which do not necessarily stop the production of an industry if not available. Conversely, the industry relying on many important, but not binding inputs is Wholesale and Retail Trade (G46) of which almost half of its inputs got rated with a score = 0.5. This makes sense given that this industry




Figure 14: Criticality scores of IHS Markit analysts. Rows are inputs (supply) and columns industries using these inputs (demand). The blue color indicates critical (score=1), red important (score=0.5) and white non-critical (score=0) inputs. Black denotes inputs which have been rated with NA. The diagonal elements are considered to be critical by definition. For industries with multiple input-ratings we took the average of all ratings and assigned a score=1 if the averaged score was at least 2/3 and a score=0 if the average was smaller or equal to 1/3.




Figure 15: (Left panel) Plotting how often an industry is rated as a critical input to other industries (x-axis) against the share of critical inputs this industry is using. The center and right panel are the same as the left panel, except for using half-critical and non-critical scores, respectively. In each plot the identity line is shown. Point sizes are proportional to gross output.

heavily relies on all these inputs, but lacking one of these does not halt economic production. This case also illustrates that a Leontief production function could starkly overestimated input

bottlenecks as Wholesale and Retail Trade would most likely still be able to realize output even if a several inputs would not be available.

In the right panel of Fig. 15 we show the same scatter plot but for non-critical inputs. 25 industries are rated to be non-critical inputs to other industries in 80% of all cases, with Household Activities (T) and Manufacture of Furniture (C31-32) being rated as non-critical in at least 96%. Industries like Other Services (R-S), Other Professional, Scientific and Technical Activities (M74-75) and Administrative Activities (N) rely on mostly non-critical inputs (>90%).

A detailed breakdown of the input- and industry-specific ratings are given in Table 9.

ISIC	Sector (abbreviated)	Input-based rankings				Industry-based rankings				n
		1	0.5	0	NA	1	0.5	0	NA	
A01	Agriculture	4	2	49	0	9	9	37	0	1
A02	Forestry	2	3	50	0	7	9	39	0	1
A03	Fishing	2	1	52	0	8	5	42	0	1
B	Mining	7	1	47	0	9	2	44	0	3
C10-C12	Manuf. Food-Beverages	5	6	44	0	14	5	36	0	1
C13-C15	Manuf. Textiles	2	5	48	0	6	2	47	0	1
C16	Manuf. Wood	3	3	49	0	8	3	44	0	1
C17	Manuf. Paper	5	10	40	0	14	11	30	0	1
C18	Media print	3	6	46	0	6	3	46	0	1
C19	Manuf. Coke-Petroleum	18	4	33	0	15	6	33	2	1
C20	Manuf. Chemical	21	10	24	0	15	6	34	0	1
C21	Manuf. Pharmaceutical	2	2	51	0	9	17	25	7	1
C22	Manuf. Rubber-Plastics	11	7	37	0	14	5	36	0	1
C23	Manuf. Minerals	8	2	44	2	7	1	47	0	1
C24	Manuf. Metals-basic	8	2	45	0	12	7	36	0	3
C25	Manuf. Metals-fabricated	12	4	39	0	5	3	47	0	1
C26	Manuf. Electronic	10	7	38	0	14	10	31	0	1
C27	Manuf. Electric	7	6	42	0	13	9	33	0	1
C28	Manuf. Machinery	10	12	32	2	5	1	49	0	1
C29	Manuf. Vehicles	4	5	46	0	14	10	31	0	1
C30	Manuf. Transport-other	2	6	47	0	12	10	33	0	1
C31-C32	Manuf. Furniture	1	1	53	0	8	4	43	0	1
C33	Repair-Installation	17	9	29	0	8	2	45	0	1
D35	Electricity-Gas	31	3	21	0	10	5	40	0	1
E36	Water	19	3	33	0	4	5	46	0	1
E37-E39	Sewage	18	3	34	0	6	8	41	0	1
F	Construction	5	3	47	0	14	9	32	0	1
G45	Vehicle trade	2	5	48	0	9	7	39	0	1
G46	Wholesale	19	3	33	0	4	25	26	0	1
G47	Retail	2	3	50	0	6	10	39	0	1
H49	Land transport	28	3	24	0	11	2	42	0	1
H50	Water transport	9	8	38	0	8	5	42	0	1
H51	Air transport	5	7	43	0	10	6	39	0	1
H52	Warehousing	12	9	34	0	9	7	39	0	1
H53	Postal	6	7	41	2	3	5	47	0	1
I	Accommodation-Food	5	3	47	0	7	6	42	0	1
J58	Publishing	1	2	52	0	10	14	31	0	1
J59-J60	Video-Sound-Broadcasting	2	2	51	0	9	5	37	7	1
J61	Telecommunications	26	11	18	0	7	5	42	2	1
J62-J63	IT	16	13	26	0	7	6	42	0	1
K64	Finance	10	19	26	0	6	3	46	0	1
K65	Insurance	6	12	36	2	6	3	46	0	1
K66	Auxil. Finance-Insurance	5	7	41	4	6	4	45	0	1
L68	Real estate	1	3	51	0	7	5	43	0	1
M69-M70	Legal	12	15	28	0	5	3	46	2	1
M71	Architecture-Engineering	6	10	39	0	4	2	49	0	1
M72	R&D	1	2	52	0	4	3	48	0	1
M73	Advertising	1	7	47	0	5	2	48	0	1
M74-M75	Other Science	1	8	46	0	4	1	50	0	1
N	Private Administration	16	16	23	0	3	2	50	0	1
O84	Public Administration	6	3	45	2	5	2	48	0	1
P85	Education	1	4	50	0	6	8	41	0	1
Q	Health	1	6	48	0	7	7	41	0	1
R_S	Other Service	1	1	50	5	4	0	51	0	1
T	Household activities	0	0	54	2	0	0	55	0	0

Table 9: Summary table of critical input ratings by IHS Markit analysts. Columns below *Input-based rankings* show how often an industry has been rated as critical (score=1), half-critical (score=0.5) or non-critical (score=0) input for other industries, or how often the input was rates as NA. Columns under *Industry-based rankings* give how often an input has been rated as with 1, 0.5, 0 or NA for any given industry. Column *n* indicates the number of analysts who have rated the inputs of any given industry. Industry *T* uses no inputs and is therefore not rated.

D Sensitivity analysis

In this appendix we perform sensitivity analysis of the economic model with respect to both supply and demand shocks (Appendices D.1 and D.2) and model parameters and assumptions (Appendices D.3 and D.4). For the latter, we follow a one-at-a-time sensitivity analysis approach (Borgonovo & Plischke 2016), in the sense that we start from the baseline scenario described in the main text and vary some assumptions while holding all other assumptions fixed to the baseline scenario (see Table 10). Further, in Appendix D.5 we show how the various scenarios compare in terms of matching sectoral unemployment data from the U.S. states of Washington and Texas, see Section 4. Finally, in Appendix D.6, we compare our model results to those of traditional input-output models, namely the Leontief and Gosh models.

Scenario names	τ	production function	hiring-firing	γ_H	γ_F	cons. function
baseline	$\tau = 10$	Leontief, important inputs non-critical	yes	0.03	0.07	muellbauer
$\tau = 1$	$\tau = 1$	Leontief, important inputs non-critical	yes	0.03	0.07	muellbauer
$\tau = 5$	$\tau = 5$	Leontief, important inputs non-critical	yes	0.03	0.07	muellbauer
$\tau = 15$	$\tau = 15$	Leontief, important inputs non-critical	yes	0.03	0.07	muellbauer
$\tau = 20$	$\tau = 20$	Leontief, important inputs non-critical	yes	0.03	0.07	muellbauer
linear	$\tau = 10$	linear	yes	0.03	0.07	muellbauer
leontief	$\tau = 10$	Leontief	yes	0.03	0.07	muellbauer
important inputs critical	$\tau = 10$	Leontief, important inputs critical	yes	0.03	0.07	muellbauer
important inputs half-critical	$\tau = 10$	Leontief, important inputs half-critical	yes	0.03	0.07	muellbauer
no hiring-firing	$\tau = 10$	Leontief, important inputs non-critical	no	0.03	0.07	muellbauer
fast labor adjustment	$\tau = 10$	Leontief, important inputs non-critical	yes	0.50	1.00	muellbauer
slow labor adjustment	$\tau = 10$	Leontief, important inputs non-critical	yes	0.01	0.02	muellbauer
fixed consumption	$\tau = 10$	Leontief, important inputs non-critical	yes	0.03	0.07	fixed
keynesian consumption	$\tau = 10$	Leontief, important inputs non-critical	yes	0.03	0.07	keynesian

Table 10: Scenarios for model parameters and assumptions considered in this sensitivity analysis. See Sections Sections D.3 and D.4 for a more detailed description of the various specifications.

D.1 First-order shocks uncertainty

Since there is substantial uncertainties in first-order shocks discussed in Section 3.5, we test how sensitive model results are with respect to the shock initialisation considered in the main text. To do this, we first randomly perturb the supply and demand shocks for every industry. More specifically, we create perturbed supply and demand shock vectors by letting

$$\bar{\epsilon}_{i,0}^S = \epsilon_{i,0}^S(1 + \psi_i^S), \tag{47}$$

and

$$\bar{\epsilon}_{i,0}^D = \epsilon_{i,0}^D(1 + \psi_i^D), \tag{48}$$

where $\psi_i^S, \psi_i^D \sim N(0, \sigma)$. We use different values for standard deviation, $\sigma \in \{0.01, 0.1, 0.2\}$, representing a normal randomization of original values by 1-20% standard deviations. We then initialise the model with the perturbed first-order shocks and run the lockdown simulations. We repeat this procedure 1,000 times and report median values, interquartile range (IQR) and the 95% confidence bounds of aggregate output values. We did not investigate perturbing other final demand $f_{i,0}^d$.

The upper left panel of Fig. 17 presents the result of this analysis. Since results are qualitatively very similar for the explored standard deviation specifications, we only show the largest perturbation case with $\sigma = 0.2$. First note that the default model result (red line) follows very closely the median result (black line). Also, the IQR is only a narrow band around the reported default values. These results are reassuring as they indicate strong robustness of the model result against uncertainty in initial shock values for a large range of simulations. This picture changes when considering the 95% confidence bounds instead. Here, the ribbon expands dramatically towards small values after around 110 time steps. This finding suggests that for a certain range of initial shocks our model would predict a substantial collapse of the economy.

Since we do not observe similar nonlinearities for the IQR, this initial shock arrangement is not particularly likely given our estimates represent reasonable expected values of the “true” shocks. Also, the economic collapse happens only after almost four months of lockdown, a much longer time horizon as considered in the simulations for the main results. Nevertheless, the results emphasizes the importance of nonlinearities in the economic system by demonstrating how related initial economic shocks can be amplified in very different ways.

The upper right and lower panels of Fig. 17 show the same simulations but using exclusively perturbations on the supply and demand side, respectively. It is immediately evident that the large confidence bounds after four months of lockdown are driven by the supply side shock uncertainty. When perturbing only demand shocks and setting supply shocks to the default values (lower panel), there is very little variance in our model prediction.

D.2 Economic impact of supply and demand shocks

We repeat the analysis in Section 5.2 of running model simulations with only parts of the initial shocks being switched on for alternative production function specifications. In the left and right panels of Fig. 16 we show simulation results for Leontief and linear production functions, respectively. We find for all production functions that supply shocks are substantially more severe than demand shocks, in particular for Leontief production.

For the Leontief production model, economic impacts on gross output are almost identical for the supply-shocks-only and baseline scenarios. There is a slightly less realized consumption when having only supply shocks present compared to both supply and demand shocks being switched on.

In the case of linear production functions there is a clearer ordering of how severe demand, supply and both shocks together impact overall economic performance. Here, impacts on gross output are smaller if only supply shocks are considered compared to the baseline case where both demand and supply shocks are switched on. This makes sense since there are no input bottlenecks in this case, making higher competition for given production levels less problematic. Nevertheless, realized final consumption is also smaller for the linear production model if only supply shocks are considered.

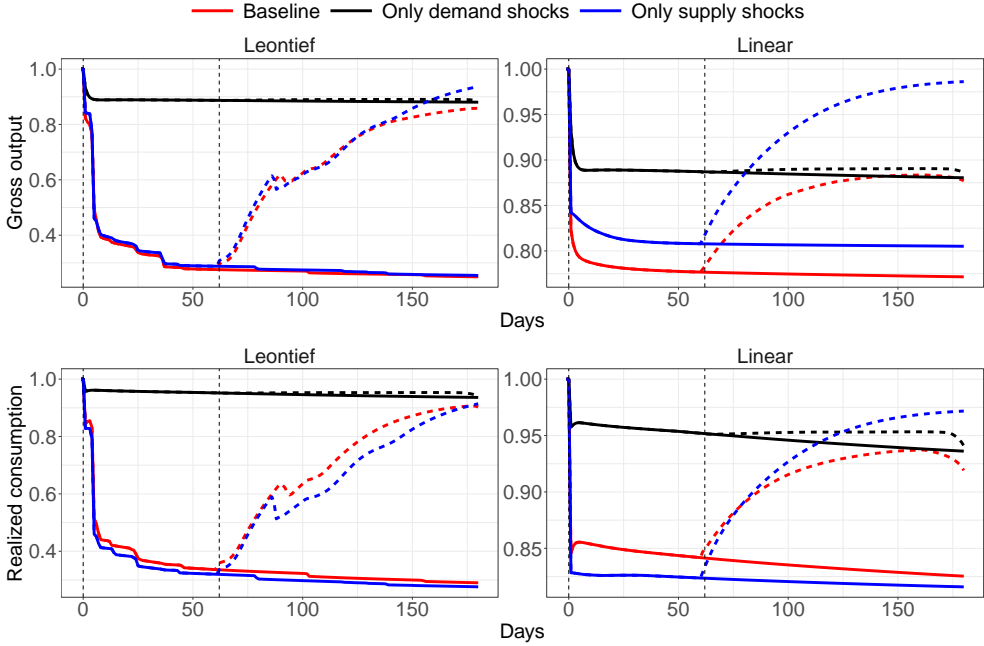


Figure 16: **Dynamic effect of supply shocks vs. demand shocks for different production functions.** As Fig. 8, but using Leontief (left panels) and linear (right panels) production functions. Normalized values of gross output (upper panels) and realized consumption (lower panels) for different shock scenarios. Baseline (red) denotes the model default setup where both supply and demand shocks are used. The blue/black line shows the case where only demand/supply shocks are switched on. The lockdown starts at $t = 0$ and ends for all industries after two months at $t = 60$.

D.3 Production function

We re-run the same simulations as in Section 5.1, but now open all industries after two months of lockdown to also compare recovery paths between different production function specifications. Fig. 18 shows the results of these simulations where the lockdown ends at $t = 62$ (vertical dashed line). We find that after six months the five recovery paths converge for different production function specifications, although the transient looks very different for an extended period of time. Note that the economy does not fully recover after six months due to the slow rebounding of pessimistic consumer expectations consumers and persistence of shocks in exports and investments (see Section 3.5).

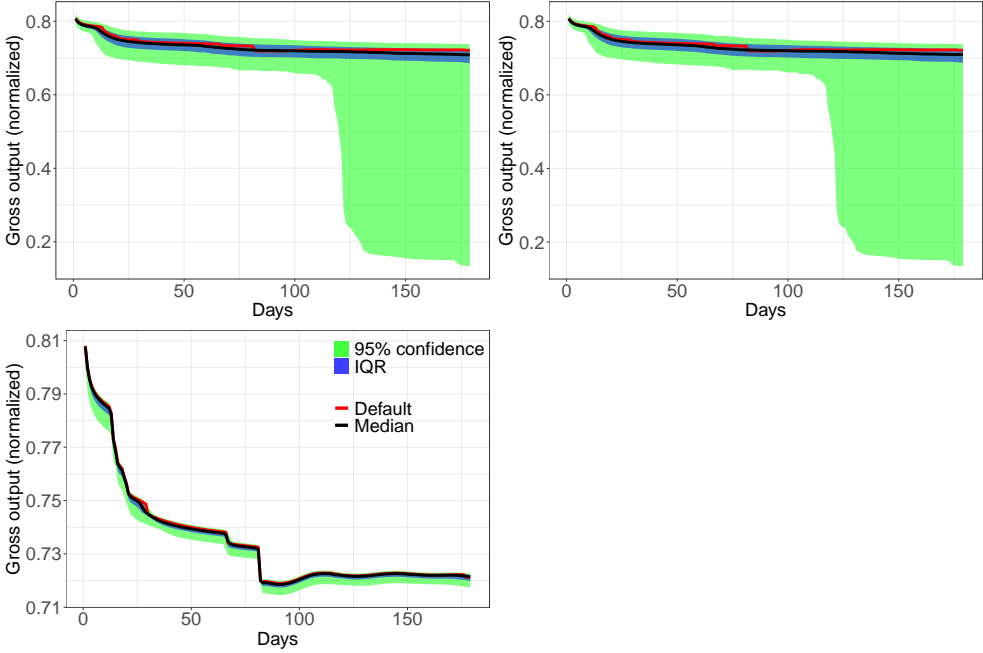


Figure 17: **Impact of uncertainty in first-order shocks.** *Upper left panel:* Perturbing initial supply and demand shocks. *Upper right panel:* Perturbing only supply shocks. *Lower panel:* Perturbing only demand shocks. The red line is the default model run reported in the main text, the black line the median of all 1,000 model runs with perturbed initial shock vectors. Green indicates the 95% quantile and blue the interquartile range.




Figure 18: **Recovery from the lockdown for four different production functions.** The same as Fig. 6, but here the lockdown ends at $t = 62$.

Covid Economics 23, 28 May 2020: 79-151

D.4 Sensitivity analysis of model parameters

To better understand how results are affected by particular model parameter choices, we conduct a series of sensitivity tests. We make ‘local’ sensitivity tests, meaning that we take the default model setup and then vary a set of parameter to investigate how simulation results are affected.

We first present sensitivity tests on inventory adjustment parameter τ which plays an important role in intermediate demand; Eq. (8). Note that a small τ represents quick adjustment behavior where firms aim to replenish run-down inventories essentially within a day. On the other hand, if τ is large, firms react slowly to changes in their input inventories, even when at risk of facing input bottlenecks.

We see in Fig. 19 how aggregate economic outcomes depend on parameter τ . We find that small values of τ , representing highly responsive firms, dampen adverse economic impacts, while negative impacts are larger if we assume higher sluggishness.

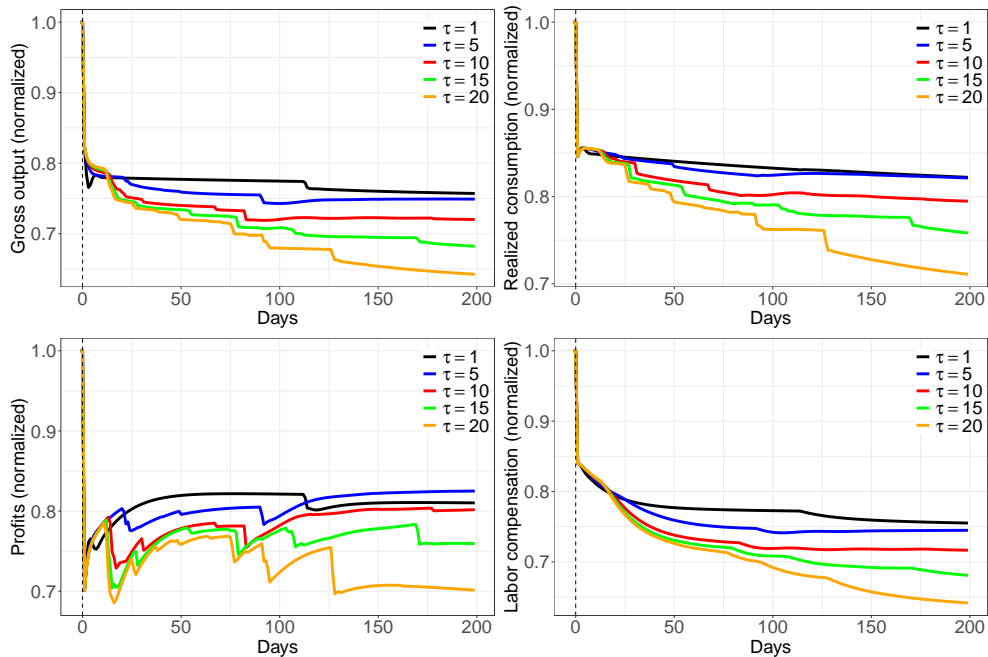


Figure 19: Model results for different choices of the inventory adjustment speed parameter τ . The baseline case is $\tau = 10$.

We also make sensitivity tests with respect to different consumption functions. We test following specifications. First, we use the default consumption function inspired by [Muellbauer \(2020\)](#) which is discussed in detail in Section 3.3. As alternative we also consider a simpler consumption function where consumers demand simply a fixed portion of their current income (i.e. have a fixed marginal propensity to consume) which for brevity we call ‘Keynesian’ consumption function. As an even simpler specification we also consider a fixed consumption function where consumers demand a fixed portion of their initial income. For the two alternative consumption functions we choose marginal propensities to consume equal to one such that all of present or initial income is consumed.

Covid Economics 23, 28 May 2020: 79-151

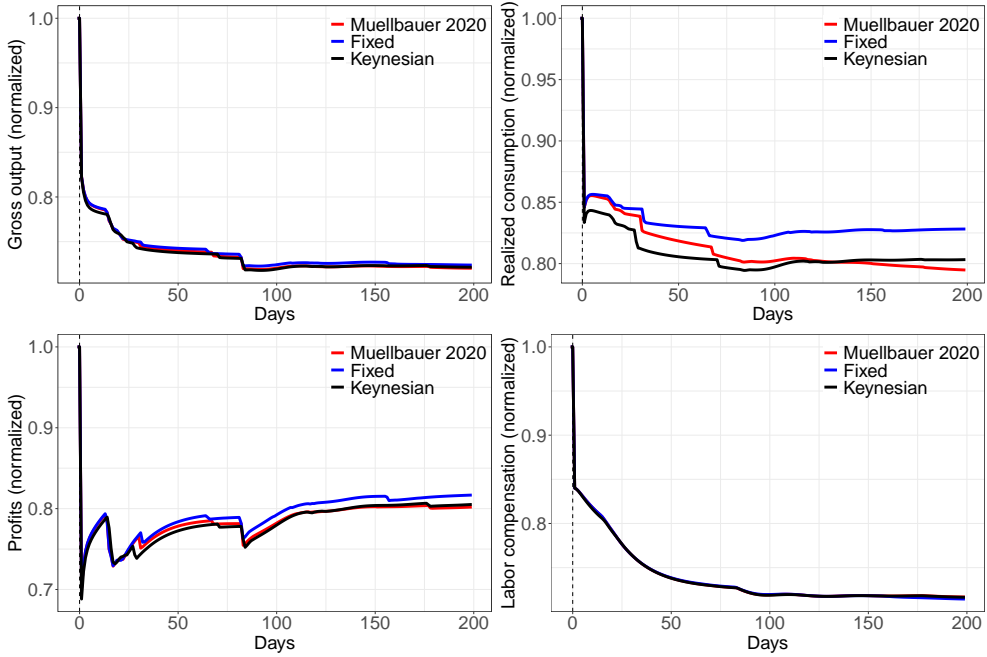


Figure 20: Model results for different consumption functions. The default case is “Muellbauer 2020”.




Figure 21: The effect of hiring and firing speed, γ_H and γ_F , on model results. The default case is $\gamma_H = 1/30$. In all simulations we used $\gamma_F = 2\gamma_H$.

Model results for alternative consumption function specifications are shown in Fig. 20. There are only negligible differences between different production functions on gross output and labor compensation. Realized consumption is slightly higher for a fixed consumption function which is not surprising and somewhat artificially achieved since here consumers demand based on comparatively large initial income values.

We also investigate how model results depend on the speed of adjustment in labor inputs. In Section 3.3 we introduced a parameter γ_H which represents how quickly firms can hire employees in case they want to ramp up their productive capacities. Values of γ_H close to one represent the case where hiring can happen very quickly, whereas values close to zero indicate that it is very hard for firms to hire new workers. Similarly, we considered an equivalent parameter γ_F for firing workers.

In Fig. 21 we show how model results are affected if different γ_H values are used as well as if hiring and firing are completely ruled out. All these simulations use $\gamma_F = 2\gamma_H$ to reflect the situation where firing of employees takes less time than hiring if allowed. We find almost no differences on gross output and realized consumption for all these cases. In line with intuition the exact specification of hiring and firing affects labor compensation and firms' profits. In case of no hiring and firing, labor compensation remains constant throughout the simulation, once the initial labor supply shock is applied. Labor compensation is smaller the easier it is for firms to fire (and hire) employees. This makes sense since firms which face production constraints other than capacity constraints will lay off employees, reducing overall labor income. The picture is reversed for profits. If there is no flexibility for firms in adjusting labor input, there is a larger negative impact on profits. The easier firms can lay off workers, the more they reduce costs on labor which they do not need to satisfy aggregate demand.

D.5 Sensitivity of comparison to empirical data

In Section 4, we compared model predictions to data coming from the U.S. states of Washington and Texas, when running the model in the baseline scenario. In this section, we consider the other scenarios outlined in Table 10 and described in the previous appendices. We do not find much difference in terms of the relative performance of each scenario when either comparing to Washington or Texas, or using the Pearson or weighted correlation coefficients. We thus report in Fig. 22 only results for Washington, using a weighted correlation coefficient to compare model predictions and empirical data.




Figure 22: For various scenarios, we show the weighted correlation coefficient between model predictions for rises in unemployment in various sectors and empirical data coming from Washington State. See Table 10 for a definition of the various scenarios.

It is immediately apparent that the performance of the various scenarios is similar, except for the cases of the basic Leontief production function and of the Leontief production function with important inputs considered as critical or half-critical. In these cases, and especially in the Leontief case, performance is substantially lower, suggesting that our modeling choice of distinguishing between critical and non-critical inputs adds realism to our model. Correlation between model predictions and empirical data is somewhat lower in case no hiring or firing takes place (in the sense that workers are only furloughed due to the epidemic shock and not due to second-order effects), but this is a clearly unrealistic assumption. Given the combined uncertainties of comparing the model to the data, and the intrinsic uncertainty in these preliminary data, it would not be wise to select an unrealistic assumption based on a small increase in empirical performance.

Therefore, our choice of the baseline reflects a balance between ability to reproduce empirical patterns and prior belief in certain assumptions/parameter values. We use poor empirical performance to exclude the Leontief, “important inputs critical” and “important inputs half-critical” scenarios. We use instead our best judgement to exclude too fast or slow adjustments of inventories and labor force, full substitution of inputs in the linear production function, and too simple consumption functions such as the fixed or “Keynesian” ones. As shown in Appendices D.3 and D.4, in any case, model results tend to depend weakly on these specific assumptions, consistently with the little ability of data to distinguish between the respective scenarios.

D.6 Comparison to traditional IO models

We also compare our model results to traditional input-output (IO) models. In particular, we compare the steady state of our model with two models, the demand-driven Leontief (Leontief 1936) and the supply-driven Gosh model (Ghosh 1958). Since these simpler IO models do not include inventory effects, we set input inventories artificially high such that they do not restrict economic production.

In the Leontief model final demand is exogenous, and under the assumption of fixed production recipes, gross output per industry is endogenously determined by multiplying demand with the Leontief inverse (Miller & Blair 2009). When considering only demand shocks, we can

write the Leontief prediction as

$$x^L = (\mathbb{I} - A)^{-1}(c^{\text{shocked}} + f^{\text{shocked}}). \quad (49)$$

We also rerun our model with all supply shocks being switched off and only considering demand shocks. We then compare the steady state results of our model with the Leontief prediction.

Fig. 23 (left panel) shows the reduction of sectoral gross output compared to the pre-shock state as barplots for our and the Leontief model. We find that our model very closely recovers the Leontief prediction in the steady state. Gross output per industry in the steady state of our model and the Leontief model have almost a correlation of one. The differences between predicted sectoral reductions in gross output are almost zero in all cases. Only for Health (Q) they differ by 2.3%, since the Leontief model would predict an increase as a result of positive demand shocks which cannot be satisfied in our model due to fixed maximum capacity constraints.

These results are very robust against using empirical inventories. It should be noted that the Leontief model is static and we are comparing the steady state of our dynamic model. Thus, modeling the transient which is relevant for the short time-scales considered in the main text is not possible with the traditional Leontief model.

We do a similar comparison with the supply-driven Gosh model. There are no fixed production recipes in the Gosh model, but fixed “allocation coefficients” $B_{ij} = Z_{ij,0}/x_{i,0}$. Here, a change in gross output is due to a change in primary inputs, i.e. represented as value added. In the notation used here we can formulate the Gosh prediction as

$$x^G = (\mathbb{I} - B^\top)^{-1}(l^{\text{shocked}} + e). \quad (50)$$

We plot the Gosh predictions and the steady state results of our model with only supply shocks turned on in the right panel of Fig. 23. We find greater differences between the Gosh and our model for the supply shocks. This should not come as a surprise, since the Gosh model builds upon a very different production function.

Rankings of sectoral declines are still very correlated (Spearman correlation of 0.91). This is higher than the correlations between our model’s steady state rankings of industries and the initial shock rankings (correlation of 0.87). Unsurprisingly, the Gosh model rankings are most similar as the ones obtained from using initial supply shocks only (correlation of 0.94).

These results are not very robust with respect to the specifications of the economic model considered here. Using empirical inventories in our model enlarges differences in model predictions tremendously.

We compared our model also to slightly more complex mixed endogenous/exogenous IO models (Dietzenbacher & Miller 2015, Arto et al. 2015) which simultaneously can take supply and demand shocks into account. Yet these models do not always guarantee positive solutions for variables such gross output and final consumption (Miller & Blair 2009, p.628). In particular when applying the large first-order supply and demand shocks of the pandemic to the UK economy, the mixed IO model does not yield feasible allocations.




Figure 23: Comparison between impacts on sectoral gross output obtained from traditional IO models with our model. The left panel compares the model steady state with the Leontief (only demand shocks are switched on), the right panel with the Gosh model (only supply shocks are switched on). Note that our model was run with artificially high inventories, no hiring and firing, $\Delta s = 1$ and no income losses ($b = 1$).

E Epidemic modelling

In this appendix we present our epidemic model where we divide contagion channels by activities. As we focus on the early stage of the epidemic, we do not explicitly model the number of recovered individuals R , although that plays a role to determine the total population M . We start denoting the number of susceptible and infected people in the pupils and students and in the non-working adults category by S^s, I^s, S^u, I^u respectively. Similarly, S_i and I_i denote the number of susceptible and infected workers of industry i . It follows that the decrease in the overall susceptible population S is given by

$$\frac{dS}{dt} = \frac{dS^s}{dt} + \frac{dS^u}{dt} + \sum_{i=1}^N \frac{dS_i}{dt}. \tag{51}$$

In what follows, we compute the rate of infection of each population category by focusing on the different channels of contagion each person is exposed to. In these computations we assume homogeneous mixing of the population, meaning that the probability that a person had contact with an individual that was infected is $\frac{I}{M}$, regardless on the channel they had contact in.

Normalizing contact-weighted shares by population As we discuss in Appendix F we have data on the share of intensity-weighted contacts in each activity of the overall population. For the derivation of the epidemiological model it is useful to renormalize these shares of weighted contacts (i.e. the β 's) by the population they come from. This is not necessary for $\beta_c(0)$ or $\beta_h(0)$, since consumption and other household interaction related contacts are spread evenly across the whole population. On the contrary, we do need to renormalize $\beta_s(0)$ by the

student and pupil population η_s , so that

$$\hat{\beta}_s(0) = \frac{\beta_s(0)}{\eta_s}. \tag{52}$$

Another way to look at the equation above is to note that $\hat{\beta}_s(0)$ is the share of intensity-weighted contacts in school *per unit* population, and to obtain the actual share of intensity-weighted contacts $\beta_s(0)$ one needs to multiply $\hat{\beta}_s(0)$ by the population share of students, η_s . Similarly, we renormalize the work intensity-contacts across the workers of different industries as follows

$$\hat{\beta}_w(0) = \frac{\beta_w(0)}{\sum_{i=1}^N \eta_i b_{i,w}}, \tag{53}$$

where the normalization includes the $b_{i,w}$ factors i.e., the heterogenous distribution of intensity-weighted contacts across industries.

In the transport channel, we must distribute the contacts across the commuter population (i.e. workers and students). To account for a density effect (see below), we assume that the number of contacts scales with the square of the number of people in public transport, and use the normalization factor

$$\hat{\beta}_T(0) = \frac{\beta_T(0)}{\left(\eta^s + \sum_{i=1}^N \eta_i\right)^2}, \tag{54}$$

where $\eta^s + \sum_{i=1}^N \eta_i$ is the pre-lockdown share of the population that commutes.

Students and pupils To simplify notation we define μ^s as the fraction of the students and pupils population attending schools, which is given by

$$\mu^s = \left(\delta_s + (1 - \delta_s) \left(g \sum_{j=1}^N \delta_{i,w} \eta_i \right) \right).$$

We know that students and pupils are exposed to infection due to school attendance, transport, consumption, and other household interaction. We assume that pupils that go to school have the same amount of contacts in school as before lockdown²⁵, while for transport we consider that the number of contacts decreases due to the reduced density of people in the bus/train. With these assumptions we obtain the following equation for the infection rate

$$\begin{aligned} \frac{dS^s}{dt} = & -\beta^* \left[\hat{\beta}_s(0) \mu^s S^s \frac{I}{M} + \hat{\beta}_T(0) \mu^s S^s \left(\mu^s \frac{I^s}{M} + \sum_{k=1}^N \delta_{k,w} \frac{I_k}{M} \right) \right. \\ & \left. + \beta_c(0) S^s \frac{I}{M} \sum_{k=1}^N \delta_k(t) b_{k,c} + \beta_h(0) S^s \frac{I}{M} ((1 - \delta_h) \kappa + \delta_h) \right], \end{aligned} \tag{55}$$

where the first two terms correspond to the infections happening at school and transport and thus only apply to the fraction μ^s of the student population that goes to school. The third and fourth term correspond to infections happening while consuming or doing other household activities and therefore apply to the whole student population. $b_{k,c}$ is the consumption related contacts, while κ is the share of social/family/friends contacts that are not avoidable by social

²⁵We make this assumption considering that a) only certain schools are open so it is unclear to what extent the density in schools has decreased b) it is possible that pupils interact more with the few pupils left in school and thus the number of contacts can remain roughly constant.

distancing. Notice that since transport is shared with both students and workers, the transport term includes both I^s and I_k . The fact that we are considering the fractions μ^s and $\delta_{k,w}$ of infected in the transport term reflects our assumption that density matters in particular in transports; note, for example, that we are not multiplying infected individuals by μ^s in the school term.

We simplify the above equation using the mean field approximation $S_i \approx \eta_i S$, $I_i \approx \eta_i I$, $S^s \approx \eta^s S$, and $I^s \approx \eta^s I$ and obtain

$$\begin{aligned} \frac{dS^s}{dt} = & -\frac{\beta^*}{M}SI \left[\hat{\beta}_s(0)\mu^s\eta^s + \hat{\beta}_T(0) \left((\eta^s\mu^s)^2 + \eta^s\mu^s \sum_{k=1}^N \delta_{k,w}\eta_k \right) \right. \\ & \left. + \beta_c(0)\eta^s \sum_{k=1}^N \delta_k(t)b_{k,c} + \beta_h(0)\eta^s ((1 - \delta_h)\kappa + \delta_h) \right]. \end{aligned} \tag{56}$$

Working population Workers are exposed to infection due to work, transport, consumption, and other household interaction. For a worker in industry i , the infection rate is

$$\begin{aligned} \frac{dS_i}{dt} = & -\beta^* \left[\hat{\beta}_w(0)\delta_{i,w}S_i b_{i,w} \frac{I}{M} + \hat{\beta}_T(0)S_i\delta_{i,w} \left(\mu^s \frac{I^s}{M} + \sum_{k=1}^N \delta_{k,w} \frac{I_k}{M} \right) \right. \\ & \left. + \beta_c(0)S_i \frac{I}{M} \sum_{k=1}^N \delta_k(t)b_{k,c} + \beta_h(0)S_i \frac{I}{M} ((1 - \delta_h)\kappa + \delta_h) \right], \end{aligned} \tag{57}$$

where we have assumed that workers that go to work make the same amount of contacts at work as before lockdown, while for transport we consider that the number of contacts decreases due to the reduced density of people in public transport. We have made explicit that the work and transport infection channels only apply to the fraction $\delta_{i,w}$ of the working population in i going to work and to the μ^s fraction of students going to school. As before, we use the mean field approximation $S_i \approx \eta_i S$, $I_i \approx \eta_i I$, $S^s \approx \eta^s S$, and $I^s \approx \eta^s I$ to simplify the equation to

$$\begin{aligned} \frac{dS_i}{dt} = & -\frac{\beta^*}{M}SI \left[\hat{\beta}_w(0)\delta_{i,w}\eta_i b_{i,w} + \hat{\beta}_T(0)\eta_i\delta_{i,w} \left(\mu^s\eta^s + \sum_{k=1}^N \delta_{k,w}\eta_k \right) \right. \\ & \left. + \beta_c(0)\eta_i \sum_{k=1}^N \delta_k(t)b_{k,c} + \beta_h(0)\eta_i ((1 - \delta_h)\kappa + \delta_h) \right]. \end{aligned} \tag{58}$$

We now sum across all N industries to obtain

$$\begin{aligned} \sum_{i=1}^N \frac{dS_i}{dt} = & -\frac{\beta^*}{M}SI \left[\hat{\beta}_w(0) \sum_{i=1}^N \delta_{i,w}\eta_i b_{i,w} + \hat{\beta}_T(0) \left(\mu^s\eta^s \sum_{i=1}^N \delta_{i,w}\eta_i + \left(\sum_{i=1}^N \delta_{i,w}\eta_i \right)^2 \right) \right. \\ & \left. + \beta_c(0) \sum_{i=1}^N \eta_i \sum_{k=1}^N \delta_k(t)b_{k,c} + \beta_h(0) \sum_{i=1}^N \eta_i ((1 - \delta_h)\kappa + \delta_h) \right]. \end{aligned} \tag{59}$$

Non-working adults By definition non-working adults are not exposed to the work or school infection channel. Furthermore, since we only consider work-commuting transport use, the non-working adults are not exposed to the transport infection channel either. It follows that the decrease in the susceptible population depends only on the consumption and other household

interaction channel

$$\frac{dS^u}{dt} = -\frac{\beta^*}{M}SI \left[\beta_c(0)\eta^u \sum_{k=1}^N \delta_{k,c}(t)b_{k,c} + \beta_h(0)\eta^u ((1 - \delta_h)\kappa + \delta_h), \right] \tag{60}$$

where we have again used the approximation $S^u \approx \eta^u S$.

Total population To get the infection rate of the overall population we substitute Eqs. (56)–(60) in equation Eq. (51). It follows that

$$\begin{aligned} \frac{dS}{dt} = & -\beta^* \frac{SI}{M} \left[\hat{\beta}_w(0) \sum_{i=1}^N \delta_{i,w} \eta_i b_{i,w} + \hat{\beta}_s(0) \mu^s \eta^s \right. \\ & + \hat{\beta}_T(0) \left(\mu^s \eta^s + \sum_{i=1}^N \eta_i \delta_{i,w} \right)^2 \\ & \left. + \beta_c(0) \sum_{k=1}^N \delta_k(t) b_{k,c} + \beta_h(0) ((1 - \delta_h)\kappa + \delta_h) \right], \end{aligned} \tag{61}$$

where we have used the fact that $\eta^s + \eta^u + \sum_{i=1}^N \eta_i = 1$. The above terms correspond to the infection due to work, school, transport, consumption, and other household interaction respectively. We can write the above equation as

$$\frac{dS}{dt} = -\beta(t) \frac{SI}{M} \tag{62}$$

where

$$\beta(t) = \beta^* \left(\beta_w(t) + \beta_s(t) + \beta_c(t) + \beta_T(t) + \beta_h(t) \right), \tag{63}$$

which is Eq. (37) of the main text. The β 's are given by

$$\beta_w(t) = \hat{\beta}_w(0) \sum_{i=1}^N \delta_{i,w} \eta_i b_{i,w} = \beta_w(0) \sum_{i=1}^N \delta_{i,w} \frac{\eta_i b_{i,w}}{\sum_{l=1}^N \eta_l b_{l,w}}, \tag{64}$$

$$\beta_s(t) = \hat{\beta}_s(0) \mu^s \eta^s = \beta_s(0) \mu^s, \tag{65}$$

$$\beta_T(t) = \hat{\beta}_T(0) \left(\mu^s \eta^s + \sum_{i=1}^N \eta_i \delta_{i,w} \right)^2 = \beta_T(0) \left(\frac{\mu^s \eta^s + \sum_{i=1}^N \eta_i \delta_{i,w}}{\eta^s + \sum_{l=1}^N \eta_l} \right)^2, \tag{66}$$

$$\beta_c(t) = \beta_c(0) \sum_{k=1}^N \delta_k(t) b_{k,c}, \tag{67}$$

and

$$\beta_h = \beta_h(0) ((1 - \delta_h)\kappa + \delta_h). \tag{68}$$

F Calibration of epidemic model

F.1 Literature review

In epidemiology, the main method to understand contact patterns is to use social contact surveys. A landmark study is the Polymod study (Mossong et al. 2008). Several other studies in the last decade have confirmed that, roughly speaking, people have about 10-20 non-casual contacts per day, mostly at home and at work. By “non-casual” contact, we mean contacts as defined by these studies, that is, either physical contact or non-physical contact defined as “a two-way conversation with three or more words in the physical presence of another person”.

The Polymod study is very interesting for us because it reports *where* contacts take place. Averaging across countries and pooling physical and non-physical contacts, 23%, 21%, 14%, 3%, and 16% are made at home, at work, at school, while travelling, and during leisure activities, respectively.

There are still significant uncertainties on the mode of transmission of SARS-CoV-2, and in particular whether it can diffuse through casual contact (whether simply ‘sharing air’ is risky, because aerosolized particles, rather than just droplets, are risky). Fortunately, there have also been a couple of studies quantifying “casual” contacts, that is, contacts between anonymous individuals but that nevertheless involve enough proximity to result in a transmission. Mikolajczyk & Kretzschmar (2008) report several studies where participants (students) were “asked about aggregate numbers of contacts on six levels of proximity: intimate contacts, close contacts (same household), direct conversation (> 2 min duration, max. 2 m distance), small group (with conversations, but less intensive than in direct conversations), larger group (seminary or lecture room) and occasional contacts (in the range of 2 m in local transportation, cinema, etc.)” The number of conversational contacts (i.e. including intimate, close and direct conversation) was sensibly below but in line with the Polymod study (6-13 contacts vs 10-20). Their Fig. 2 suggests that the number of contacts in small groups, large groups, and random contacts was roughly speaking 8, 30 and 40, with variations depending on study design. Roughly speaking, people have on average 10 close contacts per day but 80 casual (non-close) contacts.

A few studies have looked at social contact patterns to understand the diffusion of tuberculosis, which spreads very easily. Reading from their Fig. 3, the estimates of mean casual contacts per day obtained by McCreesh et al. (2019) for South Africa are about 10 for transport (combining trains and large taxis), 6 for school and work, 5 for shops (spaza shop, shebeen and mall), 2 for home, and less than 2 for church and community hall. These estimates are for the mean number of contacts per day, but McCreesh et al. (2019) also asked participants for the number of casual contact present during the visit to a location (Fig. S2), showing about 40 casual contacts in Malls and Trains. In many other categories relating to transport or shopping, the number of casual contacts is around 10-15.

In this paper, we use a study from Sweden (Strömgren et al. 2017). The study reports, for a variety of places, the likelihood that it is visited during an average day, the duration of the visit²⁶, the number of people present, and the likelihood of physical contact, see Table 11.

Category	Stromgren et al's category	<i>Visit</i>	<i>Duration</i>	<i>Crowd</i>	<i>Physical</i>	<i>Intensity</i>	
		Visit likelihood	Duration (hours)	Median number people present	Likelihood physical contact	Weighted contacts	
Work	Work	21.20	7.60	20.00	55.80	0.29	0.29
School	Pre-school	8.60	7.60	20.00	73.30	0.12	0.28
	School	12.00	7.60	20.00	71.10	0.16	
Consume	Convenience store	5.20	0.40	10.00	8.30	0.00	0.16
	Large store	24.10	0.80	21.50	18.00	0.04	
	Restaurant	9.40	1.40	30.00	30.80	0.04	
	Sports venue	11.50	2.30	34.50	53.80	0.08	
Transport	Public transport	16.30	1.00	40.00	8.30	0.06	0.06
Home	Home	95.00	18.40	1.00	73.70	0.16	0.21
	Car	58.70	0.90	1.00	25.80	0.00	
	Public urban space	6.60	1.80	20.00	28.30	0.02	
	Friends and relatives	21.00	5.10	3.00	80.10	0.03	

Table 11: The columns *Duration* and *Crowd* for the rows School and Pre-school are inferred from the equivalent number in the row Work. “Large store” is short for “Large and specialist store”. The source of raw data is Strömrgren et al. (2017). Intensity-weighted contacts are our own calculations, see text. The last column shows the values calculated for Table 3, Eqs (37)-(38) in the main text.

F.2 Calibration

We used the data from Table 11 to create an intensity-weighted number of contacts. We define for each of the 12 places²⁷.

$$Intensity_i = \frac{Visit_i \times Duration_i \times Crowd_i}{\sum_{j=1}^{12} [Visit_j \times Duration_j \times Crowd_j]} \tag{69}$$

To compute the values in Table 3 for Eqs. (37)-(38), we sum-up the relevant *Intensity* variables.

Work. To calibrate b_{iw} , we create an index based on the *physical proximity* and *exposure to infection* index of each industry, which, as explained in Appendix A, we map from O’NET data. At the occupation level, physical proximity and exposure to infection range from 0 to 100 and are described as follows.

- *Exposure to disease and infection.* O*NET assigns a score to each occupation depending on the frequency with which workers in that occupation are exposed to disease and infection *in normal times*. The scale runs from 0, indicating that the worker is never exposed to 100, indicating that the worker is exposed every day. It is important to consider that this rating was done before the pandemic, and doesn’t seem to properly take into account the properties of COVID-19

²⁶The duration of shop visits is highly consistent with the data reported by Goldfarb & Tucker (2020), who use mobile phone data for the US and report an average visit of 22 to 42 minutes across 11 categories of retail shops.

²⁷Note that we could have used the variable showing the likelihood of physical contact as proxy for the closeness of contact, as an additional factor in Eq. (69). We have done so in a robustness check and most results are similar, except for Sports Venue which becomes an even larger share of all consumption risks. We decided against using this additional variable in the current draft as we match this activity with the industry that contains cinemas, theatres, religious gatherings, etc. It is true that, like sports, these activities have a significant duration, but they are not as likely to involve physical contact.




Figure 24: **COVID-19 relative infection risk of industries.** This index is constructed by taking the average of the exposure to infection and the physical proximity index of industries and then normalizing so that they sum up to one.

- *Physical proximity.* O*NET also considers to what extent performing job tasks requires physical proximity. A score of 75 implies being moderately close (at arm’s length) and 100 implies near touching.

To obtain a score at the industry level, we aggregate occupation-level scores using employment data from the BLS, which indicates the occupational composition of each industry and then map into the WIOD classification (see Appendix A for details). Our industry-specific infection risk is the average of physical proximity and exposure to infection. That is

$$b_{iw} = \frac{1}{2}(\text{exposure to infection}_i + \text{physical proximity}_i). \tag{70}$$

Consumption. We consider that, from Table 11, there are three types of consumption activities: Shopping (Convenience stores and Large stores), Restaurants, and Sports Venue. We then map these into the WIOD but looking at the list of industries (Table 5) and assuming that all Shopping activity comes from the WIOD industry *G47: Retail*; that all Restaurant activity comes from the Industry *I:Accommodation-Food*; and all Sports Venue activity comes from the *Other Services* activity.

Transport. We consider the value from Table 11. We note that Strömghren et al. (2017) observe an important divide between rural and urban places in terms of time spent in public transports.

Home-related. In the main text, we need to consider the impact of social distancing measures on β_h , the share of contacts that are unrelated to whether industries are open or not. We assume that during lockdown, the number of contacts at home stay the same, but the number of contacts with Friends and Relatives, within a family car, or in public urban spaces fall to zero. Since Home is responsible for 76% (16/21) of the home related contacts, we take $\beta_h(\text{lockdown}) = \beta_h(t = 0) * 0.7$.

Population data. To obtain the share of population in the special industries Schools and Out-of-the-labor-force, we use the ONS Current Population Survey²⁸ According to these surveys 62% of the population is employed and 23% of the population is between 0-19 years old. Therefore, we assign 62% of the population to the $i = 1, \dots, M - 2$ working industries, 23% to the school industry²⁹ and the rest to the retired industry (unemployed and all inactive are thus assigned into this industry). From the 62% of the working population we assign them to the $i = 1, \dots, M - 2$ working industries according to the shares of employment calculated from the WIOD employment data. Finally, again using the ONS CPS, we compute that the share of 0-19 year old who are 14 or below is $g = 17/23$.

²⁸<https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/datasets/populationestimatesforukenglandandwalesscotlandandnorthernireland>

²⁹We assume that all people between 0-19 years old go to school.

G Notation

Symbol	Name
N	Number of industries
t	Time index
$t_{start_lockdown}$	Start date of lockdown
$t_{end_lockdown}$	End date of lockdown
$t_{end_pandemic}$	End date of pandemic
$x_{i,t}$	Gross output of industry i
$Z_{ij,t}$	Intermediate consumption of good i by industry j
$c_{i,t}$	Household consumption of good i
$c^d_{i,t}$	Demand of household consumption of good i
$f_{i,t}$	Non-household final demand of good i
$f^d_{i,t}$	Demand non-household final demand of good i
$l_{i,t}$	Labor compensation to workers of industry i
$\pi_{i,t}$	Profits of industry i
$e_{i,t}$	"All other" (non intermediates or labor) expenses of industry i
$\bar{l}_t, \bar{c}_t, \bar{\pi}_t, \bar{x}_t$	Total labor compensation, consumption, profits and output
$d_{i,t}$	Total demand for industry i
$O_{ij,t}$	Orders (demand from industry j to industry i)
n_j	Number of days of targeted inventory for industry j
$A_{i,j}$	Payments to i per unit produced of j (technical coefficients)
$S_{ij,t}$	Stock of material i held in j 's inventory
τ	Speed of inventory adjustment
$\theta_{i,t}$	share of goods from industry i in consumption demand
$\bar{\theta}_{i,t}$	share of goods from industry i in consumption demand (unnormalized)
\bar{c}^d_t	Aggregate consumption demand
ρ	speed of adjustment of aggregate consumption
ϵ_t	consumption exogenous shock
\bar{l}^p_t	expectations for permanent labor income
m	share of labor income used to consume final domestic goods
ξ_t	fraction of pre-pandemic labor income that households expect to retain in the long-run
ξ^L_t	fraction of pre-pandemic labor income that households expect to retain in the long-run during the lockdown
$x^{cap}_{i,t}$	industry production capacity based on available labor
$x^{inp}_{i,t}$	industry production capacity based on available inputs
$\tilde{\epsilon}^S_{i,t}$	exogenous supply shock to industry i
$\tilde{\epsilon}^D_{i,t}$	relative changes in demand for goods of industry i during lockdown
$\tilde{\epsilon}_{i,t}$	relative changes in demand for goods of industry i
$\tilde{\epsilon}_t$	aggregate consumption shock
$\Delta l_{i,t}$	desired change of labor supply of industry i
$l^{max}_{i,t}$	maximum labor supply for industry i
γ_H, γ_F	Speed of upward/downward labor adjustment (hiring/firing)
Δs	Change in saving rate
\bar{l}^*_t	Household income including social benefits

Table 12: Notation for the Economic model

Covid Economics 23, 28 May 2020: 79-151

Symbol	Name
S^s, S^u, S_i	Number of Susceptible individuals in the student, adult non-working, and working population
I^s, I^u, I_i	Number of Infected individuals in the student, adult non-working, and working population
$S = S^s + S^u + S_i$	Number of Susceptible individuals
$I = I^s + I^u + I_i$	Number of Infected individuals
R	Number of Recovered individuals
$M = (S + I + R)$	Number of individuals in the population
β^*	Force of infection
γ	Recovery rate
R_0	Reproduction number
η^s, η^u, η_i	Share of people in the student category, the adult non-working category and in industry i
μ^s	Share of the student population that attends school
β_w	Share of intensity-weighted contacts at work
β_s	Share of intensity-weighted contacts in schools
β_c	Share of intensity-weighted contacts in consumption
β_T	Share of intensity-weighted contacts in transports
β_h	Share of intensity-weighted contacts at home

Table 13: Notation for the Epidemic model

Optimal case detection and social distancing policies to suppress Covid-19¹

Stefan Pollinger²

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

This paper shows that the optimal combination of social distancing and case detection allows for complete and efficient eradication of COVID-19. The first contribution is theoretical. I show that the optimal suppression-policy is a simple function of observable sufficient-statistics, making it easily implementable. I prove that optimal social distancing is the strongest when an outbreak is detected, and then gradually relaxed. If case detection is sufficiently efficient, social distancing vanishes wholly and quickly; otherwise, it needs to stay in place until a vaccine arrives. The second contribution is quantitative. I find that, if Italy adopts digital contact tracing, total suppression costs only 0.8% of annual GDP. In sharp contrast, under the current detection efficiency, the total cost of suppression amounts to at least 14% of GDP.

1 I am especially grateful to Christian Hellwig and Nicolas Werquin for their supervision and guidance. I thank Christian Gollier, Paul Seabright, François Salanié and the participants of the COVID workshop at the Toulouse School of Economics for their comments and suggestions. All errors are mine. I plan to update the paper frequently. Please find the latest version here: www.stefanpollinger/research.

2 Ph.D. candidate at the Toulouse School of Economics.

Copyright: Stefan Pollinger

1 Introduction

What is the optimal response to a rapidly spreading and deadly infectious disease, when no vaccine or efficient medication is available? Due to the COVID-19 pandemic, it has become very urgent to answer this policy question. Despite facing the same threat, the policy responses, and therefore the economic and health outcomes, are very heterogeneous among countries. Some, such as Taiwan and Hong Kong, avoided to have an outbreak. Others, such as New Zealand, Australia and China, controlled their outbreaks with heavy social distancing and meticulous case detection using contact tracing. South Korea controlled an initial outbreak even without relying on a lock-down. Some countries, such as Germany, could not avoid applying strict lock-downs to reduce new infections, despite large efforts in testing. Most countries were forced to helplessly impose strict lock-downs, when their health care system started to collapse. Can the mentioned success stories be replicated in other countries? If yes, how? What are the conditions to do so?

This paper shows how to optimally suppress a virus when the policymaker has two tools: social distancing and case detection. Suppression pushes the viral growth rate into negative terrain, such that the virus disappears in the long run. Another possible policy response is mitigation, which I will discuss in the conclusion.¹ Social distancing reduces the growth rate of the virus by reducing the rate of social contacts between all individuals in the population. Case detection, for instance with the help of contact tracing, is the policy of actively finding infectious individuals and isolating them from the susceptible population. The policymaker has to trade-off the cost of suppression measures against the flow of death that results from infections. This trade-off is inherently dynamic because policy measures at a certain point in time affect the flow of death in the future.

The first contribution of the paper is theoretical. I show that suppression is always possible, irrespective of the efficiency of case detection, and I characterize the properties of the optimal policy. Suppose a policymaker discovers an outbreak of the virus. I prove that, in the optimum, she immediately implements social distancing measures to reverse the viral growth. As a consequence, the number of infectious reduces and converges to zero. Importantly, as the number of infectious reduces, the policymaker gradually relaxes the social distancing measures. The optimal response is instantaneous and the largest at the

¹Mitigation controls the spread of the virus until contagions stop because the population achieves herd immunity.

onset. In particular, it is never optimal to "smooth in" social distancing, a mistake made by many countries at the beginning of the COVID-19 pandemic. In the optimum, the number of infections should never rise. If a policymaker observes signs of an increasing number of infections - such as an increase in the flow of confirmed cases, symptomatic patients, hospitalizations, or death - she should immediately strengthen social distancing measures. In particular, any kind of stop and go policy is not optimal.

In the long-run, the optimal degree of social distancing depends crucially on the detection-technology. In particular, what matters is the rate of detection when prevalence is low, i.e., the number of daily detected cases relative to the overall number of currently infectious when the latter is close to zero. Note that, due to decreasing returns to scale, the rate of detection is maximal at zero prevalence. On the one hand, if the rate of detection at zero is larger than the uncontrolled growth rate of the virus, optimal social distancing measures are completely removed in the long run. It means that society is going back to normal, along with the decreasing number of overall infectious in the population. Intuitively, the smaller the number of infectious, the larger is the relative amount of control coming from case detection. In the long run, case detection is efficient enough to control the virus completely. On the other hand, if the rate of detection at zero is smaller than the uncontrolled growth rate of the virus, optimal long-run social distancing is constant and positive.

The long-run behavior of the optimal policy has stark consequences for the total cost of suppression. On the one hand, if social distancing goes to zero, the total cost of suppression is bounded. On the other hand, if the optimal policy is constant in the long run, the total cost is unbounded, even at the optimum, unless another solution for eradication, like a vaccine, is found. The result is due to the fact that the prevalence follows an exponential decay process. In the long run, the reduction in infections becomes infinitely slow. In addition, some degree of social distancing needs to stay in place. As a consequence, the total cost is unbounded. However, despite this result, suppression may still be the optimal policy to follow. A society may prefer keeping some degree of social distancing until a vaccine arrives, instead of taking the deadly path towards herd immunity. The more efficient is the tracing technology, the lower is the necessary amount, and therefore the cost, of social distancing. These results suggest that efficient tracing, at least at low infection levels, has enormous benefits. Its implementation should be a top priority for governments. Note, however, I do not take welfare losses from an eventual loss in privacy into account.

Two simple sufficient statistics characterize the optimal policy at each point in time: first, the instantaneous growth rate of the virus, and second, the instantaneous flow of costs from suppression measures and lives lost. A policy at a certain point in time is optimal if the elasticity of the current growth rate to the current flow cost is equal to one. Note that this property is somewhat surprising. In principle, the optimal policy at a certain point in time depends on the past and the future. However, the two sufficient statistics contain all relevant dynamic information. The condition gives specific and straightforward guidance on how to relax social distancing measures over time, and, in particular, on how to organize a de-confinement. To decide upon relaxing a certain confinement measure, the policymaker only needs to evaluate its relative impact on the current flow of total cost and viral growth rate. If the percentage reduction in cost is larger than the percentage increase in growth, a measure should be relaxed. All results presented so far theoretical. They hold for any set of parameter values.

The second contribution of the paper is quantitative. I calibrate the unknown functions and parameter values in the case of Italy. I assume Italy starts to follow the optimal suppression policy on May 11th. I calculate the optimal policy and its cost for three different detection scenarios: first, Italy uses fast and efficient digital contact tracing like South Korea; second, Italy uses slower and less efficient manual tracing; and third, Italy continues to detect cases at its current low rate.² I find that, using digital tracing, the total cost of suppressing COVID-19 is 0.8% of annual GDP. The strategy allows for a fast and continuing reduction of social distancing. After 1.7 months already, optimal social distancing is at such a low level that its flow-cost is only 1 % of daily GDP. Afterward, the daily cost continues to converge to zero. The virus is entirely under control, and social activity is back to a normal level well before a vaccine arrives. Additionally, the strategy is robust to a certain degree of imported cases. The number of additional casualties under this strategy would be 3,413.

When using manual contact tracing, I find that the total cost of suppression is 2.8% of the annual GDP. The flow-cost of social distancing drops below 1% of daily GDP after 4

²Digital contact tracing uses mobile phone data to identify and inform the past contacts of a confirmed infectious individual. It is particularly fast and efficient. Its maximal detection rate is 35% per day (Ferretti et al., 2020). Manual contact tracing relies on teams of tracing personal who question confirmed infectious and find their contacts manually. Its maximal detection rate is 10% per day. See Ferretti et al. (2020) for an extensive discussion. Currently, Italy detects 2% of cases per day.

months. Manual tracing is not efficient enough to allow for a total return to normality. In the long run, some degrees of social distancing need to stay in place, however, its flow-cost is only 0.1% of daily GDP. I consider two stopping points for the pandemic. First, in the optimistic case, the virus disappears when prevalence falls below one case per million. Second, in the pessimistic case, the virus disappears only when a vaccine arrives in one and a half years. The total cost to reach both stopping points is essentially the same. The number of additional casualties would be 3,444.

In stark contrast, the total cost of suppression in the no tracing scenario is 14% of annual GDP in the optimistic case and 33% of annual GDP in the pessimistic case. The reason is that optimal social distancing is very close to constant under this scenario. Its flow-cost is 24% of daily GDP. The cost needs to be paid until the virus becomes extinct. In the optimistic case, extinction takes place after 7.6 months. However, the assumption of an extinction threshold is optimistic. If the virus survives in a small subpopulation in the form of one infectious, or one new case is imported from abroad, the pandemic restarts. Extinction by vaccination is robust to these concerns. In this case, a daily cost of 24% of GDP needs to be paid until the vaccine arrives. Experts estimate the arrival time at about one and a half years from now. The number of additional casualties for both cases would be 5,985. I compare the optimal suppression policy with optimal mitigation policies in the conclusion.

Methodologically, I exploit the fact that when the number of infectious is low compared to the number of susceptible, a simple exponential growth process approximates the dynamic behavior of infections well. In the limit, when the ratio of infectious per susceptible goes to zero, the approximation is exact. Note that it is the relevant case for studying suppression because the number of infectious goes to zero. The simplification avoids the heavy SIR machinery currently used in the literature. Additionally, but not less importantly, I eliminate the time variable from the planning problem, by writing it as a function of the stock of infectious only. Using this trick is without loss of generality. As a consequence, I can solve the model with pen and paper. It is imperative because it allows me to study the dependence of the optimal policy and welfare on the unknown functions and parameters. These unknowns are: the economic cost and viral growth impact of social distancing policies, the flow of death per infection and its social cost, the uncontrolled growth rate of the virus, and the speed of tracing as a function of the overall stock of infected. Although in

principle possible to estimate, at this stage we know very little about these key determinants for optimal policy. Therefore, it is crucial to study the properties of optimal policy without making restrictive assumptions on the unknowns.

The theoretical results follow from intuitive and straightforward properties of the problem. As already mentioned, the virus follows an exponential process. In the case of suppression, the process is exponential decay. The relative speed of tracing is the highest when the number of infectious is zero. It is when the most resources for tracing per infected individual are available. If even at zero, tracing is not efficient enough to be faster than the virus, then some amount of social distancing is always necessary to keep the viral growth rate in negative terrain. The viral growth rate is negative but bounded. Therefore, the decay of the virus is infinitely slow in the limit, a basic property of exponential decay. Note that an unbounded growth rate is infeasible. In particular, ignoring tracing for a moment, it takes the same intensity and time of social distancing to reduce the number of infected from 20 Million to 10 Million as from 20 to 10. The unit cost to reduce one infection goes to infinity as infections converge to zero. Towards the end of the pandemic, it is necessary to impose social distancing on the whole population for an extended period of time, just to avoid one last transmission of the virus. Efficient tracing offers an easy solution. Assume, at zero, tracing is faster than the growth of the virus. It follows that for a low enough number of infected, the virus disappears without the use of social distancing. All that is needed are small and targeted interventions to find and remove the last cases. A simple policy bounds the total cost of suppression. Use social distancing to push infections below the critical level, and then, let tracing do its job. The optimal policy always combines tracing with social distancing. However, it cannot be more costly than the simple policy.

The exact characterization of the optimal policy at each point in time by the two sufficient statistics follows from a simple intuition as well. Consider a particular current level of infectious, and consider the cost of reducing it by one unit. This unit cost is the current flow cost from deaths and social distancing measures, multiplied by the time it takes to reduce infections by one unit. Both factors, and therefore the product, depend on the intensity of social distancing. The stricter social distancing, the higher the cost, and the lower the time. The unit cost is at its minimum when the marginal change in cost divided by the cost is equal to the marginal change in time divided by the time - a property of interior extrema of products of functions. As time is inversely proportional to the growth rate, the

same property is valid for the growth rate instead of time. What follows is the optimality condition as a relation of two simple sufficient statistics: the current flow cost and the current growth rate. The optimal total cost is simply the integral over the optimal unit costs.

Relevant Literature This paper contributes to the economic literature on optimal disease control. A large and recent literature studies mitigation policies, using variants of the SIR model augmented with economic interactions. Mitigation controls the spread of the virus until contagions stop because the population achieves herd immunity. The policy is very distinct from a suppression policies, which I study in this paper. I discuss the relation of these two broad policy-approaches in the conclusion. The mitigation literature mostly uses numerical methods to solve for the optimal policy, or to simulates the impact of certain policies of interest. See Acemoglu et al. (2020), Alvarez et al. (2020), Atkeson (2020), Gollier (2020), Gonzalez-Eiras and Niepelt (2020), Miclo et al. (2020), Piguillem and Shi (2020), Bethune and Korinek (2020), Eichenbaum et al. (2020), Farboodi et al. (2020), Jones et al. (2020). The list is far from exhaustive. Garibaldi et al. (2020) and Assenza et al. (2020) characterize the theoretical properties of the optimal mitigation policy.

A smaller part of the literature studies suppression. Gollier (2020) and Ugarov (2020) simulate the impact of a uniform and strict lock-down, with the assumption that the last cluster can be removed with tracing. Wang (2020) simulates the impact of mass testing and shows that it can lead to suppression before herd immunity. I contribute to this literature by explicitly characterizing the optimal time-variable suppression policy.

Closest to my paper are Piguillem and Shi (2020), and Alvarez et al. (2020). Piguillem and Shi (2020) numerically solve for the optimal suppression policy in a SIR model under social distancing and random testing. They assume the last cluster of the disease can be removed by tracing. Alvarez et al. (2020) conduct a similar exercise but introduce tracing. They explicitly assume a functional form for tracing. Both of these contributions are quantitative. I contribute to this literature by characterizing the optimal suppression policy as the solution to simple sufficient statistics. I derive its properties under general functional forms and parameter values. It is important because very little is know about key parameters and relevant functions influencing optimal policy. I show which properties of the tracing function have consequences for the optimal policy and its cost. The tracing function used by Alvarez et al. (2020) is infinitely efficient in the limit. This property may pushes the

quantitative analysis towards suppression instead of herd immunity as an exit strategy. The quantitative results in my paper are novel as well. They contribute to the quantitative suppression literature by explicitly considering different case detection policies.

Pueyo (2020) gives an extensive informal discussion of possible policy responses.

2 The Model

Assume there is an initial mass I_0 of infectious individuals in a susceptible population. The virus transmits from infectious to susceptible. Infectious individuals die or recover from the disease after a certain time. If the mass of infected is small compared to the mass of susceptible, the spread of the virus follows the differential equation:

$$\dot{I}_t = r^0 I_t, \quad (1)$$

I_t is the mass of infectious at time t and \dot{I}_t is its time derivative. r^0 is the uncontrolled viral growth rate. It consists of two parts: $r^0 = \beta - \theta$. β is the rate of new contagions, θ is the rate of recovery or death from an infection. Assume that $r^0 > 0$, the virus is spreading. The equation describes an exponential growth process with a growth rate of r^0 . Note that at the beginning of the process, the mass of susceptible is large compared to the mass of infectious. However, that is also the case after an extended period of effective control measures such as social distancing. Even if part of the population is immune, as long as I_t is small compared to the number of susceptible, the above approximation is valid. The fraction of immune agents will simply be captured by a lower r^0 .

The policymaker can alter the spread of the virus by using two tools: case detection, and social distancing. Assume, in particular, the policymaker is interested in suppressing the virus, i.e., I_t converges to zero.

2.1 Case Detection

Case Detection allows for quarantining a mass X of infectious individuals at each instant of time. I assume infectious individuals in quarantine do not infect any susceptible individuals. $X(I)$ is the daily flow of detected cases into quarantine. Intuitively, it is the speed of detection. It depends on the mass of infected I . Assume $X(0) = 0$; if there are no

infected, none can be detected. $X'(I) > 0$; the speed of detection increases in the number of infected. When there are more infected it is easier to find them. $X''(I) < 0$; the increase in speed is decreasing in the number of infected. I assume that the overall capacity of case detection is fixed. I leave a generalization of this assumption for future research. The detection-technology becomes overwhelmed if there are too many infected, i.e., there are decreasing returns to scale. In the limit, if I goes to infinity, $X'(I)$ goes to zero. Define the detection rate as $\frac{X(I)}{I}$. Intuitively, it gives the percentage of overall cases that are detected daily. Under the above assumptions, the detection rate is decreasing in I . It is the largest at zero. The detection rate at zero is a key parameter for the analysis. Denote it as ξ_0 :

$$\xi_0 = \lim_{I \rightarrow 0} \frac{X(I)}{I} = X'(0). \quad (2)$$

Lemma 1. .

If the detection rate at zero is larger than r^0 , then there exists a level of infections I^ , such that for all $I < I^*$, it holds that $\dot{I} < 0$. I^* is the point where $r^0 I^* = X(I^*)$.*

It means that, as soon as new infections are below the threshold I^* , I_t converges to zero without any other policy intervention. The intuition behind this result is simple. The time derivative of infections is equal to

$$r^0 I - X(I). \quad (3)$$

The first summand is the speed at which the virus grows. It is decreasing in the number of infected, and it is zero in zero. The second summand is the speed at which infections are detected. Like the speed of viral growth, it decreases in the number of infected, and it is zero in zero. However, the relative speed of growth of the virus is constant, while the relative rate of tracing increases as I decreases. It is the largest at zero. What matters is if there is an infection level I^* , at which tracing is faster than the virus. Two cases are possible:

1. $\xi_0 < r^0$, tracing is never faster than the virus. In this case, detection alone can never suppress the virus.
2. $\xi_0 > r^0$, close enough to the origin, detection is faster than the virus. In particular, this is the case for all $I < I^*$. If $I < I^*$, tracing alone suppresses the virus.

In particular, it may hold that $\xi_0 = \infty$. Detection fulfills an Inada condition. In this case, suppression follows an accelerating decay process.

2.2 Social Distancing

Assume social distancing policies are indexed by $p \in [0, 1]$. Each policy p has mass zero. Each p reduces the growth rate of the virus by $dr(p)$ and has a social cost $dc(p)$. Assume policies are indexed such that the cost benefit ratio $\frac{dc(p)}{dr(p)}$ is increasing. Also, assume that $\frac{dc(0)}{dr(0)} = 0$. Applying policies 0 to p has a growth impact of

$$r(p) = \int_0^p p'(\tilde{p})d\tilde{p}. \quad (4)$$

Assume that there are enough policies available such that $r(1) \gg r^0$. Strict enough measures allow pushing the growth rate of the virus below zero, i.e., exponential decay. Denote by p_t the fraction of policies applied by the policymaker at time t . The spread of the virus follows the process:

$$\dot{I}_t = (r^0 - r(p_t))I_t. \quad (5)$$

If $r(p_t) > r^0$ the process follows an exponential decay. Physically, for any initial level of infections I_0 , the suppression of the virus is possible by keeping $r(p_t) > r^0$. However, assume p_t is large enough but constant. It follows that \dot{I}_t goes to zero as I_t goes to zero. The smaller I_t , the slower the suppression is advancing. In the limit, the process becomes infinitely slow.

The flow cost of applying policies 0 to p is

$$c(p) = \int_0^p c'(\tilde{p})d\tilde{p}. \quad (6)$$

To summarize, the functions $r(p)$ and $c(p)$ have the following properties: they are increasing and zero at zero, $\frac{c'(p)}{r'(p)}$ is increasing and zero at zero, and $r(1) \gg r^0$. Sometimes it is more convenient to express the flow-cost as a function of r instead of p :

$$c(r) = c(p(r)). \quad (7)$$

It follows that $c(r)$ is increasing and convex. The cost as well as the marginal cost are zero in the origin: $c(0) = 0$ and $c'(0) = 0$. Note that this is an abuse of notation. I use the same letter for two different functions. Which function is meant will be clear from the context.

Assume the policy maker wants to minimize the total economic cost of reducing infections from I_0 to 0:

$$C = \min_{p(\cdot)} \int_{I_0}^0 \frac{c(p(I))}{\dot{I}(I)} dI. \tag{8}$$

The solution to the problem is an optimal control function $p^*(I)$. I integrate over the number of infections instead of over time. Therefore, I need to divide the economic flow cost by the flow of infections \dot{I} . Mathematically, I change variable from t to I in the optimal control problem.

Define the optimal unit cost of reducing an infection as

$$\frac{dC(I)}{dI} = \frac{c(p^*(I))}{-\dot{I}(I)}. \tag{9}$$

Intuitively, consider the minimal cost ΔC to reduce infections by a small amount ΔI . The optimal unit cost is $\frac{\Delta C}{\Delta I}$.

Proposition 1. .

1. *When only using social distancing, the optimal cost-minimizing policy is constant over time.*
2. *The optimal policy p^* is equal to*

$$\frac{c'(p^*)}{c(p^*)} = \frac{1}{r(p^*) - r^0}. \tag{10}$$

3. *Assume $c(r)$ is iso elastic. It follows that the optimal effect of social distancing r^* is equal to*

$$r^* = r^0 \frac{\zeta_1 - 1}{\zeta_1}, \tag{11}$$

where $\zeta_1 > 1$ is the cost-elasticity.

4. *The optimal unit cost of reducing an infection $\frac{dC(I)}{dI}$ goes to infinity as I goes to zero.*

All proofs are in the appendix. The cost-efficiency of social distancing measures decreases as I decreases, even in the optimum. The reason, as discussed above, is that the

reduction in infectious becomes infinitely slow as I goes to zero. This result is quite intuitive. Given a certain intensity of social distancing, it takes the same time to reduce infections from 10 million to one million as reducing them from 10 to 1. Suppressing the virus by social distancing is possible, but very costly. If one takes the model literally, suppression takes infinitely long, and therefore, it is infinitely costly. Note that the policy p^* is the cost-minimizing policy in an economic sense. When maximizing social welfare - which takes the social cost of the flow of death into account- the result becomes even more extreme. Below I solve for the optimal policy, taking economic and social cost into account.

A word of caution. Modeling infections as a continuous mass has its limits when it represents only a handful of cases in the population. As soon as the number of infected is low, the transmission becomes granular. The literature uses a convenient shortcut to solve this problem. It assumes that the virus dies as soon as infections fall below some critical value \underline{I} . Under this assumption, the time to suppress the virus is finite. However, the results above still hold. If \underline{I} is close to zero, the relative cost of reducing the last infections is "close to infinity," i.e., very large.

3 The Optimal Policy

Assume there is a flow cost vI coming from the mass of infections. Note that infections resolve at some rate of θ . Individuals are not infectious forever. They recover, or they die at rate θ . Denote the probability that the outcome of an infection is death by δ . Denote the statistical value of life by l . The flow cost per infection v is equal to $v = \theta\delta l$.³ The cost can be generalized to a nonlinear cost in I , accounting for congestion effects in the health care sector. The problem of the policymaker is:

$$\min_{p_t} C(p_t) = \int_0^\infty c(p_t) + vI_t dt \quad (12)$$

such that

$$\dot{I}_t = (r^0 - r(p_t))I_t - X(I_t). \quad (13)$$

For now, I neglect time discounting. This assumption simplifies the problem considerably. Time discounting is not very important for the problem of the optimal suppression policy.

³ v may be interpreted more broadly as containing all other costs caused by an infection, such as the dis-utility of being sick and chronic damages caused by the virus.

The time frame is days and daily interest rates are very low. The solution to the problem is an optimal control function p_t . Assume that at the optimum $\lim_{t \rightarrow \infty} I_t = 0$ and $\dot{I}_t < 0$ for all t . These assumptions can be verified ex-post. However, they are very intuitive. If I_t does not converge to zero the integral does not exist as the integrated cost is infinitely large. $\dot{I}_t \geq 0$ cannot be optimal as letting infections grow increases the cost from death vI and only transfers the cost of reducing infections to a later point in time. Under these assumptions, I_t is strictly decreasing in time and therefore invertible. Use the invertibility of I to eliminate time in the minimization problem (12):

$$\min_{p(\cdot)} C(p(\cdot)) = \int_{I_0}^0 \frac{c(p(I)) + vI}{\dot{I}(I)} dI, \tag{14}$$

where

$$\dot{I}(I) = (r^0 - r(p(I)))I - X(I). \tag{15}$$

The solution to the problem is a control function $p(I)$. It is the solution to a simple point wise minimization of the above integral.

Proposition 2.

There always exists a unique optimal suppression policy $p(I)$. In particular, for each amount of currently infectious I , the optimal policy solves:

$$\frac{c'(p)}{c(p) + vI} = \frac{r'(p)}{r(p) + \frac{X(I)}{I} - r^0}. \tag{16}$$

In words, a policy is optimal if at each point in time, its relative effect on the flow cost is equal to its relative effect on the viral growth rate.

3.1 The Intuition Behind Proposition 2

To better understand the intuition behind the optimality condition, it is useful to recall each mathematical step in the derivation intuitively:

The first step is the change in the variable from t to I . Integrating over time means summing the flow costs at each point in time. Integrating over I means summing the flow cost for each reduction in I . The policymaker would like to reduce new infections from I_0 to 0. It is useful to think about the reduction as of a distance to cover. In particular,

partition the distance into many small and constant intervals of ΔI . The minimization problem consists in minimizing the cost for each of these intervals. The cost to reduce new infections at I to $I - \Delta I$ depends on the flow cost and the time it takes to cross the interval:

$$\underbrace{(c(p) + vI)}_{\text{flow cost}} \times \underbrace{\Delta t(I, p)}_{\text{crossing time}}. \tag{17}$$

Note that the crossing time is a function of p and I :

$$\Delta t = \frac{\Delta I}{(r(p) - r^0)I + X(I)}. \tag{18}$$

To find the optimal policy $p(I)$, take the logarithm of the above expression and perturb the current policy p by a small amount Δp to derive a change in cost ΔC :

$$\Delta C = \left(\frac{c'(p)}{c(p) + vI} + \frac{\frac{\partial \Delta t(I, p)}{\partial p}}{\Delta t(I, p)} \right) \Delta p. \tag{19}$$

A policy is optimal if there exists no policy perturbation that reduces the cost. It is the case when the expression in brackets is equal to zero. Instead of using Δt in the condition above, it is possible to express the same condition as a function of the growth rate of the virus. Define the growth rate g as $\frac{\dot{I}}{I}$. The crossing time Δt is inversely proportional to the growth rate:

$$\Delta t(I, p) = \frac{\Delta I}{I} \frac{1}{-g(I, p)}. \tag{20}$$

Therefore, the change in cost ΔC as a function of the growth rate is:

$$\Delta C = \left(\frac{c'(p)}{c(p) + vI} - \frac{\frac{\partial -g(I, p)}{\partial p}}{-g(I, p)} \right) \Delta p. \tag{21}$$

Using the definition of the growth rate gives the expression in Proposition 2:

$$\frac{c'(p)}{c(p) + vI} = \frac{r'(p)}{r(p) + \frac{X(I)}{I} - r^0}. \tag{22}$$

3.2 The Policy Implications of Proposition 2

Two simple sufficient statistics characterize the optimal policy: the current flow-cost and the current growth rate of the virus. The policymaker only needs to consider the relative change of these two statistics to a change in policy, to evaluate the optimality of the current policy. Optimality solely depends on current variables, which is somewhat surprising. The problem is a dynamic optimization problem, and, in principle, a decision at a certain point in time needs to account for its effects on the whole future to be optimal.

In particular, the optimality condition gives specific guidance to organize a de-confinement after an extended lock-down. For relaxing a certain confinement measure, the policymaker only needs to evaluate its relative impact on the current social cost and viral growth rate. If the relative reduction in cost is larger than the relative increase in growth, a measure should be relaxed. For instance, a policymaker may want to evaluate reopening a particular sector of the economy, for example, construction. The policymaker only needs information on how many percentage points such a measure would ease the current cost of the confinement and by how many percentage points it would increase the current growth rate of the virus to make an optimal decision.

Note that how to reopen, which is which policy to reverse first, is determined by the ratio $\frac{dc(p)}{dr(p)}$. Policies with a high ratio should be relaxed first. While the question of which policy to reverse first is by no means trivial empirically, it is not very difficult to answer theoretically. The harder theoretical question is how fast to reopen, which is determined by the above optimality condition. The optimality condition is robust to complementarities between policies, both in cost and in growth impact. The optimal decision only depends on the marginal impact of the most efficient policy at a certain point in time.

3.3 The Properties of the Optimal Policy

It is simpler to use r as a control variable instead of p , to derive the properties of the optimal policy. Note that such a change in the variable is without loss of generality. The optimal policy is characterized by a function $r(I)$.

Proposition 3. .

1. *In the optimum, social distancing measures are always positive, and increasing in the*

number of infectious. Social distancing is the largest at the beginning when $I = I_0$, and then gradually released, as the number of infectious decreases:

$$r(I) > 0, \text{ for all } I > 0, \text{ and } r'(I) > 0. \tag{23}$$

2. In the limit, as I goes to zero, the optimal policy $r(I)$ has the following properties:

- If $\xi_0 \geq r^0$, social distancing goes to zero: $\lim_{I \rightarrow 0} r(I) = 0$;
- If $\xi_0 < r^0$, social distancing goes to a constant: $\lim_{I \rightarrow 0} r(I) = 2(r^0 - \xi_0) > 0$.

3. Under the optimal policy, the growth rate of I is negative: $g(I) < 0$. In the limit it is equal to $\lim_{I \rightarrow 0} g(I) = -|\xi_0 - r^0|$. In particular, the growth rate goes to $-\infty$ if $\xi_0 = \infty$.

Note that for the case $\xi_0 < r^0$, I assume a quadratic cost to derive the results above. The proposition underlines the key role of ξ_0 , i.e., the rate of detection at zero. It governs the amount of time it takes to suppress the virus and the optimal policy in the limit. If detection is efficient enough, it is possible to gradually go back to normal. However, it is not the case when case detection is not efficient enough. Note that the efficiency of detection is characterized sharply by the derivative of the flow of detections in zero. With inefficient detection, some amount of social distancing needs to stay in place "forever," with stark consequences for the total cost. "Forever" stands for the time until another solution, like a vaccine, is found. However, the efficiency of detection still matters in this case. It determines the level of necessary social distancing in the limit. The necessary level may contain only mild measures such as washing hands, wearing masks, and forbidding mass events. In the next step, I study the cost of suppression at the optimum. The above proposition already gives a preview for the case where $\xi_0 < r^0$. The optimal policy does not go to zero in the limit; therefore, the cost of applying it does not go to zero. On top of that, the time to suppress is infinite. It will follow that suppression is infinitely costly in this case.

3.4 The Cost of the Optimal Policy

In the optimum, the total cost of suppressing I_0 infectious is

$$C = \int_0^{I_0} \frac{c(r(I)) + vI}{\left(r(I) + \frac{X(I)}{I} - r^0\right) I} dI, \tag{24}$$

where $r(I)$ denotes the optimal policy. The unit cost of suppression, intuitively, the cost to suppress one more infectious, is equal to

$$\frac{dC}{dI} = \frac{c(r(I))}{\left(r(I) + \frac{x(I)}{I} - r^0\right) I} + \frac{v}{\left(r(I) + \frac{x(I)}{I} - r^0\right)}. \quad (25)$$

It consists of two parts: an economic unit cost, which comes from the taken suppression policies, and a social unit costs, which comes from the flow of death.

Proposition 4.

Case 1, $\xi_0 > r^0$:

- *As I converges to zero, the economic unit cost of suppression converges to zero, and the social unit cost of suppression converges to $\frac{v}{\xi_0 - r^0} \geq 0$. In particular, if $\xi_0 = \infty$, also the social unit cost is zero.*
- *The total cost of suppression is bounded.*

Case 2, $\xi_0 < r^0$:

- *As I converges to zero, the economic unit cost of suppression converges to infinity, and the social unit cost of suppression converges to $\frac{v}{r^0 - \xi_0} > 0$.*
- *The total cost of suppression is infinite.*

The case $\xi_0 < r^0$ assumes a quadratic cost. The proposition underlines the importance of the properties of case detection when I goes to zero. If the rate of detection is high enough, it is possible to gradually go back to normal, which bounds the total cost of suppressing the virus. If the rate is not high enough, the total cost is infinite. It comes from the fact that social distancing measures need to stay in place forever. The reason is that the suppression becomes infinitely slow in the limit. However, this does not mean suppression is not a good idea. The necessary long-run distancing measures may be very mild, and therefore worth enduring. Even if $\xi_0 < r^0$, its size still matters, because it determines the amount of social distancing necessary in the long run. It may still be cheaper to suppress the virus and wait for a vaccine than to use another solution, such as herd immunity. Especially, suppression avoids the risk that the virus mutates and becomes endemic.

4 Quantitative Results

To implement an optimal policy, it is necessary to know the functions $c(p)$, $r(p)$ and $X(I)$, as well as the parameters r^0 and v . To estimate these functions should be a top priority for future research. In particular, the properties of $X(I)$ for small values of I are especially important. They drive the optimal policy at the end of the suppression process, where social distancing loses its effectiveness. The calibration exercise in this chapter uses a rough approximations of these functions. The exercise is useful to obtain rough estimates of the cost of different policy options. On top of that, it illustrates the qualitative results from the former sections.

I calibrate my model in the case of Italy.

4.1 Data Sources

I use frequently updated epidemiological data from the Institute for Health Metrics and Evaluation (IHME) at the University of Washington. They provide a time series of confirmed cases, as well as estimates for the real number of daily infections for many countries. Their estimates are based on Murray et al. (2020). I use data from Italy and South Korea.

4.2 Calibration

4.2.1 Parameters Literature

I use the following parameters from the literature as a starting point for my calibration:

Parameter	Symbol	Value	Source
Mortality rate	δ	0.01	Alvarez et al. (2020)
Time of contagiousness	$\frac{1}{\theta}$	6 days	Fernández-Villaverde and Jones (2020)
Value of statistical life	vs_l	$20 \frac{\text{GDP}}{\text{capita}}$	Alvarez et al. (2020)
Uncontrolled growth rate	r^0	0.14	Ferretti et al. (2020)
Max. rel. speed digital tracing	ξ_0^d	0.35	Ferretti et al. (2020)
Max. rel. speed manual tracing	ξ_0^m	0.1	Ferretti et al. (2020)
GDP loss strict lock-down	c_{LD}	0.5	Gollier (2020)

Table 1: Parameters Literature

4.2.2 The Cost Function

I use a direct relation between the cost of social distancing, measured as lost GDP, and the reduction in the viral growth rate r . I assume the function is iso-elastic:

$$c(r) = \zeta_0 r^{\zeta_1}, \quad (26)$$

where $\zeta_0 > 0$ and $\zeta_1 > 1$. Note that, under the assumption of an iso-elastic cost, and neglecting the value of lives lost as well as tracing, it holds that the optimal r solves

$$\zeta_1 = \frac{r}{r - r^0}. \quad (27)$$

From the the 10th of March until the 26th of April the Italian government imposed a nation wide lock-down. A lock-down is a strict form of social distancing. Any no-essential social contact is forbidden. A large part of the population is forced to stay at home, i.e., a "stay-at-home order". Going outside is permitted only if absolutely essential. To calibrate ζ_1 , I assume the strict lock-down in Italy was close to optimal. Note that Italy did not use much tracing during the time of the lock-down. The value of lives lost is small compared to the lost GDP. The assumption of optimality from the part of the government is strong. However, in many countries such as France, Spain, the UK, and Germany, we have seen very similar intensities of lock-downs. This is consistent with Equation (27). Note that the optimal intensity of r does not depend on the level of infections. It only depends on ζ_1 , which parametrizes the convexity of the cost. Once a country discovers an outbreak, it should hit hard to reduce new infections. If tracing is infeasible in the short term, and the number of death is relatively small, Equation (27) is a good approximation for the optimal policy. The intensity of r does only depend on the convexity of the cost ζ_1 . Note that the convexity should be similar between countries. The more convex the cost $c(\cdot)$, the more it costs to implement a hard lock-down. The similar intensities in between different countries suggest that governments followed, at least approximately, the optimal lock-down strategy. A different interpretation of the optimality assumption is that it makes the results consistent with the strict lock-down. In this case, the implied optimal policies are consistent with the observed past behavior of the government, even if this past behavior was not optimal.

Under the assumption that the intensity of the lock-down was optimal, it is informative

about the convexity of its cost. Using the epidemiological data from Murray et al. (2020), I estimate the growth rate under the Italian lock-down at $g_{LD} = -0.036$. I use the estimated number of new infections from the peak on March 11th until the most recent estimates on May 11th. Together with an uncontrolled growth rate of $r^0 = 0.14$ (see Ferretti et al. (2020)), I calculate the growth reduction from the lock-down at $r_{LD} = r^0 - g_{LD} = 0.176$. Using Equation (27), it implies an elasticity of $\zeta_1 \approx 5$. Following Gollier (2020), I assume a strict lock-down implies a daily GDP loss of around $c_{LD} = 50\%$. It implies a parameter $\zeta_0 \approx 3000$.

4.2.3 The Tracing Function

I use the following tracing function:

$$X'(I) = \left(\xi_0^{-\frac{1}{\alpha}} + \xi_1 I \right)^{-\alpha}, \quad (28)$$

and $X(0) = 0$. The function fulfills the necessary properties of a tracing function, i.e., it is zero at zero, increasing and convex. $\xi_0 > 0$ is the value of the function for $I = 0$. Note that it is equal to $\lim_{I \rightarrow 0} \frac{X(I)}{I}$, i.e., the relative speed of tracing at zero. The parameter $\xi_1 > 0$ controls the behavior of the function for large values of I . $\alpha \geq 0$ controls how fast $X'(I)$ goes from ξ_0 to its behavior for large I . Note that this function is quite general and contains some intuitive tracing functions as special cases. For $\alpha = 0$ it reduces to a constant returns to scale tracing function: $X(I) = \xi_0 I$. In particular, if ξ_0 is equal to the daily flow of tests, it is equal to tracing under random testing. When ξ_0 goes to infinity, the function reduces to a power function as used in Alvarez et al. (2020). The disadvantage of a power function is that $X'(0) = \infty$ by assumption. This assumption is unrealistic. It makes tracing overly efficient at the end of the epidemic.

To calibrate the parameters, I distinguish two cases: digital tracing and manual tracing. I use micro estimates to calibrate the function for both cases. I use results from Ferretti et al. (2020). This epidemiological paper estimates by how much optimal contact tracing can reduce daily new infections. They compare digital contact tracing with manual contact tracing. Ferretti et al. (2020) estimate that, under optimal conditions, digital contact tracing can find infectious individuals at a rate of $\xi_0^d = 35\%$ per day. It means that the stock of currently infectious can be reduced by 35% in one day. Manual contact tracing is much slower. Because of unavoidable delays, the authors argue that optimal manual contact trac-

ing achieves a rate of $\xi_0^m = 10\%$ per day. I use these estimates as values for ξ_0 in the two cases. I assume that, if a country uses its full resources to find the last cases, tracing achieves its optimal rate. However, as soon as the caseload grows, the system becomes overwhelmed, and the efficiency of tracing decreases.

I calibrate ξ_1 such that at a prevalence level of 10%, i.e., 10% of the population is infected at the same time, $X'(I) = \xi_0/10000$, which is close to zero. It means that at a prevalence level of 10%, the system is so overloaded that any further increase in the number of infected will not lead to more traced cases. To calibrate α , I use estimates of the fraction of traced cases from Italy and South Korea. Using data from Murray et al. (2020), I estimate that Korea, using digital tracing, at a prevalence of 60 infected per million, found 20% of total cases daily. I assume the number of confirmed cases is equal to the number of traced cases. Note that the estimated rate is not too far from the theoretical limit of 35%. It implies that, for digital tracing, $\alpha_d = 1.2$. For manual tracing, I use the same procedure using data from Italy. Recently, at an estimated prevalence level of 1000 per million, Italy manages to confirm 2% of the total cases daily. It implies that, for manual tracing, $\alpha_m = 1.4$.

4.2.4 Remaining Parameter Values

As already mentioned, I use $r^0 = 0.14$ as in Ferretti et al. (2020).

To estimate the current prevalence in Italy, I use the data from Murray et al. (2020). I use May 11th as a starting date because it is the last date of available observations. Murray et al. (2020) estimate daily new infections. I use new infections to calculate the current stock of infectious by summing the infections over the 6 preceding days. I assume an infected stays infectious for $1/\theta = 6$ days, following Fernández-Villaverde and Jones (2020). I find a current level of prevalence for Italy of $I_0 = 0.001$.

To estimate the cost of the flow of death, I assume that an infectious dies with probability $\delta = 0.01$, following Alvarez et al. (2020). I assume that an infectious dies $1/\theta = 6$ days after being infected. Note that, in general, this is not true. However, because I do abstract from time discounting, this assumption is without loss of generality. Following Alvarez et al. (2020), I use a value of statistical life of 20 times the annual output per capita. It

follows that $v = vsl * 365 * \delta * (1/\theta) = 14.4$. It means that if the whole population was infected, society incurs a flow-cost from casualties of around 14 times its daily GDP. Note that, because of the low prevalence level, my results are very insensitive to the assumptions related to mortality.

4.2.5 Summary Relevant Parameters

Parameter	Symbol	Value	Matched Moment or Source
Factor cost function	ζ_0	3000	GDP loss lock-down
Cost-elasticity	ζ_1	5	Lock-down intensity
Max. rel. speed digital tracing	ξ_0^d	0.35	Ferretti et al. (2020)
Max. rel. speed manual tracing	ξ_0^m	0.10	Ferretti et al. (2020)
Scalability digital tracing	α_d	1.2	Confirmed cases Korea
Scalability manual tracing	α_m	1.4	Confirmed cases Italy
Initial Prevalence	I_0	0.1%	Estimate for Italy May 11th
Flow value of casualties	v	14.4	Alvarez et al. (2020)
Uncontrolled growth rate	r^0	0.14	Ferretti et al. (2020)

Table 2: Relevant Parameters

4.3 Results

I take the current level of prevalence in Italy as given and analyze the optimal suppression policy for three different tracing scenarios:

1. Italy continues to isolate infectious at the current rate of 2% per day. I refer to this case as no tracing.
2. Italy adopts an optimal manual contact tracing strategy.
3. Italy adopts an optimal digital contact tracing strategy like South Korea.

To compare the three scenarios, I compare the intensity of the optimal social distancing measures, their implied flow costs, the time it takes to reach certain thresholds in daily cost, as well as the total cost.

Cost or intensity measure	No tracing	Manual tracing	Digital tracing
Optimal intensity r on May 11th	0.15	0.15	0.13
Implied daily cost [daily GDP]	24 %	24 %	13%
Time until cost drops to 10% of daily GDP	never	35 days	10 days
Time until cost drops to 1% of daily GDP	never	4 months	1.7 month
Daily cost in the limit	24%	0.1%	0
Time to reach extinction threshold 1 ppm	7.6 months	15 months	3.7 months
Total cost until extinction [annual GDP]	14.3%	2.7%	0.8%
Total additional death until extinction	5,985	3,444	3,413
Total cost until vaccine [annual GDP]	33 %	2.8%	0.8%

Table 3: Comparison Scenarios

The optimal reduction in the intensity of social distancing is such that r goes from 0.17 under the lock-down, to 0.15 (no tracing/manual tracing), and 0.13 (digital tracing). Some degree of easing is optimal. The reason is that identifying a fraction of the contagious takes over some of the burdens to keep viral growth at an optimal level. This modest reduction in social distancing already has an important impact on economic cost. It reduces from 50% of daily GDP under the lock-down, to 24 % under no tracing/manual tracing, and 13% under digital tracing. It means that, on May 11th, it is possible to ease the lock-down by around half, measured in lost daily GDP. An immediate switch to the Korean strategy would allow for an easing of a factor of 4.

The cost of social distancing drops over time because the number of infectious reduces and social distancing is gradually relaxed in the optimum. Under digital tracing, the cost drops below 10% of daily GDP after ten days already. Under manual tracing, it takes 35 days to reach this point. However, under no tracing, this point is never reached. The optimal intensity is almost constant and stays close to 0.15. The cost of 24% of daily GDP needs to be paid until the virus disappears. The time until the daily cost drops below one % is 4 months under manual tracing vs. 1.7 month under digital tracing. Not that at this point the crisis is de-facto over as the economy returns to an activity level very close to normal. In the long run, the cost reduces further to 24% for no tracing (it is practically constant), to 0.1% for manual tracing, and 0% for digital tracing. The numerical result confirms the theoretical results. Only efficient tracing with $\xi_0 > r^0$ allows the society to go back to a normal activity level. If tracing is inefficient, i.e., $\xi_0 \ll r^0$ relatively strong and costly social distancing measures need to stay in place. Mild efficiency, implies that measures have to stay in place in the long run, however, they are mild and not very costly. For instance,

this may corresponds to the case where society only imposes restrictions on mass events and general hygiene measures such as mask wearing.

Next, to compare the total cost of the different strategies, following Piguillem and Shi (2020), I assume the virus disappears when prevalence falls below an extinction threshold of 1 infectious per million inhabitants. Piguillem and Shi (2020) use a threshold of 10 per million. I use a more conservative threshold because, currently, South Korea already reached a prevalence of 6 per million, and the virus is not extinct.

The differences in cost between the different strategies are enormous. No tracing takes 7.6 months and costs 14.3% of annual GDP. Note that this cost is in addition to the already incurred cost due to the strict lock-down. Manual tracing is slower but much less costly. The reason is that social distancing is gradually relaxed in the optimum. It reaches a limit where its cost is only 0.1% of daily GDP. The total cost is 2.7% of annual GDP. This cost is still substantial. The cheapest option is digital tracing. The virus disappears in 3.7 months. Social distancing is relaxed quickly and substantially, well before that date. The total cost is only 0.8% of annual GDP. Note that this cost is by an order of magnitude smaller than estimates for the total cost under optimal mitigation strategies. Acemoglu et al. (2020) and Gollier (2020) evaluate mitigation strategies with age-dependent social distancing measures. They find a total cost of mitigation in the range of 7 to 13 % of total GDP.

However, my estimate is somewhat optimistic as it relies on the assumption that the virus disappears when prevalence falls below a prevalence of 1 per million. Note that this assumption is very common in the literature and implicit in quantitative models. There are two problems with this assumption. First, even if the assumption is correct, there is always the possibility that a new case is imported after social distancing is relaxed. In this case, the pandemic starts from the beginning. Second, the assumption may not be correct, even if the system is closed and an import of cases can be avoided. Some cases may always survive in a subset of the population. Therefore, it is safer to assume that the virus does never really disappear by itself, prevalence converges to zero but never reaches the limit.

To compare costs in the case where there is no extinction threshold, I assume the pandemic is over after 1.5 years because a vaccine arrives. I assume that this arrival comes in the form of a one-time, unanticipated shock. Under the no-tracing and the manual tracing

strategy, the daily cost of social distancing is close to the long-run value after passing a prevalence of 1 ppm. It means that this long-run cost needs to be paid from that point until the vaccine arrives. Under this conservative assumption, the total cost for the two tracing strategies stays essentially unchanged. The reason is that, after passing the threshold of 1 per million, optimal social distancing is low and not very costly anymore. This result shows the big advantage of tracing, and especially efficient tracing, with $\xi_0 > r^0$. Tracing controls the virus and social distancing can be relaxed, completely, or to a very low level. Even imported cases are not a problem, because tracing removes them. Under digital tracing, the society converges back to normal and the arrival of the vaccine is not even necessary. Under manual tracing, the society incurs a trivial permanent cost of 0.1% of daily GDP. The picture looks very different for the no tracing strategy. A relatively strong - and therefore costly (24% of daily GDP) - amount of social distancing needs to be sustained until the vaccine arrives, to avoid a new outbreak. The total cost is 33 % of annual GDP.

5 Conclusion

This paper characterizes the optimal policy to suppress COVID-19. I find that a complete and efficient eradication of COVID-19 is possible at a reasonable economic cost of 0.8% of annual GDP. The optimal suppression policy is easily implementable. However, some crucial questions are still unanswered. In particular, is it more efficient to use mitigation or suppression?

Mitigation controls the spread of the virus until contagions stop because the population achieves herd immunity. The problem with this policy is that a very large part of the population has to get infected with the virus, which leads to an important number of lives lost. On top of that, the strategy bears the risk that immunity vanishes, or that the virus mutates. In both cases, the virus could become endemic, i.e., circulate in the human population for a very long time. In contrast, suppression pushes the viral growth rate into negative terrain, such that the virus disappears in the long run. The policy avoids the infection of a large part of the population and the risk that the virus becomes endemic is very low.

If the current number of infectious individuals is sufficiently low, and case detection is efficient enough, the answer to this question is undoubtedly suppression. The same is true if the value of lives lost is large enough. However, for all other cases, it becomes much

harder to take an optimal decision. Moreover, a decision needs to be taken. The two policy responses dictate a very different optimal time path of infections. A mitigation policy lets infections grow at some point, because the virus needs to reach a large enough part of the population. Optimal suppression never lets infections grow. The policy maker stands at a crossroad and needs to decide which path to take. The total cost of either of them is still very uncertain. It depends crucially on: the cost and viral growth impact of social distancing policies, the speed of tracing, especially at low infection levels, the statistical value of life, and the capacity of the health care system and its impact on mortality rates. All of these variables are highly uncertain. Only the precise estimates of the mentioned unknowns can give a definite answer to the question.

However, the calibration exercise in this paper can give rough guidance on how to answer the question. I find that the total cost of suppression is 0.8% of annual GDP when using digital contact tracing and 2.8% of GDP when using manual tracing. In comparison, the cost-estimates of an optimal mitigation strategy range from 7% (Gollier, 2020), to 14% (Acemoglu et al., 2020), to around 30% (Alvarez et al., 2020). The two lower numbers rely on the assumption that is possible to shelter the most vulnerable part of the population. None of these estimates take the risk that the virus could become endemic into account. Additionally, mitigation strategies imply a much higher number of casualties. The cost-estimates depend strongly on the statistical value of life. The comparison suggests that suppression is the most cost-efficient strategy. It is certainly the strategy that reduces the number of casualties.

Curiously, it is easier to find the exact optimal amount of social distancing at each point in time when following suppression, than to decide on the optimal broad direction of policy. The policymaker only needs to turn to the econometrician and the epidemiologist - they can estimate the local impact of a policy change on the flow cost and the viral growth rate - and apply the condition derived in this paper.

References

- Acemoglu, D., Chernozhukov, V., Werning, I., and Whinston, M. D. (2020). A multi-risk sir model with optimally targeted lockdown. Technical report, National Bureau of Economic Research.
- Alvarez, F. E., Argente, D., and Lippi, F. (2020). A simple planning problem for covid-19 lockdown. *Covid Economics*, 14.
- Assenza, T., Collard, F., Dupaigne, M., Fève, P., Hellwig, C., Kankanamge, S., Werquin, N., et al. (2020). The hammer and the dance: Equilibrium and optimal policy during a pandemic crisis. Technical report, Toulouse School of Economics (TSE).
- Atkeson, A. (2020). What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios. Technical report, National Bureau of Economic Research.
- Bethune, Z. A. and Korinek, A. (2020). Covid-19 infection externalities: Trading off lives vs. livelihoods. *Covid Economics*, 11.
- Eichenbaum, M. S., Rebelo, S., and Trabandt, M. (2020). The macroeconomics of epidemics. Technical report, National Bureau of Economic Research.
- Farboodi, M., Jarosch, G., and Shimer, R. (2020). Internal and external effects of social distancing in a pandemic. *Covid Economics*, 9.
- Fernández-Villaverde, J. and Jones, C. (2020). Estimating and simulating a sird model of covid-19 for many countries, states, and cities.
- Ferretti, L., Wymant, C., Kendall, M., Zhao, L., Nurtay, A., Abeler-Dörner, L., Parker, M., Bonsall, D., and Fraser, C. (2020). Quantifying sars-cov-2 transmission suggests epidemic control with digital contact tracing. *Science*, 368(6491).
- Garibaldi, P., Moen, E. R., and Pissarides, C. A. (2020). Modelling contacts and transitions in the sir epidemics model. *Covid Economics*, 5.
- Gollier, C. (2020). Cost-benefit analysis of age-specific deconfinement strategies. christian golliertoulouse school of economics. Technical report, University of Toulouse-Capitole.
- Gonzalez-Eiras, M. and Niepelt, D. (2020). On the optimal lockdown during an epidemic. *Covid Economics*, 7.

- Jones, C. J., Philippon, T., and Venkateswaran, V. (2020). Optimal mitigation policies in a pandemic: Social distancing and working from home. Technical report, National Bureau of Economic Research.
- Miclo, L., Spiro, D., and Weibull, J. (2020). Optimal epidemic suppression under an icu constraint. *arXiv preprint arXiv:2005.01327*.
- Murray, C. J. et al. (2020). Forecasting the impact of the first wave of the covid-19 pandemic on hospital demand and deaths for the usa and european economic area countries. *medRxiv*.
- Piguillem, F. and Shi, L. (2020). Optimal covid-19 quarantine and testing policies.
- Pueyo, T. (2020). Coronavirus: The hammer and the dance. Technical report, blog post, Medium. URL <https://medium.com/@tomaspueyo/coronavirus-the-hammer-and-the-dance-be9337092b56>.
- Ugarov, A. (2020). Inclusive costs of npi measures for covid-19 pandemic: Three approaches. *medRxiv*.
- Wang, Y. (2020). An analytical sir model of epidemics and a sustainable suppression policy: Testing. *Available at SSRN 3573979*.

A Appendix

A.1 Proofs

A.1.1 Proof Proposition 1

PROOF:

Guess that $\dot{I} < 0$ in the optimum and verify ex-post. The change of variable from t to I is valid as $I(t)$ is invertible. Minimize the integral point-wise to get the first order condition

$$\frac{c'(p)}{c(p)} = \frac{1}{r(p) - r^0}, \quad (29)$$

which proves point two. The condition does not depend on I which proves point 1. For the discussion of existence and uniqueness consider the proof of Proposition 2. Consider the cost $c(r)$ as a function of r instead of p , and assume it is iso-elastic:

$$c(r) = \zeta_0 r^{\zeta_1}, \quad (30)$$

with $\zeta_0 > 0$ and $\zeta_1 > 1$. Point 3 follows from using $c(r)$ in the FOC. Point 4 follows from taking the limit in the definition of the optimal unit cost.
qed.

A.1.2 Proof Proposition 2

The minimum of the integral

$$\min_{p(\cdot)} C(p(\cdot)) = \int_0^{I_0} -\frac{c(p(I)) + vI}{(r^0 - r(p(I)))I - X(I)} dI \quad (31)$$

is at the point-wise minimum of each integrand. Note that I swapped the bounds. Change policy variable from p to r . For each I , the integrand is equal to

$$\frac{c(r) + vI}{\left(r - r^0 + \frac{X(I)}{I}\right) I}. \quad (32)$$

Note that $\dot{I} < 0$ by assumption. Therefore, the denominator has to be positive, which is the case when $r > r^0 - \frac{X(I)}{I}$. Also, $r \geq 0$ by definition. There are two cases:

First, if $r^0 - \frac{X(I)}{I} > 0$ it holds that $r \in \left(r^0 - \frac{X(I)}{I}, \infty\right)$.

Second, if $r^0 - \frac{X(I)}{I} \leq 0$ it holds that $r \in [0, \infty)$.

Note that the integrand is finite, positive, and continuous for any interior r .

Lemma 2. *There exists a minimum of the integrand and it is interior.*

PROOF:

Case 1, $r^0 - \frac{X(I)}{I} > 0$:

It follows that $r \in \left(r^0 - \frac{X(I)}{I}, \infty\right)$. If r goes to the left limit, the integrand goes to infinity.

If r goes to the right limit, the integrand goes to infinity as well. To see that, take the limit:

$$\lim_{r \rightarrow \infty} \frac{c(r) + vI}{\left(r - r^0 + \frac{X(I)}{I}\right) I} = \lim_{r \rightarrow \infty} \frac{c(r)}{r} = \lim_{r \rightarrow \infty} \frac{c'(r)}{1} = \infty. \quad (33)$$

The integrand is finite, positive, and continuous for any interior r . It follows that there exists an interior minimum.

Case 2, $r^0 - \frac{X(I)}{I} < 0$:

It follows that $r \in [0, \infty)$. At the left boundary, the integrand is equal to $\frac{v}{\frac{X(I)}{I} - r^0}$. The minimum cannot be at zero, because the integrand is strictly decreasing in zero:

$$\frac{c'(0) \left(0 - r^0 + \frac{X(I)}{I}\right) I - (c(0) + vI) I}{\left(0 - r^0 + \frac{X(I)}{I}\right)^2 I^2} = \frac{-v}{\left(-r^0 + \frac{X(I)}{I}\right)^2} < 0. \quad (34)$$

If r goes to the right limit, the integrand goes to infinity. The argument is as in case 1. The integrand is finite, positive, and continuous for all r . It follows that there exists an interior minimum.

Case 3, $r^0 - \frac{X(I)}{I} = 0$:

It follows that $r \in (0, \infty)$. As case 1. The integrand goes to infinity at both boundaries.

qed.

Any interior extremum fulfills the first order condition:

$$\frac{c(r) + vI}{\left(r - r^0 + \frac{X(I)}{I}\right) I} \left(\frac{c'(r)}{c(r) + vI} - \frac{1}{r + \frac{X(I)}{I} - r^0} \right) = 0 \tag{35}$$

Cancel the left factor and change the choice variable back from r to p to get the optimality condition in Proposition 2.

Each interior extremum is a strict minimum. To see that, rearrange the first derivative of the integrand to:

$$\frac{1}{\left(r - r^0 + \frac{X(I)}{I}\right)^2 I} \left(-(c(r) + vI) + c'(r) \left(r + \frac{X(I)}{I} - r^0 \right) \right) \tag{36}$$

Take the derivative to get the second order condition. Note that it is equal to $a'b + ab'$ where a is the first factor above and b is the second factor. If the FOC holds, b is zero. Also, a is always positive. The sign of the SOC only depends on the sign of b' . $b' = c''(r) \left(r + \frac{X(I)}{I} - r^0 \right)$, which is strictly greater than zero.

As each minimum is a strict minimum, and the function is continuous, there can only be one minimum. In particular, it fulfills the first order condition.

qed.

A.2 Proof Proposition 3

Lemma 3. .

Consider the case where $\lim_{I \rightarrow 0} \frac{X(I)}{I} \geq r^0$. It follows that:

1) The optimal policy $r(I)$ converges to zero as I converges to zero:

$$\lim_{I \rightarrow 0} r(I) = 0. \tag{37}$$

2) For small I the optimal policy $r(I)$ is approximately equal to

$$r(I) \approx - \left(\frac{X(I)}{I} - r^0 \right) + \sqrt{\left(\frac{X(I)}{I} - r^0 \right)^2 + 2 \frac{v}{c''(0)} I}. \tag{38}$$

In particular, $r(I) > 0$ for $I > 0$.

3) For small I the growth rate under the optimal policy $g(I)$ is approximately equal to

$$g(I) \approx -\sqrt{\left(r^0 - \frac{X(I)}{I}\right)^2 + 2\frac{v}{c''(0)}I}. \tag{39}$$

4) For

$$\lim_{I \rightarrow 0} \frac{X(I)}{I} = \infty, \text{ it holds that } \lim_{I \rightarrow 0} g(I) = -\infty. \tag{40}$$

The decay of the virus is accelerating as I approaches zero.

PROOF:

The optimality condition with r as the policy variable writes

$$\frac{c'(r)}{c(r) + vI} = \frac{1}{r + \frac{X(I)}{I} - r^0}. \tag{41}$$

Taylor approximate the function $c(r)$ in the origin:

$$c(r) \approx \frac{1}{2}c''(0)r^2. \tag{42}$$

Use the approximation in the optimality condition to solve for Equation (38), which proofs point 2). Note that $r(I)$ is the solution of a quadratic equation. The second solution can be discarded as it violates $\dot{I} < 0$. Point 1) follows from taking the limit in Equation (38). Point 3) follows from using the definition of the growth rate. Point 4) follows from taking the limit in Equation (39).

qed.

Lemma 4.

Consider the case where $\lim_{I \rightarrow 0} \frac{X(I)}{I} = \xi_0 < r^0$. Assume that the cost function is quadratic: $c(r) = \frac{1}{2}c''(0)r^2$ It follows that:

1) As I converges to zero, the optimal policy $r(I)$ converges to:

$$\lim_{I \rightarrow 0} r(I) = 2(r^0 - \xi_0). \tag{43}$$

In particular, if there is no test and trace $\xi_0 = 0$, and

$$\lim_{I \rightarrow 0} r(I) = 2r^0. \quad (44)$$

3) The optimal policy $r(I)$ is equal to

$$r(I) = r^0 - \frac{X(I)}{I} + \sqrt{\left(r^0 - \frac{X(I)}{I}\right)^2 + 2\frac{v}{c''(0)}I}. \quad (45)$$

3) The implied optimal growth rate $g(I)$ is equal to

$$g(I) = -\sqrt{\left(r^0 - \frac{X(I)}{I}\right)^2 + 2\frac{v}{c''(0)}I}. \quad (46)$$

In particular $r(I) > 0$ for all I .

4) Under the optimal policy $r(I)$ the growth rate converges to

$$\lim_{I \rightarrow 0} g(I) = -(r^0 - x^0). \quad (47)$$

PROOF:

As above. However, the cost function is quadratic by assumption and not by approximation.

qed.

Lemma 5. .

The optimal policy $r(I)$ is strictly increasing in I :

$$r'(I) > 0. \quad (48)$$

PROOF:

The optimal policy solves

$$\frac{c'(r(I))}{c(r(I)) + vI} = \frac{1}{r(I) + \frac{X(I)}{I} - r^0}. \quad (49)$$

Differentiate with respect to I to obtain

$$r'(I) = \frac{v - c'(r(I)) \frac{dX(I)}{dI}}{-g(r(I), I)c''(r(I))}. \tag{50}$$

The expression is positive as $v > 0$, $c'(\cdot) > 0$, $\frac{dX(I)}{dI} < 0$, $g(r(I), I) < 0$ and $c''(\cdot) > 0$.

A.3 Proof Proposition 4

Lemma 6. .

If $\xi_0 = \infty$, the total unit cost of suppression goes to zero as the mass of infectious goes to zero.

PROOF:

$\frac{dC}{dI}$ is the unit cost at the optimum. It is smaller or equal to the unit cost under any other policy that satisfies $\dot{I}(I) < 0$. In particular, take the policy $\tilde{r}(I) = 0$ for all $I < I^*/2$. It follows that

$$0 \leq \frac{c(r(I))}{\left(r(I) + \frac{X(I)}{I} - r^0\right) I} + \frac{v}{\left(r(I) + \frac{X(I)}{I} - r^0\right)} \leq \frac{v}{\frac{X(I)}{I} - r^0}. \tag{51}$$

Take the limit on both sides to obtain the result.

qed.

Lemma 7. .

If $\xi_0 > r^0$, the economic unit cost of suppression goes zero, and the social unit cost from the flow of death goes to a constant, as the mass of infectious goes to zero.

PROOF:

Use the same argument as above. In the limit

$$0 \leq \lim_{I \rightarrow 0} \frac{c(r(I))}{\left(r(I) + \frac{X(I)}{I} - r^0\right) I} + \frac{v}{\xi_0 - r^0} \leq \frac{v}{\xi_0 - r^0}, \tag{52}$$

which proves the result.

qed.

Lemma 8. .

If $\xi_0 > r^0$, the total cost of suppression is bounded at the optimum.

PROOF:

The total cost of suppression at the optimum is smaller or equal to the total cost of suppression under any other policy that satisfies $\dot{I}(I) < 0$. In particular, take the policy $\tilde{r}(I) = 0$ for $I \leq I^*/2$ and $\tilde{r}(I) = r^0$ for $I > I^*/2$. It follows that

$$\int_0^{I_0} \frac{c(r(I)) + vI}{\left(r(I) + \frac{X(I)}{I} - r^0\right) I} dI \leq \int_0^{I^*/2} \frac{v}{\frac{X(I)}{I} - r^0} dI + \int_{I^*/2}^{I_0} \frac{c(r^0) + vI}{X(I)} dI \quad (53)$$

Both integrals exist, which gives the result.

qed.

Lemma 9. .

If $\xi_0 < r^0$, and the cost function is quadratic, the economic unit cost of suppression goes to infinity and the social unit cost goes to a constant as the mass of infectious goes to zero.

PROOF:

Take the definition of the total unit cost and take the limit. Use the results from Lemma 4:

$$\lim_{I \rightarrow 0} \frac{c(r(I))}{\left(r(I) + \frac{X(I)}{I} - r^0\right) I} + \lim_{I \rightarrow 0} \frac{v}{\left(r(I) + \frac{X(I)}{I} - r^0\right) I} = \frac{c(2(r^0 - \xi_0))}{r^0 - \xi_0} \lim_{I \rightarrow 0} \frac{1}{I} + \frac{v}{r^0 - \xi_0}. \quad (54)$$

qed.

Lemma 10. .

If $\xi_0 < r^0$, and the cost function is quadratic, the total cost of suppression is infinite even at the optimum.

PROOF:

Take the expression for the total cost and take the optimality condition to get

$$C = \int_0^{I_0} \frac{c(r(I)) + vI}{\left(r(I) + \frac{X(I)}{I} - r^0\right) I} dI = \int_0^{I_0} \frac{c'(r(I))}{I} dI \quad (55)$$

The optimal policy is increasing and larger than zero in zero; therefore

$$C \geq \int_0^{I_0} \frac{c'(r(0))}{I} dI = \infty \quad (56)$$

qed.

Exposure to the Covid-19 stock market crash and its effect on household expectations¹

Tobin Hanspal,² Annika Weber³ and Johannes Wohlfart⁴

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

We survey a representative sample of US households to study how exposure to the Covid-19 stock market crash affects expectations and planned behavior. Wealth shocks are associated with upward adjustments of expectations about retirement age, desired working hours, and household debt, but have only small effects on expected spending. We provide correlational and experimental evidence that beliefs about the duration of the stock market recovery shape households' expectations about their own wealth and their planned investment decisions and labor market activity. Our findings shed light on the implications of household exposure to stock market crashes for expectation formation.

- 1 This paper was previously circulated under the title "Income and Wealth Shocks and Expectations During the COVID-19 Pandemic". We are grateful for helpful comments from Carola Binder, Francesco D'Acunto, Michalis Haliassos, Claus Kreiner, Christine Laudenbach, Chris Roth, Sonja Settele, Michael Weber, Rudiger Weber and seminar participants at Copenhagen and St. Gallen. We thank the Economic Policy Research Network (EPRN) for financial support. The activities of the Center for Economic Behavior and Inequality (CEBI) are funded by the Danish National Research Foundation. We received ethics approval from Goethe University Frankfurt. The survey instructions and the online appendix can be found at the following link: <https://sites.google.com/site/tobinhanspal/survey>.
- 2 Department of Finance, Accounting and Statistics, WU Vienna University of Economics and Business.
- 3 Department of Finance, Goethe University Frankfurt.
- 4 Department of Economics and CEBI, University of Copenhagen, Danish Finance Institute.

Copyright: Tobin Hanspal, Annika Weber and Johannes Wohlfart

1 Introduction

A major part of the wealth of households in the United States and in other industrialized countries is invested in the stock market. While historically investing in the stock market has provided a premium over the long run, it exposes households' savings to volatility and to the risk of market crashes. Most recently, the spread of the COVID-19 coronavirus pandemic and the policy measures put in place to contain the virus have sent major stock markets around the world plummeting, with the S&P500 losing an unprecedented third of its value during the sharp drop of stock prices in February and March 2020. How do households adjust their plans about spending, investments, and labor supply in response to wealth losses during such a crash? And how do beliefs about the stock market recovery affect individuals' expectations about their own wealth and plans? Answering these questions is crucial for understanding the implications of households' exposure to stock market crashes for the vulnerability of different socioeconomic groups and for the formation of household expectations, which are central to economic models and important determinants of household behavior (Armona et al., 2018; Bachmann et al., 2015; Bailey et al., 2017; Coibion et al., 2019a, 2020b; D'Acunto et al., 2019a; Giglio et al., 2020a; Kuchler and Zafar, 2019).

In this paper we shed light on these issues using a survey on a sample of more than 8,000 US households, representative in terms of age, gender, income, and region, which we conducted in April 2020. We elicit the value of participants' wealth holdings in retirement accounts and in other financial accounts as of January 2020, as well as the capital losses they incurred in those accounts as a result of the drop in stock prices. We then measure respondents' expectations regarding the stock market and their own financial prospects, and elicit their planned decisions with respect to stock investments, spending and labor supply. The survey includes an experimental section, in which random subsets of respondents receive information on the duration of the recovery in the case of a historical stock market crash (the Black Monday crash in 1987, the burst of the Dot-com bubble in 2000, or the 2007-2009 Financial Crisis). These treatments generate

exogenous variation in our respondents' expectations about the recovery of the stock market from the current crash. Our survey allows us to study how exposed households adjust their plans about investment, spending, debt and labor supply in response to a stock market crash, and how beliefs about the recovery causally shape these plans and people's expectations about their own household wealth. At the same time, our survey offers a comprehensive real-time snapshot of household finances and expectations during the COVID-19 pandemic in the US.¹

We start by quantifying the exposure of different groups of the population to the February/March 2020 stock market crash. US households report median financial wealth losses of \$1,750 and mean losses of \$30,415 at the time of our survey in early April 2020. Relative losses of financial wealth strongly increase in net wealth and income, and are strongest for those in middle age. These differences can largely be explained by differences in the share invested in stocks before the onset of the crisis. Across groups, wealth shocks tend to be negatively correlated with household income shocks experienced during the early stages of the pandemic, which are strongest among the poorest and younger households and almost zero for those with high incomes or wealth and for older households. Wealth shocks due to the stock market crash therefore counteract the role of income shocks in the effect of the pandemic on overall inequality of available economic resources.

How did households adjust their decisions and plans regarding investment, spending, household debt and labor supply in the medium-term in response to the pandemic more generally and to wealth shocks in particular? About 50 percent of households who were invested in the stock market at the onset of the crisis made active adjustments to their stock investments since the beginning of the crash, with about equal shares of respondents increasing and decreasing the stock share in their overall financial wealth. Thus, households did not exhibit a systematic tendency to rebalance their portfolios in response to the decrease in their stock share due to the crash. Moreover, 36 percent of respondents

¹By contrast, data from long-running surveys such as the PSID or from some administrative data sources may only become available with a lag, potentially extended by the shutdown of large parts of the economy, society and public administration.

report that the coronavirus crisis increases their expectations about household debt at the end of 2020 and 44 and 53 percent report that the crisis increases their expectations about their retirement age and desired working hours over the coming years.

Shocks to stock wealth inside and outside of retirement accounts are strongly correlated with upward adjustments in expected desired working hours and retirement age. A ten percent shock to retirement financial wealth is associated with a four percentage points higher tendency to report upward adjustments in retirement age. This suggests that households plan to make up for losses experienced during a crash by increasing labor supply, in line with a key mechanism in portfolio choice models with human capital (Bodie et al., 1992; Gollier, 2002). We find evidence of only small changes in expected household spending in response to wealth shocks, with a \$1 shock to retirement financial wealth being associated with a \$0.02 reduction in spending. By contrast, income shocks experienced during the pandemic have strong effects on expected spending, with an average reduction in expected spending in 2020 of \$0.45 for each \$1 shock to income. This is consistent with the view that retirement wealth holdings are less liquid and not used to finance current spending. In addition, households hit by wealth shocks tend to be better insured against shocks due to higher savings and easier access to credit.

We next turn to the role of households' beliefs about the further development of the stock market. Respondents who personally experienced losses during past crashes, Democrats, and women expect the stock market to take more time to recover to pre-crisis levels and expect significantly lower returns over the coming year.²

Finally, we exploit the information experiment embedded in our survey to examine the causal effects of people's expectations about the stock market recovery on their economic outlook for their own household and their planned economic decisions. When respondents are provided with information on the duration of a longer (shorter) historical stock market crash, this causes them to be more pessimistic (optimistic) about the development of the stock market in the coming years compared to respondents in control groups who

²This is in line with previous literature highlighting the importance of these factors during more tranquil times (D'Acunto et al., 2020; Kuchler and Zafar, 2019; Malmendier and Nagel, 2011).

have not received information. This suggests that households had not been fully informed about historical facts they consider relevant for the further development of the stock market, pointing to a role for information frictions in households' stock market expectations (Abel et al., 2007; Alvarez et al., 2012). Moreover, respondents update their expectations about their own wealth, their investment plans, and their long-term labor market activity in response to the information. We also find strong correlations between expected recovery duration and these outcomes in OLS regressions. These findings suggest that, next to incurred wealth shocks, expectations about the stock market going forward play an important causal role in shaping households' outlook regarding their own wealth and decisions.

We contribute to a literature studying the formation of households' subjective stock market expectations and their association with economic choices (Ameriks et al., 2019; Amromin and Sharpe, 2014; Das et al., 2017; Dominitz and Manski, 2007; Giglio et al., 2020a; Greenwood and Shleifer, 2014; Malmendier and Nagel, 2011; Vissing-Jorgensen, 2003). Giglio et al. (2020b) document that investor beliefs about the 1-year ahead stock market return declined following the February-March 2020 stock market crash, while expectations over the 10-year horizon remained stable. Guiso et al. (2018) and Weber et al. (2013) study the development of beliefs and risk-taking following the Financial Crisis 2008. We contribute to this literature by providing evidence on how stock market expectations affect individuals' economic outlook and plans in both financial and non-financial domains following a crash. Methodologically, we add to the literature on subjective stock market expectations by applying an information experiment. Such experiments have previously been used to study household expectations about inflation (Armantier et al., 2016; Binder and Rodrigue, 2018; Cavallo et al., 2017; Coibion et al., 2020a, 2019b, 2018), house prices (Armona et al., 2018; Fuster et al., 2019) and GDP growth (Roth and Wohlfart, 2019). To the best of our knowledge, our results provide the first direct *causal* evidence on the role of subjective return expectations in shaping individuals' planned stock investment behavior in a real-world setting, and the first evidence on the role of financial market expectations in shaping households' long-term plans about labor market activity.

Our study also adds to previous work studying households' responses to changes in their stock market wealth. Several studies document the finding that households are unlikely to actively rebalance their portfolios to counteract passive changes to their portfolio allocation (Brunnermeier and Nagel, 2008; Calvet et al., 2009). Di Maggio et al. (2019) and Bräuer et al. (2020) estimate small marginal propensities to consume (MPC) out of passive changes in households' stock wealth.³ We provide real-time evidence on how stock wealth losses during a crash affect plans about trading, spending, debt as well as long-term labor market activity.

Finally, we contribute to a rapidly expanding literature on the economic and financial consequences of the spread of the coronavirus. Coibion et al. (2020b) study the effect of lockdowns on households' beliefs about inflation, unemployment, and mortgage interest rates, as well as their consumer spending. Binder (2020) examines how beliefs about inflation and unemployment respond to information about the Fed's interest rate response to the coronavirus crisis. Fetzner et al. (2020) study how perceptions of pandemic risk factors shape people's economic sentiment. Gormsen and Koijen (2020) use data on the aggregate equity market and dividend futures to quantify how investors' expectations about economic growth evolve in response to the outbreak of the virus and subsequent policy responses. Relatedly, Dietrich et al. (2020) provide survey evidence on households' perception of the effect of the coronavirus on US GDP growth. Bu et al. (2020) document a sharp decrease in risk-taking stemming from changes in beliefs after the onset of the coronavirus pandemic among survey respondents in China. Others study the impact of the coronavirus shock on labor markets (Adams-Prassl et al., 2020; Bick and Blandin, 2020; Coibion et al., 2020c) and on consumer spending (Andersen et al., 2020; Baker et al., 2020; Cox et al., 2020). We contribute to this literature by providing the first evidence on how financial wealth shocks during the coronavirus crisis affect households' medium-term plans about investment, spending, debt and labor market activity.

The remainder of the paper is structured as follows. In section 2 we describe the

³We also relate to a literature making use of survey data on subjective beliefs to study the consumption response to changes in economic resources more generally (Christelis et al., 2019; Fuster et al., 2020; Jappelli and Padula, 2015).

survey and the sample. Section 3 provides descriptive evidence on the distribution of wealth losses across the population. In section 4 we examine how wealth shocks affect households' economic decisions and plans. In section 5 we provide correlational and experimental analysis of the role of people's expectations about the further development of the stock market in their economic plans. Section 6 discusses implications of our findings and concludes.

2 Survey design and data

In this section we provide details on the structure and administration of our survey, as well as the characteristics of our sample.⁴

2.1 Survey design

Our survey starts with a set of questions on demographics such as age, gender and household income. The respondents then answer questions on the value of i) their retirement accounts and ii) the value of all financial assets they held outside of their retirement accounts at the end of January 2020. We ask them explicitly to think of the value of their assets before the start of the current crisis. To ease cognitive strain we ask our respondents to indicate the brackets into which the values of their assets fell instead of asking them for exact estimates.⁵ Respondents then report the percent shares of financial assets in retirement accounts and of financial assets in other accounts that were invested in stocks or stock mutual funds at the end of January. Finally, they estimate by what percent the value of their retirement accounts and the value of their other financial accounts changed as a result of the stock market developments since the beginning of the crisis until the day of the survey. The survey continues with questions on whether respondents lost their job since the beginning of the year, and whether their net household income in the first quarter of 2020 was higher or lower than they had expected before the crisis,

⁴The wording of the survey questions is available at <https://sites.google.com/site/tobinhanspal/survey>

⁵One concern might be that individuals are imperfectly informed about their retirement wealth. This concern is arguably mitigated by the fact that we conducted our survey at the beginning of April. Pension plan providers usually send out wealth statements to clients on a quarterly basis, so respondents should have received at least one such statement in the weeks prior to the survey. Moreover, plans about spending, investment and labor supply should be affected by *perceived* shocks to respondents' wealth, which is what we measure in our survey.

and by what percent it was higher or lower.

Respondents then proceed to the short experimental part of the survey. They are randomly allocated into one of seven groups. Respondents in arms *FinCrisisInfo* and *FinCrisisControl* are asked to estimate the number of years it took the stock market to recover from the drop during the Financial Crisis in 2007. Only respondents in arm *FinCrisisInfo* are then provided with the actual number of years it took the stock market to reach its pre-crisis peak (5 1/2 years). Similarly, respondents in arms *DotComInfo* and *DotComControl* and in arms *BlackMondayInfo* and *BlackMondayControl* report prior estimates and respondents in the respective treatment arm receive information on the recovery duration from the burst of the Dot-com bubble in 2000 (7 years) and the Black Monday stock market crash in 1987 (2 years), respectively.⁶ Although asking respondents to estimate the number of years could have framing effects, such framing effects would likely occur with any method of eliciting these beliefs, and such effects should be common across treatment arms. Finally, respondents in the *PureControl* arm are not shown any questions on priors or information and immediately proceed to the next part of the survey. Online appendix Table A1 provides an overview of the treatment and control arms.

Next, all respondents report their beliefs about the recovery of the US stock market. They report the calendar year in which they expect the stock market to recover to its January 2020 level, as well as their agreement on three qualitative statements on the severity of the recent drop in stock prices on 7-point scales.⁷ Respondents are also asked in which year they expect their own household's net wealth to recover to its pre-crisis level, including an option that their net wealth will never recover. Finally, the respondents allocate probabilities across eight intervals into which the US stock market return over the next 12 months might fall, which are mutually exclusive and collectively exhaustive.⁸

⁶The information treatments included a dynamic figure contrasting the respondent's prior belief with the information. Online appendix Figure A2 plots an example survey screen for the *FinCrisisInfo* information treatment.

⁷Specifically, respondents are asked to what extent they agree or disagree with the following statements: "The outbreak of the coronavirus will keep US stock prices below their January 2020 levels for many years."; "The outbreak of the coronavirus has set the level of the stock market back by many years."; "The US stock market will have recovered by the end of the year 2020."

⁸Specifically, respondents report the percent chance they assign to each of the following brackets of aggregate stock returns: less than -30 percent, between -30 and -15 percent, between -15 and -5 percent, between -5 and 0 percent, between 0 and 5 percent, between 5 and 15 percent, between 15 and 30 percent,

The survey continues with a set of questions on respondents' expectations about their own economic and financial situation as well as their decisions. Specifically, respondents answer a qualitative question on the financial prospects of their household, and questions on whether they expect the total spending and the total net income of their household to be higher or lower in 2020 as compared to 2019, and by what percent they expect it to be higher or lower. Those who report an expected reduction in their household income also forecast the year in which they expect their household income to have recovered. The participants then respond to qualitative questions on whether the current crisis affects their expectations about their retirement age, their desired working hours in the next years, as well as their outstanding household debt at the end of 2020. Finally, those who held any equity in the beginning of 2020 are asked whether they have made any active adjustments to the share of their financial assets invested in stocks or stock mutual funds, and whether they plan to do so over the next weeks. The survey ends with additional background questions on topics such as stock investment experience or the value of real estate and debt holdings at the beginning of the year.

Our design with seven survey arms has important advantages. On the one hand, we can study the causal effect of information about past crashes on expectations and plans by comparing individuals who have reported priors and received information about a particular crash with those who only have reported priors (e.g. comparing the *FinCrisisInfo* and *FinCrisisControl* arms). On the other hand, we can use the pure control group, who has not received questions or information on past crashes, to provide descriptive evidence that is not affected by drawing people's attention to past crashes. Throughout the analysis, all descriptive figures on survey questions asked after the experimental stage are restricted to the pure control group. All non-experimental regressions using such questions as outcomes restrict the sample to the four control arms to increase power. In the appendix we show versions of these tables using only the pure control group, however.

greater than 30 percent.

2.2 Data

Survey administration We collaborated with the survey company *Lucid*, which is widely used in economic and financial research. The survey was conducted between 6th and 13th April 2020. The US stock market had partially recovered at the time of the survey, but still showed drastic losses of close to 20 percent compared to its pre-crisis level, and the number of initial jobless claims had escalated (Figure A1).⁹ Participants were recruited from the provider's online panel and then completed the survey on our own platform. They proceeded to the main survey after initial screening according to demographics in order to achieve representativeness in terms of observables.¹⁰ In total, 8,156 respondents completed our survey. We drop 162 respondents in the top and bottom percentiles of the response time, as very short or very long response times may indicate inattention to the survey. We also remove 547 respondents who refused to answer any of our questions on financial wealth holdings, as these questions are used extensively throughout the analysis. This leaves us with a sample of 7,447 respondents, who completed the survey within 16.6 minutes on average (13.7 minutes at the median).

Sample characteristics Table 1 shows summary statistics of our sample, including a comparison with targets from the 2018 American Community Survey (ACS). The composition of our sample is close to the general population in terms of gender (52 percent females vs 51 percent in the ACS), mean age (48.3 years compared to 47.6 years in the ACS) and median gross household income in the previous year (\$62,500 vs \$65,700 in the ACS), as well as Census region of residence. As it is common in online samples, a slightly larger fraction of our respondents have a Bachelor's degree compared to the general population (38 percent in our sample vs 31 percent in the ACS).

Integrity of the randomization Our sample is well-balanced across the seven arms of the experimental part of the survey for a set of key characteristics (see Table A2). To rule

⁹An advantage of the survey date is that it should give a more accurate picture of the longer-term wealth shocks due to the crash going beyond the very drastic short-term effects as of mid-March. We cannot meaningfully exploit variation within the one-week survey period, given that different population groups were targeted in the course of the week in order to achieve a representative sample.

¹⁰Respondents received a small reward for participating in the survey.

Table 1: Descriptive statistics

	Our sample					ACS		
	Mean	SD	P10	Median	P90	Obs.	Mean	Median
Female	0.51	0.50	0	1	1	7,447	0.51	
Age (years)	48.64	16.29	26	49	70	7,447	47.60	
- 18-24 years (d)	0.10	0.30	0	0	0	7,447	0.12	
- 25-34 years (d)	0.16	0.36	0	0	1	7,447	0.18	
- 35-44 years (d)	0.19	0.39	0	0	1	7,447	0.16	
- 45-54 years (d)	0.17	0.38	0	0	1	7,447	0.16	
- 55-65 years (d)	0.20	0.40	0	0	1	7,447	0.18	
- 65 years and older (d)	0.19	0.39	0	0	1	7,447	0.19	
Bachelor's degree or higher (d)	0.39	0.49	0	0	1	7,447	0.30	
Some college (d)	0.40	0.49	0	0	1	7,447		
High school (d)	0.19	0.39	0	0	1	7,447		
Married (d)	0.53	0.50	0	1	1	7,447		
Separated (d)	0.13	0.34	0	0	1	7,447		
Widowed (d)	0.05	0.21	0	0	0	7,447		
Household income (gross, USD)	80,952	57,246	20,000	62,500	175,000	7,417	91,673	65,700
- <15,000 (d)	0.09	0.29	0	0	0	7,417	0.10	
- 15,000-25,000 (d)	0.08	0.28	0	0	0	7,417	0.08	
- 25,000-50,000 (d)	0.21	0.41	0	0	1	7,417	0.20	
- 50,000-75,000 (d)	0.17	0.38	0	0	1	7,417	0.17	
- 75,000-100,000 (d)	0.14	0.35	0	0	1	7,417	0.13	
- 100,000-150,000 (d)	0.17	0.38	0	0	1	7,417	0.16	
- 150,000-200,000 (d)	0.07	0.26	0	0	0	7,417	0.07	
- >200,000 (d)	0.05	0.23	0	0	0	7,417	0.09	
Census Region (d)								
- West	0.22	0.41	0	0	1	7,447	0.24	
- Midwest	0.25	0.43	0	0	1	7,447	0.21	
- Northeast	0.21	0.41	0	0	1	7,447	0.18	
- South	0.33	0.47	0	0	1	7,447	0.38	
Democrat	0.39	0.49	0	0	1	7,447		
Republican	0.37	0.48	0	0	1	7,447		
Employment situation (d)								
- Employed	0.53	0.50	0	1	1	7,447	0.62 ¹⁾	
- Self-employed	0.07	0.25	0	0	0	7,447		
- Unemployed	0.07	0.25	0	0	0	7,447	0.03	
- Out of labor force	0.33	0.47	0	0	1	7,447	0.35	
Retired	0.21	0.41	0	0	1	7,447		
Other	0.12	0.33	0	0	1	7,447		
Financial wealth (USD)	208,528	315,752	0	55,000	662,500	7,447		
Retirement wealth (USD)	120,738	180,658	0	17,500	575,000	7,447		
Other financial wealth (USD)	87,790	159,373	0	7,500	325,000	7,447		
Real estate wealth (USD)	223,761	338,990	0	150,000	625,000	7,447		
Debt outstanding (USD)	70,827	137,519	0	7,500	250,000	7,447		
Household net wealth (USD)	357,326	536,009	-11,500	132,500	1,150,000	7,447		
Stock investor (d)	0.61	0.49	0	1	1	7,447		
% Equity in fin. wealth (%)	39.14	33.23	0	40	91	5,835		
Inv. experience > 10 yrs. (d)	0.55	0.50	0	1	1	4,567		
Credit constrained (1-5)	2.53	1.45	1	2	5	7,447		

Notes: The table shows summary statistics for the 7,447 respondents in the final sample. Stock market experience is elicited for stock investors only. The share of equity in total financial assets is conditional on positive financial asset holdings as of January 2020. Observation numbers for some wealth items vary due to item non-response. ¹⁾ includes self-employed.

out any concerns, we include a set of control variables not only in our non-experimental but also in our experimental estimations.

Variable definitions The survey elicits levels of household income, assets and liabilities by asking respondents to indicate the respective value bin. Shocks to households' financial wealth and net income during the first quarter of 2020 as well as expected differences in household net income and spending in 2020 compared to 2019 are elicited as numerical entries in percentage terms. In order to reduce the impact of outliers in these variables in our analysis, for each variable, we set the top and bottom 2 percent of the distribution to missing. When calculating changes in financial wealth components and income in dollar terms, we first translate percentage changes into dollar terms by multiplying respondents' reported percentage changes and base levels, and then trim the top and bottom 2 percent of the resulting dollar distribution, respectively.¹¹ Finally, all dummy outcomes in our regressions are coded as either 0 or 100 in order to bring them on the same scale as independent variables referring to percent changes.

3 Descriptive evidence: Exposure to the COVID-19 stock market crash

The main goal of our survey was to examine how exposure to a stock market crash affects households' expectations about investment, spending and labor supply. However, our survey also offers a comprehensive real-time snapshot of the financial situation of households in the US during the early stages of the COVID-19 pandemic. In this section we describe how wealth shocks from the COVID-19 stock market crash are distributed across the population and how they are correlated with income shocks.

Unconditional wealth shocks Panel A of Figure 1 displays the average unconditional percent change (top row) in the value of household financial assets across groups, where those with no financial assets are coded as having experienced a shock of zero. Wealth losses due to the stock market crash are strongly increasing along the net wealth dis-

¹¹Our results are not sensitive to the exact choice of the cutoff. Results are available upon request.

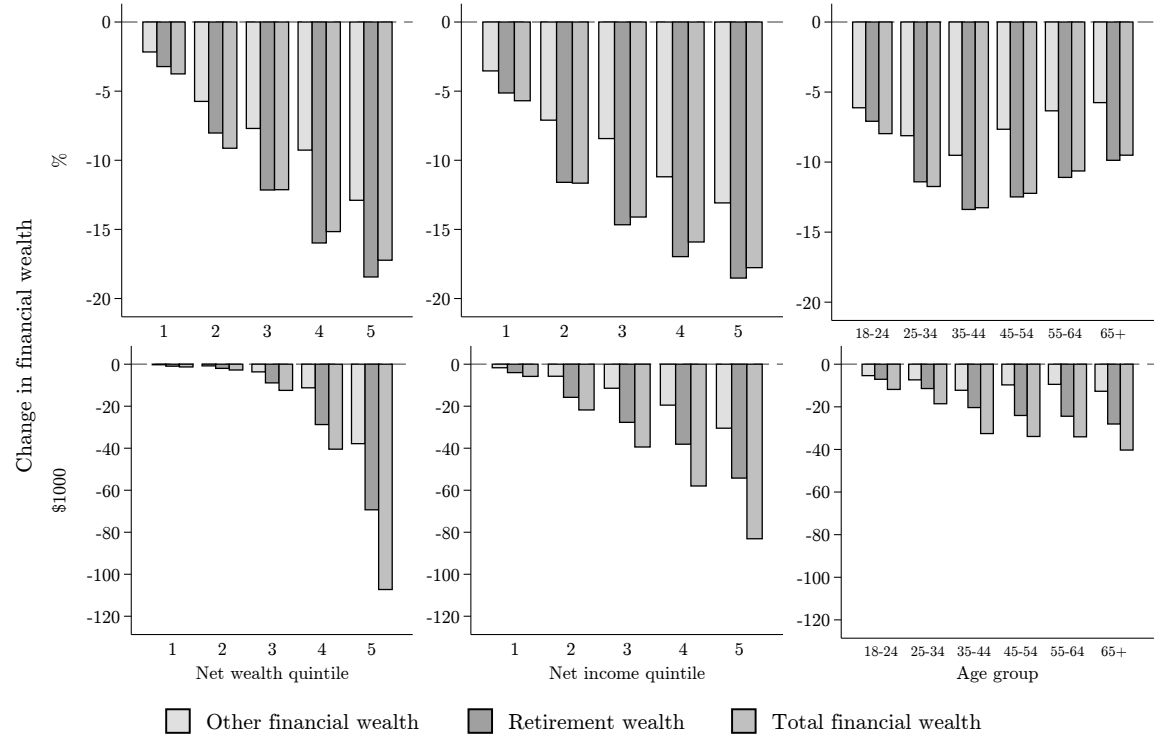
tribution (left column), with overall financial losses amounting to 4 percent of pre-crisis financial wealth in the lowest quintile and to about 17 percent in the highest quintile. There is a similar gradient of wealth losses along the pre-crisis net income distribution (middle column). The distribution of dollar losses (bottom row) is naturally much more skewed along the net wealth and income distribution, reflecting the strong inequality in financial asset holdings across groups (as shown in Figure A3). Unconditional wealth losses in dollar terms amount to \$30,415 at the mean and \$1,750 at the median, and average \$1,311 in the lowest and \$107,275 in the highest net wealth quintile.

The right column of Figure 1 displays unconditional capital losses by age group. Percent changes in financial wealth are most pronounced for those aged between 25 and 54 (net capital losses of between 13 and 14 percent), and are markedly lower for younger individuals (8 percent) and for older individuals (11 percent for those aged 55-64 and 10 percent for those above 65). Wealth shocks in dollar terms increase in age, reflecting increasing wealth accumulation over people's working life. Across net wealth, income and age groups, both absolute and percent losses are larger for holdings in retirement accounts (e.g., 401Ks or IRAs) than for holdings outside of retirement accounts, largely reflecting higher wealth (see Figure A3) and higher stock shares (see Figure A4) inside retirement accounts.

Figure A5 shows that unconditional wealth shocks are strongly increasing in educational attainment, are stronger for men, and less pronounced for those retired or part-time employed compared to those full-time employed as of January 2020. Given that we ask about household wealth, one should interpret these patterns according to individual-level characteristics with caution.


Figure 1: Wealth and income shocks across groups

Panel A: Financial wealth shocks



Covid Economics 23, 28 May 2020: 188-262

Panel B: Household net income shocks



Notes: This figure displays the change in the value of financial assets due to the February/March 2020 stock market drop until the survey date in percentage terms and in USD (Panel A) and unexpected changes in net household incomes during the first quarter of 2020 in percentages and USD (Panel B), by quintile of the pre-crisis net wealth distribution (left column), by quintile of the pre-crisis net income distribution (middle column) and by age group (right column). Changes in the value of household financial assets are displayed separately for financial assets outside of retirement accounts (other financial wealth), for financial assets in retirement accounts, and for the combined value of financial assets. Changes in value of financial assets are net capital losses for the majority of respondents, and net capital gains for a small fraction of respondents. We trim reported shocks to income and financial wealth at the 2nd and 98th percentiles. The sample is the full sample without missings in the relevant survey questions. Note that the average percent reduction in overall financial wealth can be larger than both average percent reductions for the individual components. This is due to the fact that we coded those with no wealth in a given category as having experienced a shock of zero percent in that category. These cases occur particularly in groups with lower wealth holdings.

Covid Economics 23, 28 May 2020: 188-262

COVID ECONOMICS
 REAL TIME PAPERS

Conditional wealth shocks The patterns in the distribution of unconditional losses in financial wealth reflect differences across groups in i) the fractions of households with no financial wealth before the crisis, who did not incur any losses, ii) the stock share in financial assets, which differs substantially across groups (see Figure A4), iii) the types of risky assets households invest in, or iv) the tendency to realize losses across groups. While our survey data are not granular enough to address iii), we explore i) and ii) in more detail by studying conditional wealth losses across groups. In subsection 4.1 we also address iv) by studying active adjustments to stockholdings across groups.

Figure A6 reproduces Figure 1 for the sample of households who report positive holdings of financial assets inside or outside of retirement accounts as of January 2020.¹² As before, percent financial losses are increasing in net wealth and income, and are hump-shaped in age. However, the patterns are substantially less pronounced than before. Due to differences in the value of financial assets, patterns of dollar changes in wealth across net wealth and income groups remain largely unchanged. Figure A7 restricts the sample further to households investing in stocks or stock mutual funds as of January 2020, which makes the patterns in percent losses across income and wealth groups almost uniform, while the age pattern remains. The last column of Figure A7 highlights that wealth losses are strongly increasing in the pre-crisis portfolio equity share. Finally, Figure A8 plots losses by equity share bin separately for different groups. Mean losses conditional on portfolio equity share are almost equal across groups. This highlights that conditional on holding positive financial wealth, differences in capital losses largely seem to be due to differences in portfolio shares invested in stocks and stock mutual funds. This becomes particularly evident when we compare mean experienced losses between the highest (5) and lowest (1) quintiles. As shown in Table A3, the difference in unconditional wealth losses between the most and least wealthy households of 13 percentage points shrinks to 7 percentage points if we condition on holding positive financial wealth, and declines further to 1.4 percentage points among those with positive investment in equities.

¹²The values in the figure are conditional on positive overall financial wealth holdings and thereby hold the sample fixed across the three bars. The patterns look similar if we condition on positive retirement and non-retirement financial wealth holdings separately.

Income shocks How does the distribution of wealth shocks across groups compare with the distribution of income shocks during the early stages of the pandemic? Figure 1 Panel B displays average shocks to household net income in the first quarter of 2020 across groups. We calculate these shocks based on a survey question asking respondents by what percent their household income in the first quarter was higher or lower than they had expected before the crisis. Strikingly, income shocks exhibit the opposite pattern compared to wealth shocks, with income losses being strongest for households in the bottom net wealth or income quintiles (7 and 6 percent respectively), and gradually becoming less severe, with those in the highest quintiles losing almost no income (top row). There is also a strong age gradient in income losses, with younger households being affected most severely and older households being more insulated. We convert these shocks into dollar changes using the approximation that, before the onset of the crisis, households had expected a quarter of their 2019 income for the first quarter of 2020 (bottom row). Unconditionally, respondents report to have lost \$844 of net household income over the first quarter 2020, on average. Income losses average \$536 in the lowest and \$1,047 in the highest net income quintile. While the income and wealth gradients naturally reverse, the age pattern remains fairly similar as for relative income losses.

Figure A9 provides evidence on the distribution of job losses, the main drivers of shocks to household incomes, among the respondents in our sample. A striking 26% of our respondents report that they have lost their job from January 2020 until the time of our survey in early April, in line with other recent evidence (Adams-Prassl et al., 2020; Bick and Blandin, 2020; Coibion et al., 2020c). Job losses broadly follow the same patterns as income losses. They are more prevalent for lower net wealth, for lower income and for younger households. Women, individuals with lower education, and part-time workers are also more likely to have lost their jobs.

Taken together, these findings imply that income losses and wealth losses during the early stages of the COVID-19 pandemic tend to be negatively correlated across groups. More generally, households' exposure to stock market crashes is concentrated among groups of households who tend to be less affected by income shocks during recessions

(Hoynes et al., 2012). This is in line with earlier findings documenting that wealth inequality tends to decline during recessions, at least in the short run (Kuhn et al., 2019).¹³

Result 1. *Shocks to financial wealth due to the COVID-19 stock market crash are strongly increasing in net wealth and income, and strongest for those in middle age. These patterns are largely due to differences in the fraction of households with no financial assets and in the stock share in financial assets across groups. Wealth shocks tend to be negatively correlated with income shocks during the early stage of the pandemic across groups.*

4 Main results: Effects of the COVID-19 stock market crash on behavior, expectations and plans

In the previous section we have explored how different groups of households were affected by the February/March 2020 stock market crash. We now turn to our main findings on how households adjust their plans about investment, spending, debt and labor supply in response to the pandemic in general, and to wealth shocks in particular.

4.1 Changes in risk-taking across groups

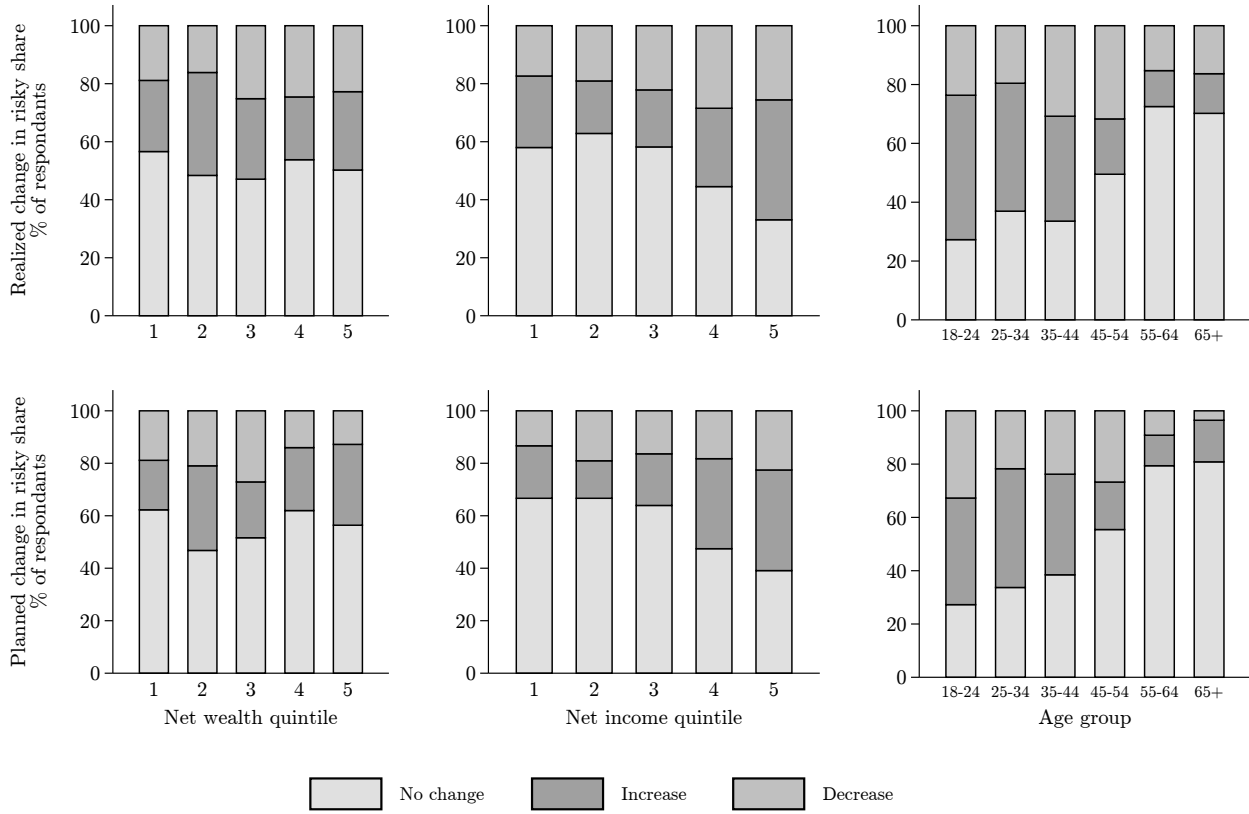
Which groups make adjustments to the share of financial assets invested in stocks or stock mutual funds in response to the crash? Our survey asks all respondents who report positive stockholdings as of January 2020 whether they have actively increased or decreased their overall portfolio share invested in equities (combining retirement and other accounts) since the onset of the crisis, and whether they plan to make active adjustments in the weeks following the survey. The wording of these questions is such that respondents should abstract from passive changes to the equity share due to changes in market prices.

Figure 2 plots the fractions of stockholders who have made active adjustments to their portfolio equity share as a result of the coronavirus crisis across demographic groups

¹³Given the concentration of stock ownership in households at the top of the distribution, decreasing stock prices tend to reduce wealth inequality across households. However, differences in the relative speed of recovery of equity and housing markets led to a spike in wealth inequality in the aftermath to the 2007-2009 Financial Crisis (Kuhn et al., 2019).

(top row). Approximately 50 percent of pre-crisis stockholders have made no active adjustments to the share of their wealth invested in stocks since the onset of the crisis, in line with earlier evidence showing that many households do not rebalance passive changes in their asset allocation (Brunnermeier and Nagel, 2008; Calvet et al., 2009). The remaining stockholders were slightly more likely to actively increase (27.9 percent) than to decrease (22.8 percent) their portfolio share in equities. Households from lower wealth and income quintiles and those in older age groups are less likely to have made active changes to their portfolio. Interestingly, while the tendency to realize sales was rather uniform across groups, households higher up in the income distribution and those in younger age cohorts were more likely to actively increase their exposure to the stock market. The bottom row of Figure 2 shows that *planned* active changes in risk-taking over the next few weeks exhibit very similar patterns as realized adjustments in risk-taking.

Figure 2: Realized and planned adjustments to stock share across groups



Notes: The top row of this figure displays the fractions of pre-crisis stockholders in different groups reporting that they made no active change, actively increased, or actively decreased the share of their overall financial assets (including retirement and non-retirement accounts) that is invested in the stock market since the beginning of the crisis, while the bottom row plots the percent of respondents who stated that they are planning to make no change, increase, or decrease their investment in the following weeks. The fractions are plotted by quintile of the pre-crisis net wealth distribution (left), quintile of the pre-crisis net income distribution (middle), and age group (right). The sample consists of all pre-crisis stock investors in the pure control group, who have not received any questions or information on past crashes before answering to the questions on investment behavior.

Covid Economics 23, 28 May 2020: 188-262

What drives households' tendency to make adjustments to the share of their portfolio held in equities in the context of a stock market crash? In Table 2 we regress dummy variables indicating realized or planned active changes in risk-taking on a set of covariates.¹⁴ Stronger negative income shocks are associated with a stronger tendency to reduce stock investments, potentially due to liquidity needs. By contrast, larger financial losses are associated with a greater likelihood to plan to increase the portfolio equity share, consistent with portfolio rebalancing (Calvet et al., 2009) or a tendency to make up for paper losses due to loss aversion (Imas, 2016).¹⁵ Respondents who held a higher share of their wealth inside retirement accounts as of January 2020 are less likely to adjust their risk-taking, in line with stronger inertia in retirement accounts (Agnew et al., 2003; Ameriks and Zeldes, 2004; Biliias et al., 2010; Madrian and Shea, 2001). Having made losses in the stock market during the Financial Crisis 2007-9 is associated with a substantially higher tendency to plan and realize sales during the February/March 2020 crash, and a lower tendency to plan and realize purchases. The patterns are less pronounced for experiences during the earlier stock market crashes following the burst of the Dot-com bubble in 2000 or the Black Monday in 1987. These findings are in line with recency bias documented by the literature on the role of experiences in financial risk-taking (Andersen et al., 2019; Laudenbach et al., 2020; Malmendier and Nagel, 2011), and suggest that losing wealth during a stock market crash may have the negative long-run consequence of a greater tendency to sell stocks following market downturns. Finally, men are more

¹⁴Our baseline set of controls includes the respondent's gender, age category, dummies for being married, separated or divorced, or widowed (single being the omitted category), dummies for highest educational attainment of highschool, some college or associate degree, or college degree or higher (below highschool being omitted), dummies for being self-employed, retired, unemployed or other labor market status (in paid employment omitted), a dummy for being the main earner in the household, a z-scored measure of the extent to which the respondent is involved in financial decision-making in the household, dummies for Republicans and for other party affiliation (Democrat being omitted), the logs of net household income, of financial wealth inside and outside of retirement accounts, of all real estate wealth, and of total household debt, a z-scored measure of perceived borrowing constraints, the share of financial wealth invested in stocks and stock mutual funds, a dummy for stock market participation, stock investment experience in years, as well as dummies for Census region and date of the survey. Table 2 uses only respondents in the four control groups, who have not received any information. Table A4 replicates the table using only respondents in the pure control group, who have not received any questions or information on past crashes.

¹⁵We cannot study the relationship between financial shocks and realized adjustments to risk-taking due to a potential reverse causality problem. Specifically, earlier or later realization of losses directly affects the capital losses households incurred during the pandemic.

likely to make adjustments to their portfolios, but there are no patterns according to education or political affiliation.

Overall, we find that investors were equally likely to reduce or increase their exposure to the stock market, although there is significant variation across groups. This is in line with Giglio et al. (2020b), who document that while respondents on average downward revised their short-run expectations about stock returns and GDP growth during the crash, they remained optimistic about the long-run outlook, and that disagreement across investors increased over the crash. In section 5 we explore our respondents' expectations about the future performance of the stock market and their role in driving plans about investments and other economic decisions.

Result 2. *About half of investors make active adjustments to their risky portfolio share during the COVID-19 crisis, with about equal fractions increasing or decreasing their risky share.*

4.2 Effects of shocks on plans about spending, debt and labor supply

How do US households' adjust their expectations about spending, debt, and labor market activity in response to wealth shocks during the pandemic, and how does the role of wealth shocks compare to that of income shocks?

4.2.1 Expected spending growth

The top row of Figure 3 plots households' expected nominal spending growth for the entire year 2020 compared to their spending in 2019 across groups. All groups on average report negative expected spending growth for 2020. Expected cuts to spending are most pronounced in the lowest net wealth quintile and in the middle of the income distribution, at about -7 percent. Individuals in age groups between 45 and 54, and between 55 and 64 report the strongest expected reduction in spending of about -10 and -8 percent, respectively. A large part of the average drop in spending is plausibly due to the shutdown of wide parts of society and the economy and the associated reduced consumption possibilities (Coibion et al., 2020b; Cox et al., 2020).

Table 2: Determinants of realized and planned adjustments to stock share

	Changed stock share	Increased stock share	Decreased stock share	Plan change stock share	Plan incr. stock share	Plan decr. stock share
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Net income (%)	-0.036 (0.052)	0.137** (0.054)	-0.173*** (0.056)	-0.062 (0.055)	0.100 (0.062)	-0.161*** (0.058)
Δ Retirement fin. wealth (%)				-0.090 (0.088)	-0.158* (0.084)	0.068 (0.073)
Δ Other fin. wealth (%)				-0.269*** (0.090)	-0.214** (0.088)	-0.055 (0.077)
Ln(Total fin. wealth)	2.836*** (0.873)	1.896** (0.810)	0.939 (0.824)	0.635 (0.879)	0.827 (0.861)	-0.191 (0.822)
Stock share in ret. wealth	-0.049 (0.034)	0.000 (0.029)	-0.050* (0.028)	-0.082** (0.034)	0.042 (0.031)	-0.124*** (0.025)
Stock share in ot. fin wealth	0.010 (0.033)	0.020 (0.028)	-0.011 (0.028)	0.026 (0.034)	0.042 (0.031)	-0.016 (0.026)
Share ret. in tot. fin. wealth	-0.151*** (0.047)	-0.103** (0.042)	-0.048 (0.044)	-0.173*** (0.047)	-0.177*** (0.045)	0.004 (0.039)
Any loss fin. crisis	7.627*** (2.538)	3.634 (2.493)	3.993* (2.343)	7.134*** (2.525)	-0.788 (2.552)	7.922*** (2.222)
Big loss fin. crisis	1.199 (3.152)	-7.420*** (2.681)	8.619*** (3.051)	-6.637** (3.011)	-7.127** (2.829)	0.490 (2.503)
Any loss dot-com	3.371 (2.572)	-1.338 (2.359)	4.709* (2.419)	3.412 (2.497)	2.768 (2.449)	0.644 (2.121)
Big loss dot-com	4.381 (4.443)	7.365* (4.023)	-2.985 (4.241)	8.189* (4.445)	7.820* (4.605)	0.369 (3.793)
Any loss Black Monday	8.085** (3.312)	0.340 (2.576)	7.745*** (2.996)	5.669* (3.112)	2.839 (2.752)	2.829 (2.349)
Big loss Black Monday	-2.557 (6.099)	2.844 (4.796)	-5.401 (5.411)	-0.854 (6.063)	-5.163 (5.254)	4.309 (4.515)
Male	11.591*** (2.350)	7.224*** (2.194)	4.367** (2.138)	11.618*** (2.418)	9.003*** (2.364)	2.615 (1.993)
At least bachelor	1.938 (10.936)	9.179 (10.930)	-7.240 (11.352)	-4.709 (10.091)	16.862 (10.723)	-21.572* (11.961)
Republican	-1.158 (2.184)	1.710 (2.089)	-2.868 (2.058)	3.080 (2.182)	3.661* (2.168)	-0.581 (1.923)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	.213	.148	.045	.27	.132	.119
Observations	2,148	2,148	2,148	1,999	1,999	1,999

Notes: This table shows OLS estimates of the determinants of realized and planned adjustments of the overall portfolio equity share. The outcomes are dummies indicating whether the respondent's household has made any change, has increased or has decreased the share of stocks and stock mutual funds in overall financial assets since the beginning of the stock market drop (columns 1-3) and dummies indicating plans to change, increase or decrease the equity share in overall financial assets in the weeks after the survey (columns 4-6), all coded as 0 or 100. All specifications are based on the four control arms, which have not received any information, using only respondents who report positive stockholdings as of January 2020. All specifications control for shocks to income, trimmed at the 2nd and 98th percentiles, dummies for having lost any wealth or substantial wealth during past stock market crashes, gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, log of total financial wealth, of real estate wealth, and of debt, share of total financial wealth in retirement accounts, borrowing constraints, stock market participation, stock shares in retirement and other accounts, investment experience, Census region, survey date, and the survey arm. Columns 4-6 additionally control for shocks to retirement and other financial wealth, trimmed at the 2nd and 98th percentile. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

Table 3: Effects of wealth and income shocks on expected behavior and plans

	Exp. spend. growth (%)	Exp. spend. growth (\$)	Exp. spend. growth (\$)	Incr. exp. debt	Incr. exp. desired hours	Incr. exp. retirement age
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Retirement fin. wealth (%)	0.054 (0.036)			-0.206*** (0.073)	-0.302*** (0.089)	-0.398*** (0.096)
Δ Other fin. wealth (%)	-0.044 (0.037)			-0.432*** (0.073)	-0.401*** (0.090)	-0.257*** (0.094)
Δ Net income (quarterly, %)	0.171*** (0.027)			-0.320*** (0.042)	-0.328*** (0.049)	-0.221*** (0.054)
Δ Retirement fin. wealth (\$)		0.024*** (0.009)	0.021** (0.009)			
Δ Other fin. wealth (\$)		-0.001 (0.012)	-0.004 (0.012)			
Δ Net income (quarterly, \$)		0.655*** (0.104)				
Δ Net income (annual, \$)			0.454*** (0.071)			
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	.041	.053	.084	.215	.146	.095
Observations	3,761	3,565	3,535	3,845	2,377	2,377

Notes: This table shows estimates of the association of shocks to the respondent's household financial wealth and net income with expected economic decisions. The outcomes are expected growth of yearly household spending from 2019 to 2020 in percent, trimmed at the 2nd and 98th percentiles (column 1); expected household spending growth in dollars, trimmed at the 2nd and 98th percentiles (columns 2-3); and dummies indicating whether the coronavirus crisis increases the respondent's expectations about outstanding household debt by the end of 2020 (column 4), expected desired working hours over the next years (column 5, only if in labor force) or expected retirement age (column 6, only if in labor force), all coded as 0 or 100. Dollar changes in columns 2 and 3 are constructed from survey questions for retirement and other financial wealth and for income (assuming that the respondent expected a quarter of her 2019 income in the first quarter of 2020), and for spending from the survey question on percent changes and estimates of the level of spending of different groups in 2019 from the CEX. Columns 1, 2, 4, 5 and 6 show simple OLS estimates. Column 3 shows 2SLS estimates, where the respondent's expected dollar change in household income from 2019 to 2020 is instrumented with the unexpected shock to household income over the first quarter of 2020. All specifications are based on the four control arms, which have not received any information. All specifications control for gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the equity share in total financial assets, investment experience, Census region, survey date, and the survey arm. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

What are the roles of financial wealth and income shocks in households' expectations about their spending? Table 3 column 1 regresses expected percent spending growth on percent shocks to retirement and non-retirement financial wealth and percent shocks to household income in the first quarter of 2020, as well as the baseline set of controls.¹⁶ While income shocks are strongly associated with an expected spending reduction, there is no significant relationship between wealth shocks and expected changes in spending. To facilitate the interpretation of magnitudes, columns 2 and 3 translate all variables into dollar changes. Column 2 uses the realized shock to household income during the first quarter, while column 3 uses the dollar shock to expected annual household income for 2020 using a 2SLS procedure in order to bring outcome and independent variable to the same scale.¹⁷ While we find a significant but small expected MPC of 2 cents for a one dollar shock to financial assets in retirement accounts, the MPC out of unexpected income shocks is much larger at 45 cents for each dollar shock to annual income. Shocks to financial wealth outside retirement accounts have no significant effect on expected changes in spending. Figure A10 uses binned scatter plots based on the specification in column 2 to illustrate the strong relationship between income shocks and expected spending growth and the small effect of wealth shocks.

In Figure A11 we examine heterogeneity in the effect of shocks to expected annual income on expected spending. Income shocks have the most pronounced effect for households with below median age, with below median incomes, or with no liquid assets, and for credit-constrained households, in line with the idea that these groups are more likely to exhibit hand-to-mouth behavior. Moreover, income shocks have particularly strong

¹⁶Table 3 uses only respondents in the four control groups, who have not received any information. Table A5 replicates the table using only respondents in the pure control group, who have not received any questions or information on past crashes.

¹⁷The dollar changes in financial wealth are calculated from survey questions on levels and percent changes. The quarterly income change in column 2 is calculated from the survey question on the unexpected percent shock to household income in the first quarter and total 2019 household income, assuming that the respondent had expected its household to earn a quarter of its total 2019 income in the first quarter of 2020. For the annual dollar shock to income in column 3 we use the first quarter dollar income shock to instrument the total expected dollar change in 2020 household income compared to 2019, which is calculated from survey questions on the expected percent change and the level in 2019. The expected dollar change in annual household spending is calculated from questions on expected percent change in spending from 2019 to 2020 and CEX estimates on the levels of annual spending of different groups to proxy spending in 2019.


effects on expected spending of individuals who expect their incomes never to recover, in line with predictions from standard life-cycle models on the differential effects of permanent and transitory income shocks. Figure A12 plots effects of shocks to retirement financial wealth or other financial wealth on expected spending across groups. We find small effects of shocks to retirement wealth among older or retired households, with a \$1 shock to retirement wealth being associated with a \$0.04 reduction in expected spending among retired individuals, likely as these households consume or plan to consume from retirement assets. Overall, the small size of the effects of wealth shocks is in line with previous literature documenting small MPCs out of changes in stock wealth (Bräuer et al., 2020; Di Maggio et al., 2019). These findings are consistent with the view that for many households retirement wealth is less liquid and therefore less likely to be used to finance spending. In addition, shocks to financial wealth are most pronounced among households with access to liquidity and credit, and may therefore lead to much smaller adjustments in spending than income shocks during the pandemic.

How do these patterns relate to other findings on the consumption responses to the COVID-19 pandemic? The average reduction in expected percent spending growth in our sample is lower than the spending cuts documented in other recent work (Andersen et al., 2020; Baker et al., 2020; Coibion et al., 2020b). Moreover, Cox et al. (2020) document that households' initial spending responses were uncorrelated with income shocks. While these studies examine immediate spending responses at the onset of the pandemic, we provide evidence on expected spending growth over the entire year 2020. One way to reconcile these findings is that while the pandemic first evolved as a supply-side shock, households expect it to unfold as a demand-side shock in the course of the year.

4.2.2 Expected household debt

Survey participants also report whether their expectations about outstanding household debt by the end of 2020 are altered by the current crisis. Overall, 36 percent of respondents report that they expect their household to have more debt outstanding by the end of the year as a result of the current crisis, while 11 percent say they expect lower debt. The bottom row of Figure 3 shows that the fractions of households expecting

Figure 3: Changes in expected spending and debt across groups



Notes: This figure displays the average percent change in expected total nominal household spending in 2020 compared to 2019 (top row) and the percent of respondents reporting that the current crisis increases their expected outstanding household debt by the end of the year 2020 (bottom row), by quintile of the pre-crisis net wealth distribution (left), by quintile of the pre-crisis net income distribution (middle) and by age group (right). The sample consists of respondents in the pure control group, who have not received any questions or information on past crashes before answering to the questions on spending and debt.

higher debt are substantial across groups, but more pronounced among those with lower net wealth, income, or those in younger age groups.

Table 3 column 4 shows that both wealth and income shocks are associated with a significantly higher tendency to report upward adjustments in expected outstanding household debt at the end of 2020. This suggests that households tend to smooth shocks to economic resources during the pandemic by taking out more debt or by postponing the repayment of debt.¹⁸ The top row of Figure A13 shows that income shocks are associated with upward adjustments in expected household debt particularly among those with lower incomes or higher age. We find no significant heterogeneity in the effect of wealth shocks

¹⁸Households may choose to take out a loan in order to make ends meet after being hit by income shocks. Wealth shocks could have a direct effect on household debt levels if households postpone the repayment of debt such as mortgages in response to wealth shocks, or if they expect to make smaller down payments on planned major purchases. Alternatively, wealth shocks could make people more pessimistic about the overall situation of their household going forward through experiential learning (Kuchler and Zafar (2019); Malmendier and Nagel (2011); see Table 4).


on expectations about debt.

4.2.3 Expected labor supply

In addition, 53 percent and 44 percent of respondents in the labor force as of January report that the current crisis increases their desired working hours in the next years or their expected retirement age, respectively. As shown in Figure 4, upward adjustments in expected labor supply are pronounced across groups. However, those in lower net wealth or income quintiles or in younger age groups are more likely to increase their expectations about desired working hours, while increases in expected retirement age are more frequent in the middle of the wealth and income distributions and among older respondents. Naturally, younger households have more time to make up for lost wealth and income and may therefore be less likely to adjust their retirement expectations. These findings point to an increase in labor supply in the US in the coming years.

Coibion et al. (2020b) document that many workers who lost their job in early 2020 dropped out of the labor force by retiring *early*, particularly older individuals. Consistent with their findings, some of our respondents report *downward* adjustments to their expected retirement age due to the crisis, and the propensity to do so is twice as high for people who report to have lost their jobs over the crisis (6 percent vs 3 percent), and highest for newly unemployed of age 55 and higher.

Figure 4: Changes in expected labor market activity across groups



Notes: This figure displays the percent of respondents who report that they have upward adjusted their expectations about desired working hours in the next years (top row) or their retirement age (bottom row) due to the current crisis (bottom row), by quintile of the pre-crisis net wealth distribution (left), by quintile of the pre-crisis net income distribution (middle) and by age group (right). The sample consists of respondents in the pure control group, who have not received any questions or information on past crashes before answering to the questions on expected labor market activity.

Covid Economics 23, 28 May 2020: 188-262

Moreover, Table 3 columns 5 and 6 show that both wealth and income shocks are associated with a significantly higher tendency to report upward revisions of expected desired working hours in the next years and expected retirement age. For instance, a one percentage point larger shock to retirement wealth is associated with a 0.43 percentage point higher likelihood of upward adjusting expected retirement age, while a one percentage point larger shock to net household income during the first quarter has an effect of 0.19 percentage points. Together, the average shocks in our sample to financial wealth inside and outside retirement accounts, and to net household income of -11, -8, and -5 percent predict a 15 percentage points higher probability of upward adjusting the expected retirement age. This implies that incurred wealth and income shocks can account for one third of the overall increase in expected retirement age due to the coronavirus crisis.¹⁹ Figure A15 displays these regressions in the form of binned scatter plots, highlighting that our findings are not driven by outliers. As illustrated in Figures A13 and Figure A14, we find no systematic heterogeneity in the effect of shocks on expected labor supply by economic resources and across demographic groups. Overall, these findings indicate that households plan to make up for wealth and income losses experienced during the crisis by working more in the coming years. Moreover, the pronounced effects of wealth shocks suggest that household exposure to the stock market can lead to swings in labor supply in response to stock market fluctuations, supporting a key mechanism in models of portfolio choice with human capital (Bodie et al., 1992; Boerma and Heathcote, 2019; Gollier, 2002).

Taken together, our third main result is the following:

Result 3. *Larger wealth and income shocks are associated with greater adjustments to planned economic activity. Income shocks strongly affect expected spending, while wealth shocks only have minor effects. However, households plan to make up for lost income and wealth by increasing their desired working hours over the coming years and by increasing their retirement age. Income and wealth shocks are also associated with upward adjustments in expected household debt.*

¹⁹Given measurement error in the shock variables and potential non-linearities, this can be interpreted as a lower bound.

5 Main results: Effects of expected stock market recovery duration on expectations and plans

In the previous section we have explored how capital losses incurred during the crash affect households' expectations about their medium- to long-term economic outcomes. In this section we study the formation of respondents' expectations about the further development of the stock market, and how those expectations causally shape their outlook for their own wealth and economic plans.

5.1 Descriptive evidence on beliefs about recovery

The survey asks respondents in which calendar year they expect the US stock market to have recovered to its pre-crisis level of January 2020. Respondents who report capital losses or income losses during the first quarter of 2020 also report the calendar year in which they expect their own wealth or income to have recovered, including options that they expect their wealth or income never to recover. The wording of the questions is agnostic about whether respondents expect further decreases in the stock market or their own wealth or whether they believe those outcomes to be on an increasing path at the time of the survey.

As shown in the left column of Figure A16, respondents who have made financial losses estimate that it will take 1.68 years for the stock market and 1.58 years for their own household wealth to recover to pre-crisis levels, and these patterns are fairly uniform across groups, aside from younger respondents expecting a longer stock market recovery duration.²⁰ Households who incurred income shocks expect their incomes to take 1.74 years to recover on average, with the lowest income, youngest, and oldest groups of respondents predicting a longer income recovery duration. The right column of Figure A16 documents that the fraction of respondents expecting their own financial wealth never to recover is highest among those with low net wealth or low net income, as well as among older respondents.

²⁰ Respondents who expect their wealth or income never to recover are excluded from the left column. The figure is based on respondents in the *PureControl* arm, who were not asked about nor received any information on the recovery duration in a previous crash.

5.2 Determinants of stock market and own wealth expectations

What is driving households' expectations about the stock market and the development of their own wealth after a crash? Table 4 explores the role of different factors previous literature has identified as playing a crucial role in expectation formation.²¹

We start with the role of personal experiences, which have been shown to be an important determinant of expectations about the stock market (Malmendier and Nagel, 2011), house prices (Kuchler and Zafar, 2019) or inflation (D'Acunto et al., 2019b; Goldfayn-Frank and Wohlfart, 2019; Malmendier and Nagel, 2016). Individuals who have experienced more negative income shocks expect the stock market to take more time to recover (column 1), expect lower stock returns (column 2), and perceive a higher probability of extreme negative stock market returns of below -30 percent (column 4) and a lower probability of very high stock returns (column 5). These patterns are somewhat weaker for financial wealth losses in the recent crash. Naturally, individuals who were hit harder expect a longer recovery duration for their own financial situation (columns 6-7). However, they are also more likely to expect a further worsening of their household's financial situation over the next year (column 8).

We also study the role of personal experiences made in historic crashes. Having experienced losses in the stock market during the Financial Crisis 2007-9 is associated with more pessimistic expectations about the stock market and regarding the development of own wealth, while the patterns are less consistent for losses incurred in earlier crashes such as the burst of the Dot-com bubble or the Black Monday. This is in line with the previously documented recency bias in the role of personal experiences in macroeconomic expectation formation (Kuchler and Zafar, 2019; Malmendier and Nagel, 2011). These findings highlight that personal experience seems to be an important driver of individuals' expectations in the time following a market crash. Moreover, this evidence offers an explanation for the more pronounced tendency to reduce stock investments among those

²¹Table 4 reports multivariate regressions of these expectations on a set of covariates. It uses only respondents in the four control groups, who have not received information. Table A6 replicates the table using only respondents in the pure control group, who have not received any questions or information on past crashes.

Table 4: Determinants of expectations about the stock market and own wealth

	Stock recovery duration	Stock return: Mean	Stock return: SD	Stock return <-30%	Stock return >30%	Wealth recovery duration	Exp. wealth never to recover	Household financial prospects
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Fin. wealth (%)	-0.005 (0.004)	-0.012 (0.026)	-0.045*** (0.013)	-0.014 (0.030)	-0.077* (0.039)	-0.034*** (0.003)	-0.159*** (0.054)	0.005*** (0.002)
Δ Net income (%)	-0.010*** (0.003)	0.085*** (0.017)	-0.001 (0.009)	-0.100*** (0.023)	0.067*** (0.022)	-0.005*** (0.002)	-0.088** (0.038)	0.007*** (0.001)
Any loss fin. crisis	-0.007 (0.120)	-1.679** (0.804)	0.602 (0.414)	1.289 (1.006)	-1.504 (1.162)	0.237*** (0.084)	4.460*** (1.691)	-0.077* (0.046)
Big loss fin. crisis	0.302** (0.150)	-1.965* (1.128)	-0.542 (0.515)	2.302 (1.451)	-0.404 (1.593)	0.156 (0.128)	-0.540 (2.166)	-0.260*** (0.062)
Any loss dot-com	-0.069 (0.125)	-0.362 (0.861)	1.088** (0.428)	0.261 (1.117)	-1.162 (1.206)	0.012 (0.106)	0.223 (1.745)	-0.009 (0.049)
Big loss dot-com	-0.215 (0.219)	-1.588 (1.662)	-0.860 (0.705)	4.313* (2.359)	0.608 (2.182)	-0.258* (0.155)	3.862 (3.267)	-0.040 (0.098)
Any loss Black Monday	0.282* (0.154)	-1.478 (1.046)	-0.700 (0.520)	-0.900 (1.354)	-3.558** (1.530)	0.239* (0.122)	-0.998 (2.049)	-0.057 (0.054)
Big loss Black Monday	0.206 (0.289)	0.649 (1.976)	-0.596 (0.813)	0.403 (2.798)	1.511 (2.551)	0.280 (0.248)	5.851 (4.212)	-0.042 (0.104)
Male	-0.640*** (0.103)	0.218 (0.709)	0.693** (0.345)	-1.167 (0.913)	0.082 (1.047)	-0.184*** (0.059)	-1.000 (1.404)	0.186*** (0.037)
At least bachelor	-0.069 (0.317)	0.279 (2.083)	-1.090 (1.003)	-5.728* (3.022)	-1.814 (2.820)	-0.303** (0.145)	-10.735** (4.637)	0.093 (0.104)
Republican	-0.657*** (0.095)	5.844*** (0.680)	-0.808** (0.337)	-3.577*** (0.862)	7.224*** (1.010)	-0.222*** (0.064)	-4.639*** (1.361)	0.240*** (0.036)
Stock investor	-0.445** (0.194)	0.088 (1.309)	1.314* (0.682)	-0.400 (1.676)	-0.482 (1.865)	0.469*** (0.125)	2.493 (2.754)	0.137* (0.072)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	.092	.048	.054	.024	.037	.304	.081	.088
Observations	3,918	3,918	3,918	3,918	3,918	3,614	3,918	3,918

Notes: This table shows OLS estimates of the determinants of respondents' expectations about the stock market and their own wealth. The outcomes are the expected duration of the recovery of the US stock market to its pre-crisis level of January 2020 in years (column 1); mean and standard deviation as well as probabilities assigned to extreme return realizations based on the respondent's reported probability distribution over the one year-ahead stock market return (columns 2-5); the expected recovery duration of the respondent's household net wealth (column 6); a dummy indicating whether the respondent thinks her household net wealth will never recover, coded as 0 or 100 (column 7); and a categorical measure of the respondent's subjective household financial prospects, z-scored using the mean and standard deviation in the sample (column 8). All specifications are based on the four control arms, which have not received any information. All specifications control for shocks to income and financial wealth, trimmed at the 2nd and 98th percentiles, dummies for having lost any wealth or substantial wealth during past stock market crashes, gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the equity share in total financial assets, investment experience, Census region, survey date, and the survey arm. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

who have lost wealth during past crashes (see section 4.1).

Men predict shorter recovery durations of the market and of their own wealth and are significantly more optimistic about their household's financial prospects, in line with previously documented gender gaps in macroeconomic expectations (D'Acunto, 2020; D'Acunto et al., 2020). Finally, Republicans expect the recovery to be 0.7 years shorter compared to Democrats, they predict a six percentage point higher stock return, and are more optimistic about their own household's financial situation. These patterns are in line with earlier findings documenting strong partisan bias in reported survey expectations (Mian et al., 2018) and a partisan gap in stock investment following the presidential election of Donald Trump (Meeuwis et al., 2019).²² Given the pronounced heterogeneity in expectations according to political affiliation it seems surprising that we did not detect significant differences between Republicans and Democrats in active adjustments to risk-taking in section 4.1.

Result 4. *Experienced losses in past crashes, gender as well as political affiliation are important determinants of beliefs about the recovery from the February/March 2020 stock market crash.*

5.3 Learning from information about past crashes

Our survey includes a short experimental section in which respondents report their prior beliefs about the duration of the recovery in the case of a historical stock market crash, and random subsets of respondents receive information on the actual recovery duration. We use this experimental setup i) to shed light on the role of beliefs about past crashes in shaping respondents' expectations in the current situation; and ii) to provide causal evidence on the role of stock market expectations in shaping respondents' outlook for their own situation and their planned economic behavior.

Stock market crashes are rare events, and can have a variety of different origins, ranging from corrections to the value of firms or industries to problems in the housing

²²Another driver of this result may be geographic and urban heterogeneity in social distancing and exposure to the coronavirus pandemic. Related research suggest that differences in exposure affect expectations and outcomes, and exposure and combative measures vary significantly geographically (Baker et al., 2020; Bu et al., 2020; Coibion et al., 2020b; Kuchler et al., 2020).

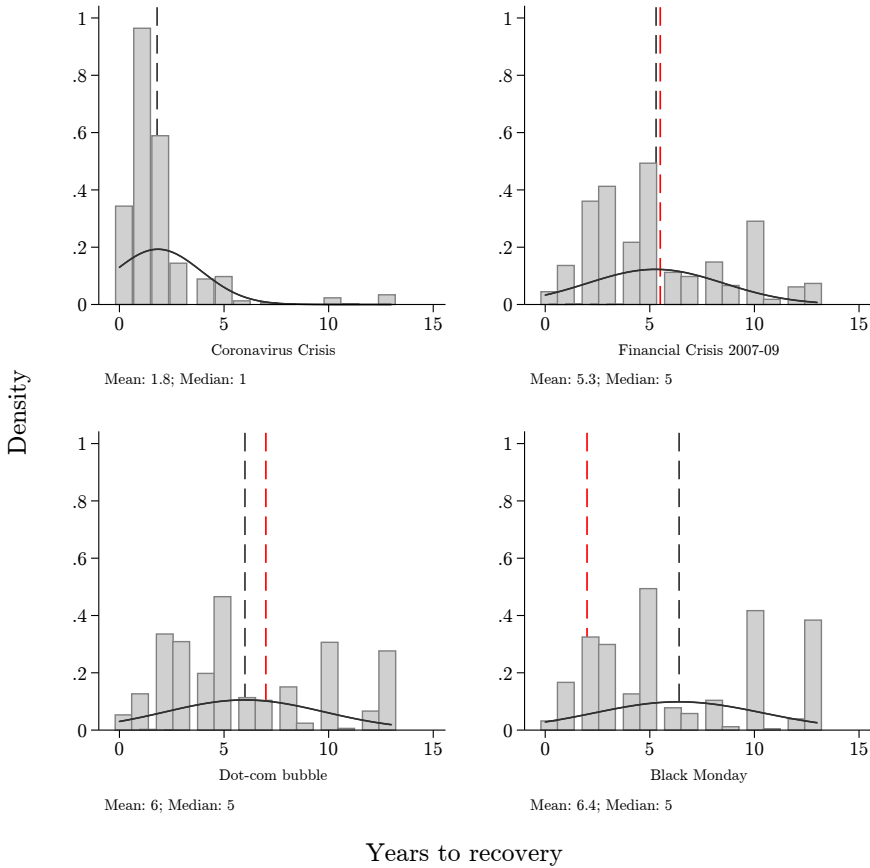
market or shocks to the real economy. Given the unprecedented speed and strength of the current crash, and given its origin in the first world-wide pandemic for more than 100 years, the historical database for predicting the further development of the stock market is arguably limited. Do our respondents believe that the current crash is “unique” in the sense that it is not comparable to previous crashes, or do they consider facts about historical crashes to be relevant for the current situation?

Figure 5 displays beliefs about the recovery duration from the current crash and from past crashes using respondents in the relevant survey arms. 65.8 percent underestimate the duration of recovery in the case of the Financial Crisis 2007-9 (5 1/2 years, top-right) and 63.2 percent do so for the Dot-com bubble in 2000 (7 years, bottom-left), but a majority of 79.4 percent overestimate the duration of recovery from the Black Monday crash 1987 (2 years, bottom-right). Respondents in the pure control group, who have not received any questions or information on past crises, predict a recovery duration of 1.9 years for the current crash. However, given differences in the scales on which these beliefs are elicited, one should interpret these differences with caution.²³

Given these patterns in prior beliefs, the information that random subsets of our respondents receive can be seen as pessimistic (in the cases of the longer recovery durations of the Financial Crisis 2007-9 or the Dot-com bubble 2000) or as optimistic (in case of the shorter recovery duration following the Black Monday crash). How do respondents change their beliefs about the current situation when provided with information on the length of recovery from past crashes? In Table 5 we regress respondents’ post-treatment expectations about the stock market on dummy variables indicating whether they have received information. Panels A, D and G use all respondents in the relevant arms. Panels B, E and H restrict the analysis to respondents who underestimate actual historic recovery durations in the case of the “pessimistic” Financial Crisis and Dot-com bubble treatments, or to respondents who overestimate the time until recovery in the “optimistic” Black Monday treatment. Panels C, F and I use only over- or underestimators who re-

²³Specifically, beliefs about historical crashes are elicited asking for number of years, while beliefs about the current situation are elicited asking for calendar year. The different elicitation scales have important advantages for our experimental analysis, as they mitigate concerns related to numerical anchoring.

Figure 5: Beliefs about durations of current and historical stock market recoveries



Notes: This figure displays respondents' subjective beliefs about the duration of the recovery of the US stock market in years for the Coronavirus crisis (top left), the Financial Crisis of 2007-2009 (top right), the Dot-com bubble (bottom left) and the crisis following Black Monday on October 19, 1987 (bottom right). The sample for the Coronavirus crisis consists of the pure control sample, where respondents did not receive any questions or information on past crashes before answering the question on expected recovery duration from the current crash. For the Financial Crisis, the Dot-com bubble and Black Monday it consists of the control and treatment samples in the relevant arms that answered questions on the corresponding crash. The expected duration is elicited prior to the respondent receiving information about the true duration (red dashed line). The mean estimate of the recovery duration is displayed as the black dashed line. Recovery duration is winsorized at 13 years in each subfigure.

port positive stockholdings as of January 2020. The “pessimistic” treatments providing information on the Financial Crisis or the Dot-com bubble increase respondents’ expected recovery duration by between 1.3 and 2.3 years, while the “optimistic” treatment providing information on the Black Monday crash reduces the expected recovery duration by about one year (column 1). Given a standard deviation of expected recovery beliefs of 3.2 years, the economic magnitude of these effects is substantial.

The treatments also lead to shifts between 0.1 and 0.3 standard deviations in respondents’ extent of agreement to verbal statements describing the severity of the current crash (columns 2-4).²⁴ The Financial Crisis and the Black Monday treatments move respondents’ expected stock returns by up to -3 and up to 2 percentage points, respectively (column 5), and change the subjective probabilities assigned to extreme return realizations accordingly (columns 7-8). The size of these effects amounts to about half of the strong partisan gap in expectations documented in Table 4. Most of the coefficient estimates increase in absolute size when restricting the sample to over- or under-estimators, although we lack the power to meaningfully explore differences in effect sizes across groups.²⁵

Taken together, the strong effects of information on respondents’ expectations about the stock market highlight that households continue to form expectations based on their beliefs about stock market developments in the past, even in very unique and unprecedented situations. Moreover, these findings imply that information about historical stock market developments had not been fully incorporated into respondents’ prior expectations, pointing to an important role of information frictions in the formation of households’ stock market expectations. This is consistent with models in which information is costly to acquire or to process (Abel et al., 2007; Alvarez et al., 2012), which may result in a lack of preparation particularly for rare events (Maćkowiak and Wiederholt, 2018).

²⁴Specifically, the treatments change respondents’ agreement on 7-point categorical scales (which we z-score using the mean and standard deviation in our sample) with the following statements: “The outbreak of the coronavirus will keep US stock prices below their January 2020 levels for many years.” (column 2); “The outbreak of the coronavirus has set the level of the stock market back by many years.” (column 3); “The US stock market will have recovered by the end of the year 2020.” (column 4).

²⁵The experimental findings on the effect of shocks to beliefs about past crashes are mirrored in correlations between priors about historical recovery durations and current stock market expectations using respondents in the control groups, which are reported in Table A7.

Table 5: Effects of information on stock market expectations: Experimental first stage

	Stock recovery duration	Agree: Recovery many yrs.	Agree: Set back many yrs.	Agree: Recover 2020	Stock return: Mean	Stock return: SD	Stock return <-30%	Stock return >30%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: All								
Info Fin. Crisis 2007	1.304*** (0.108)	-0.207*** (0.041)	-0.179*** (0.042)	0.237*** (0.042)	-2.561*** (0.795)	0.014 (0.394)	2.040* (1.055)	-2.239** (1.134)
Panel B: Underestimators								
Info Fin. Crisis 2007	1.645*** (0.115)	-0.317*** (0.051)	-0.311*** (0.051)	0.287*** (0.050)	-3.026*** (0.947)	0.086 (0.482)	1.271 (1.266)	-3.724*** (1.322)
Panel C: Under. & Stocks > 0								
Info Fin. Crisis 2007	1.503*** (0.134)	-0.332*** (0.064)	-0.308*** (0.066)	0.293*** (0.063)	-2.226* (1.145)	-0.216 (0.567)	1.308 (1.517)	-2.520 (1.541)
Panel D: All								
Info Dot-com 2000	1.758*** (0.123)	-0.101** (0.041)	-0.145*** (0.042)	0.199*** (0.043)	-1.615** (0.795)	0.220 (0.400)	1.238 (1.063)	-1.238 (1.107)
Panel E: Underestimators								
Info Dot-com 2000	2.261*** (0.134)	-0.207*** (0.053)	-0.231*** (0.053)	0.211*** (0.054)	-1.545 (0.988)	0.216 (0.497)	1.123 (1.363)	-1.509 (1.342)
Panel F: Under. & Stocks > 0								
Info Dot-com 2000	2.071*** (0.170)	-0.221*** (0.068)	-0.260*** (0.069)	0.164** (0.068)	-0.977 (1.158)	-0.086 (0.610)	0.742 (1.550)	-0.987 (1.526)
Panel G: All								
Info Black Monday 1987	-0.928*** (0.104)	0.107*** (0.041)	0.107*** (0.041)	0.062 (0.042)	1.486* (0.766)	0.247 (0.382)	-2.733*** (0.941)	-0.575 (1.132)
Panel H: Overestimators								
Info Black Monday 1987	-1.276*** (0.125)	0.169*** (0.045)	0.192*** (0.046)	-0.047 (0.046)	2.016** (0.856)	0.296 (0.434)	-3.179*** (1.075)	0.037 (1.226)
Panel I: Over. & Stocks > 0								
Info Black Monday 1987	-1.108*** (0.157)	0.140** (0.060)	0.162*** (0.061)	-0.079 (0.062)	1.939* (1.029)	0.184 (0.536)	-3.016** (1.210)	0.234 (1.449)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS estimates of the effect of being shown information on the duration of a historical stock market crash on respondent's expectations about the stock market. The outcomes are the expected duration of the recovery of the US stock market to its pre-crisis level of January 2020 in years (column 1); agreement on 7-point scales to statements describing the severity of the current stock market crash, z-scored using the mean and the standard deviation in the sample (columns 2-4); mean and standard deviation as well as probabilities assigned to extreme return realizations based on the respondent's reported probability distribution over the one year-ahead stock market return (columns 5-8). Panels A-C are based on the treatment and control arms including information or questions on the Financial Crisis 2007. Panels D-F are based on the treatment and control arms including information or questions on the burst of the Dot-com bubble 2000. Panels G-I are based on the treatment and control arms including information or questions on the Black Monday 1987. Panels A, D and G are based on the full sample in the corresponding arms. Panels B, E and H are based only on under-estimators (for Financial Crisis and Dot-com bubble) or over-estimators (for Black Monday) of the length of the recovery from the crash. Panels C, F and I are based only on under-estimators or over-estimators who participated in the stock market before the current crisis. All specifications control for the respondent's prior belief about the recovery duration following the corresponding crash as well as for gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the equity share in total financial assets, investment experience, Census region and survey date. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

5.4 Expected stock market recovery and own outlook and plans

Do expectations about the further development of the stock market have similar effects on households' economic plans as capital losses already incurred during the crash (see section 4.1)? Our randomized provision of information about past crashes generates exogenous variation in our respondents' recovery expectations. We exploit this setting to shed light on the causal effects of households' stock market expectations on their outlook regarding their own wealth and their plans about investment, spending, debt and labor supply.

In Table 6 we regress different outcomes on respondents' expected recovery duration of the stock market and our baseline set of control variables. First, the table shows OLS estimations using respondents in all control groups, who have not received any information (Panels A and B). Second, the table shows 2SLS estimations, where the respondents' expected recovery duration is instrumented with the dummy for the relevant information treatment assignment, as well as the corresponding OLS estimates in the relevant subsamples (Panels C-H). Panels B-H restrict the sample to stockholders as of January 2020, and Panels C-H are restricted to the majorities of respondents who overestimated (Panels C-F) or who underestimated (Panels G-H) the duration of recovery from the corresponding historical crash.²⁶ In addition, Table A9 displays OLS estimates for different subgroups using all control groups.²⁷

Expectations about own wealth Respondents' beliefs about the recovery duration of the stock market are strongly correlated with their expectations about their own wealth (Table 6 column 1). Among stockholders, a one year longer expected stock market recovery translates into a 0.45 years longer expected recovery of respondents' own wealth and a 0.09 standard deviations reduction in people's financial prospects for their household (column 2). For comparison, already incurred financial wealth losses of 11 percent (the

²⁶This increases the strength of our first stage estimates and ensures that the monotonicity assumption (that the first stage shifts all respondents' beliefs in the same direction) holds. Table A8 presents reduced form estimates of the effects of the information treatments on wealth expectations and plans.

²⁷We are not powered to conduct IV estimations on subsamples due to the smaller sample available for each instrument.

Table 6: Effects of expected stock market recovery on own outlook and plans: OLS and 2SLS estimates

	Wealth recovery duration	Household financial prospects	Plan incr. stock share	Plan decr. stock share	Exp. spend. growth	Incr. exp. debt	Incr. exp. desired hours	Incr. exp. retirement age
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: All control groups (OLS)								
Expected stock recovery duration (years)	0.223*** (0.017)	-0.070*** (0.006)	-0.584** (0.282)	0.530* (0.280)	-0.250* (0.138)	1.063*** (0.258)	1.303*** (0.356)	1.727*** (0.362)
Observations	3,903	4,260	2,599	2,599	4,097	4,260	2,682	2,682
Panel B: All control groups & Stocks (OLS)								
Expected stock recovery duration (years)	0.455*** (0.029)	-0.088*** (0.009)	-0.584** (0.282)	0.530* (0.280)	-0.546*** (0.178)	0.517 (0.342)	2.050*** (0.453)	2.065*** (0.464)
Observations	2,297	2,599	2,599	2,599	2,515	2,599	1,915	1,915
Panel C: Fin. Cris. Und.&Stocks (OLS)								
Expected stock recovery duration (years)	0.407*** (0.043)	-0.113*** (0.017)	-0.300 (0.556)	0.517 (0.625)	-0.693* (0.372)	2.102*** (0.714)	3.079*** (0.894)	4.184*** (0.789)
Panel D: Fin. Cris. Und.&Stocks (2SLS)								
Expected stock recovery duration (years)	0.487*** (0.072)	-0.091** (0.042)	1.844 (1.491)	2.096 (1.413)	-0.844 (0.852)	2.063 (1.936)	5.770** (2.321)	5.704** (2.377)
First-stage F-stat	130.984	125.897	131.394	131.394	129.047	125.897	108.424	108.424
Observations	781	879	879	879	854	879	675	675
Panel E: Dot-com Und.&Stocks (OLS)								
Expected stock recovery duration (years)	0.429*** (0.046)	-0.073*** (0.014)	-0.190 (0.537)	0.337 (0.477)	-0.402 (0.309)	0.426 (0.552)	0.870 (0.707)	0.958 (0.715)
Panel F: Dot-com Und.&Stocks (2SLS)								
Expected stock recovery duration (years)	0.519*** (0.073)	0.025 (0.033)	0.553 (1.215)	-1.428 (1.075)	0.371 (0.639)	0.081 (1.362)	-0.639 (1.779)	-1.924 (1.867)
First-stage F-stat	119.547	148.127	150.857	150.857	150.266	148.127	112.237	112.237
Observations	740	839	839	839	823	839	625	625
Panel G: Black Mon. Over.&Stocks (OLS)								
Expected stock recovery duration (years)	0.531*** (0.040)	-0.081*** (0.012)	-1.680*** (0.415)	1.192*** (0.409)	-1.196*** (0.266)	0.635 (0.534)	1.124 (0.727)	0.837 (0.718)
Panel H: Black Mon. Over.&Stocks (2SLS)								
Expected stock recovery duration (years)	0.341*** (0.102)	-0.087 (0.054)	-4.298** (2.058)	3.773** (1.820)	-1.118 (1.136)	4.532* (2.453)	1.220 (3.600)	3.828 (3.716)
First-stage F-stat	48.703	50.007	48.322	48.322	54.393	50.007	27.102	27.102
Observations	903	1,033	1,033	1,033	993	1,033	756	756
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS and 2SLS estimates of the effect of respondents' expected stock market recovery duration on their expectations about their own financial situation and behavior. The outcomes are the expected recovery duration of the respondent's household net wealth (column 1); a categorical measure of the respondent's subjective household financial prospects, z-scored using the mean and standard deviation in the sample (column 2); dummies indicating plans to increase or decrease the equity share in overall financial assets in the weeks after the survey (columns 3-4, only for stockholders); expected growth of yearly household spending from 2019 to 2020 in percent, trimmed at the 2nd and 98th percentiles (column 5); dummies indicating whether the coronavirus crisis increases the respondent's expectations about outstanding household debt by the end of 2020 (column 6), expected desired working hours over the next years (column 7, only if in labor force) or expected retirement age (column 8, only if in labor force). All dummy outcomes are coded as 0 or 100. Panels A and B are based on the four control arms, which have not received any information. In Panels D, F and H, we use the relevant information treatment dummy as instrument for expected recovery duration. Panels C-D are based on the treatment and control arms including information or questions on the Financial Crisis 2007, Panels E-F are based on the Dot-com arms, and Panels G-H are based on the Black Monday arms. Panels C-F are restricted to under-estimators and Panels G-H to over-estimators. Panels B-H are restricted to those who participated in the stock market before the current crisis. All specifications control for gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the equity share in total financial assets, investment experience, Census region and survey date. The specifications on planned stock trading in columns 3 and 4 also control for realized trading since the onset of the crisis. Panels C-H also control for the respondent's prior belief about the recovery duration following the corresponding crash. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

Covid Economics 23, 28 May 2020: 188-262

mean losses in our sample), are associated with a longer expected recovery of own wealth by 0.39 years and a reduction in household's financial prospects by 0.06 of a standard deviation (see Table 4). The 2SLS estimates exploiting the experimental variation are mostly highly significant and of similar size as the OLS estimates. Table A9 columns 1-2 show that stock market recovery expectations play a significantly larger role for the wealth expectations of older stockholders, those with lower net wealth and for men. These findings indicate that beliefs about the further development of the stock market play a substantial causal role in shaping households' wealth expectations, particularly among those who have less time during their working life to make up for these losses.

Investment plans Do people's expectations about the further development of the stock market also affect their planned investment behavior? A substantial literature has studied correlations between subjective expectations and stock market participation or the equity portfolio share (Ameriks et al., 2019; Amromin and Sharpe, 2014; Dominitz and Manski, 2007; Giglio et al., 2020a; Hudomiet et al., 2011; Vissing-Jørgensen, 2002). While our survey only contains self-reported investment plans, the randomized information provision allows us to provide, to the best of our knowledge, the first causal evidence on the role of subjective return expectations in (planned) investment decisions. Respondents who expect a longer recovery are 0.58 percentage points more likely to plan to increase the share of their portfolio invested in equities and 0.53 percentage points less likely to plan a reduction (Table 6 columns 3 and 4). The results of the IV estimations exploiting the "pessimistic" Financial Crisis and Dot-com instruments are insignificant. However, when we use the "optimistic" Black Monday instrument we estimate significant causal effects of expected recovery duration on plans to increase and to decrease the share invested in stocks by -4.3 and by 3.7 percentages points, respectively. Average tendencies to plan increases or decreases are 28 and 23 percent in our sample, highlighting that expectations seem to play an important role in shaping investment plans following a crash. Table A9 columns 3-4 show that the association of expectations and investment plans seems to be fairly uniform across groups. Future research could link survey and administrative data to examine whether investment decisions are more elastic to beliefs during times

of market turmoil than during more tranquil times, when the role of beliefs seems to be moderate (Giglio et al., 2020a).

Expected spending and debt How do people's expectations about the stock market recovery affect their plans in other domains? Expectations about the duration of the stock market recovery are negatively correlated with respondents' expectations about their spending growth (Table 6 column 5). However, none of the causal estimates from the 2SLS regressions are significant. This is in line with our earlier finding that financial wealth shocks incurred during the February/March 2020 stock market crash are not reflected in significant changes in expected spending growth (see section 4.2 and Figure A10). More pessimistic expectations about the stock market recovery are associated with a significantly higher tendency to report upward adjustments in expected household debt for the end of 2020 (column 6). However, this correlation turns insignificant once we restrict the sample to stockholders. Among the causal estimates only the Black Monday treatment gives a marginally significant estimate. We interpret this as mixed evidence for an effect of stock market expectations on expectations about household debt

Expected labor market activity Portfolio choice models including human capital predict that households should adjust their labor supply in response to wealth fluctuations (Bodie et al., 1992; Boerma and Heathcote, 2019; Gollier, 2002). Among stockholders, a one year increase in expected stock market recovery duration is associated with 2 percentage point increases in the tendencies to upward adjust expectations about desired working hours over the next years (Table 6 column 7) and about retirement age (column 8). Moreover, using the "pessimistic" Financial Crisis instrument we find significant causal effects of about 5 percentage points on the tendency to upward adjust these expectations.²⁸ The economic magnitudes of these effects is substantial, corresponding to the effects of having experienced a shock to retirement financial wealth of 21 or 13 percentage points, respectively. Table A9 columns 7-8 highlight that the effects of stock market expectations on planned labor market activity are particularly pronounced among those with lower net

²⁸While we find no significant effects using the other instruments, we note that also the OLS estimates in the relevant arms are insignificant.

wealth, suggesting that these households' long-term plans did not include a buffer for the case of large wealth losses during a crash. To the best of our knowledge our experimental findings are the first direct causal evidence on the role of expectations about financial markets in shaping people's long-term plans regarding work. This highlights that rising exposure to the stock market among households can have important consequences for households' long-term planning and for labor markets.

Taken together, these results suggest that not only wealth shocks incurred during the crash, but also beliefs about the performance of the stock market in the next years play an important role in shaping US households' expectations about their own economic situation and plans. Our fifth main result is the following:

Result 5. *Households' beliefs about the duration of recovery from the stock market crash are strongly correlated with expectations about their own wealth, their planned investment behavior and their long-term expectations about labor market activity. Results from IV estimations exploiting randomized information provision suggest that part of these effects are causal.*

5.5 Robustness

Cross-learning Respondents update their expectations about their own wealth and their economic plans in response to the provided information, plausibly through direct effects working through their stock market expectations. Alternatively, there could be cross-learning in the sense that respondents may update their beliefs about overall GDP growth and labor markets in response to the information. We view such cross-learning as a natural by-product of changes in expectations induced by random information provision. For instance, changes in stock return expectations in panel data from existing surveys tend to be associated with changes in GDP growth expectations (Amromin and Sharpe, 2014; Giglio et al., 2020a), raising the question whether it would be a meaningful exercise to change people's expectations about stock returns, while holding fixed their expectations about growth. However, we do not believe that cross-learning about GDP growth or labor markets is the main driver behind our findings. First, we find mostly insignificant

treatment effects when we restrict our sample to non-stockholders. Second, we find only minor effects on respondents' income expectations due to our information treatments. These results are unreported for brevity but available upon request.

Numerical anchoring One concern about our experimental findings could be unconscious numerical anchoring on the provided information (Cavallo et al., 2017; Coibion et al., 2019b). We believe that in our setting this concern is likely much less severe than in other settings because the response scales of the post-treatment questions are different to the scale of the provided information (calendar year, 7-point agreement scale, or density distribution instead of number of years). Moreover, previous studies have documented only small changes of reported survey expectations in response to the provision of irrelevant numerical anchors (Coibion et al., 2019b; Roth and Wohlfart, 2019).

Experimenter demand effects Relatedly, our experimental findings could be driven by experimenter demand effects, i.e. by subjects guessing the experimental hypothesis and reporting posterior beliefs such as to confirm with the hypothesis. We think that our experimental findings are unlikely driven by demand effects for three reasons: i) our study is fully based on a between-subject design, where no question is asked twice in the survey (i.e. both before and after the treatment, as in within-subject designs), arguably mitigating demand effects; ii) at the end of our survey we explicitly asked our respondents to report their beliefs about the purpose of the study, and less than 10 respondents suspected the survey to contain some form of experimental treatment (Table A10); iii) experimenter demand effects have been shown to be of limited importance in comparable online surveys (de Quidt et al., 2018).

6 Implications and conclusion

With increasing stock market participation, households around the world have become more exposed to stock market downturns. We have conducted a survey on a representative sample of more than 8,000 US households, which offers a comprehensive real-time snapshot of US households' finances and expectations about the future in the time following one such crash. We document that shocks to households' financial wealth due to

the COVID-19 stock market decline tend to be negatively correlated with income shocks experienced during the early stages of the pandemic. While about half of all stock owners made adjustments to their investments in the course of the crash, there was no systematic tendency to rebalance portfolios in response to the passive reduction in equity portfolio shares. Financial wealth shocks are associated with adjustments in expectations about household debt, retirement age and desired working hours, but have no substantial effect on expected spending. Finally, beliefs about the recovery of the stock market causally shape individuals' expectations about their own wealth and their plans regarding investment, debt and labor market activity in the future.

Our findings highlight that exposure to stock market downturns is concentrated among groups who tend to be less exposed to income shocks and job losses during recessions. Moreover, when households experience shocks to their retirement wealth during a stock market crash, they plan to make up for it by increasing their labor supply in the following years and by postponing their retirement age. Similarly, their expectations about the stock market recovery directly shape their expectations about own labor market activity. This implies that households who invest their retirement wealth in stocks accept fluctuations in their long-term expectations about retirement age and working life, in line with a key mechanism in portfolio choice models including human capital (Bodie et al., 1992; Gollier, 2002). Households who are unwilling to accept such fluctuations may be reluctant to invest in stocks, contributing to the widely-documented non-participation in the stock market across groups (Guiso and Sodini, 2013; Haliassos and Bertaut, 1995). Moreover, since for older households it is more difficult to make up for wealth losses by extending labor supply, this mechanism can explain reductions in the equity share as people age. At a macro level, our results suggest that increasing household exposure to the stock market may generate a link from financial market developments to medium-term swings in labor supply.

Our results have several more specific implications for the economic and financial consequences of the COVID-19 pandemic. First, in order to adequately gauge the short-run impact of the current crisis on inequality in overall economic resources one should

consider both income shocks and wealth shocks. Second, our findings on substantial increases in expected retirement age or desired working hours suggest that there will be an increase of labor supply in the US after the lockdowns are lifted, as households are trying to make up for the lost wealth and income. This could put downward pressure on wages and further aggravate economic hardships for those in the bottom of the distribution. Third, beliefs about the recovery of the stock market seem to be central to individuals' subjective economic prospects and expected decisions, indicating that policymakers may stimulate the economic recovery after the lockdown by managing these expectations.

References

- Abel, Andrew B, Janice C Eberly, and Stavros Panageas**, "Optimal Inattention to the Stock Market," *American Economic Review*, 2007, 97 (2), 244–249.
- Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh**, "Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys," *Working Paper*, 2020.
- Agnew, Julie, Pierluigi Balduzzi, and Annika Sundén**, "Portfolio Choice and Trading in a Large 401(k) Plan," *American Economic Review*, March 2003, 93 (1), 193–215.
- Alvarez, Fernando, Luigi Guiso, and Francesco Lippi**, "Durable Consumption and Asset Management with Transaction and Observation Costs," *American Economic Review*, 2012, 102 (5), 2272–2300.
- Ameriks, John and Stephen P Zeldes**, "How do Household Portfolio Shares Vary with Age?," *Working Paper*, 2004.
- , **Gábor Kézdi, Minjoon Lee, and Matthew D Shapiro**, "Heterogeneity in Expectations, Risk Tolerance, and Household Stock Shares: The Attenuation Puzzle," *Journal of Business & Economic Statistics*, 2019, pp. 1–27.
- Amromin, Gene and Steven A Sharpe**, "From the Horse's Mouth: Economic Conditions and Investor Expectations of Risk and Return," *Management Science*, 2014, 60 (4), 845–866.
- Andersen, Asger Lau, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan**, "Consumer Responses to the COVID-19 Crisis: Evidence from Bank Account Transaction Data," *Working Paper. CEPR Covid Economics 7*, 2020, pp. 88–114.
- Andersen, Steffen, Tobin Hanspal, and Kasper Meisner Nielsen**, "Once Bitten, Twice Shy: The Power of Personal Experiences in Risk Taking," *Journal of Financial Economics*, 2019, 132 (3), 97–117.

- Armantier, Olivier, Scott Nelson, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar**, “The Price Is Right: Updating Inflation Expectations in a Randomized Price Information Experiment,” *Review of Economics and Statistics*, 2016, 98 (3), 503–523.
- Armona, Luis C, Andreas Fuster, and Basit Zafar**, “Home Price Expectations and Behavior: Evidence from a Randomized Information Experiment,” *The Review of Economic Studies*, 2018, 86 (4), 1371–1410.
- Bachmann, Rüdiger, Tim O Berg, and Eric R Sims**, “Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence,” *American Economic Journal: Economic Policy*, 2015, 7, 1–35.
- Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel**, “The Economic Effects of Social Networks: Evidence from the Housing Market,” *Journal of Political Economy*, 2017, 126 (6), 2224–2276.
- Baker, Scott, R.A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis**, “How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic,” *Working Paper. CEPR Covid Economics 18*, 2020, pp. 73–108.
- Bick, Alexander and Adam Blandin**, “Real Time Labor Market Estimates During the 2020 Coronavirus Outbreak,” *Working Paper*, 2020.
- Bilias, Yannis, Dimitris Georgarakos, and Michael Haliassos**, “Portfolio Inertia and Stock Market Fluctuations,” *Journal of Money, Credit and Banking*, 2010, 42 (4), 715–742.
- Binder, Carola**, “Coronavirus Fears and Macroeconomic Expectations,” *Review of Economics and Statistics*, 2020.
- **and Alex Rodrigue**, “Household Informedness and Long-Run Inflation Expectations: Experimental Evidence,” *Southern Economic Journal*, 2018.
- Bodie, Zvi, Robert C. Merton, and William F. Samuelson**, “Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model,” *Journal of Economic Dynamics and Control*, 1992, 16 (3), 427 – 449.
- Boerma, Job and Jonathan Heathcote**, “Illiquid Wealth and the Timing of Retirement,” *Working Paper*, 2019.
- Bräuer, Konstantin, Andreas Hackethal, and Tobin Hanspal**, “Consuming Dividends,” *Working Paper*, 2020.
- Brunnermeier, Markus K and Stefan Nagel**, “Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals,” *American Economic Review*, 2008, 98 (3), 713–36.
- Bu, Di, Tobin Hanspal, Yin Liao, and Yong Liu**, “Risk Taking During a Global Crisis: Evidence from Wuhan,” *Working Paper. CEPR Covid Economics 5*, 2020, pp. 106–146.

- Calvet, Laurent E., John Y. Campbell, and Paolo Sodini, “Fight or Flight? Portfolio Rebalancing by Individual Investors*,” *The Quarterly Journal of Economics*, 02 2009, 124 (1), 301–348.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia, “Inflation Expectations, Learning and Supermarket Prices: Evidence from Field Experiments,” *American Economic Journal: Macroeconomics*, 2017, 9 (3), 1–35.
- Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten van Rooij, “Wealth Shocks and MPC Heterogeneity,” *Working Paper*, 2019.
- Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Maarten van Rooij, “How Does Consumption Respond to News about Inflation? Field Evidence from a Randomized Control Trial,” *Working Paper*, 2019.
- , – , – , and Michael Weber, “Forward Guidance and Household Expectations,” *Working Paper*, 2020.
- , Yuriy Gorodnichenko, and Michael Weber, “Monetary Policy Communications and Their Effects on Household Inflation Expectations,” *Working Paper*, 2019.
- , – , and – , “The Cost of the COVID-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending,” *Working Paper. CEPR Covid Economics 20*, 2020, pp. 1–51.
- , – , and – , “Labor Markets during the Covid-19 Crisis: A Preliminary View,” *Working Paper. CEPR Covid Economics 21*, 2020.
- , – , and Saten Kumar, “How Do Firms Form Their Expectations? New Survey Evidence,” *American Economic Review*, September 2018, 108 (9), 2671–2713.
- Cox, Natalie, Diana Farrell, Peter Ganong, and Fiona Greig, “The Initial Household Spending Response to COVID-19: Evidence from Credit Card Transactions,” *Working Paper*, 2020.
- D’Acunto, Francesco, “Identity and Choice under Risk,” *Working Paper*, 2020.
- , Daniel Hoang, and Michael Weber, “Managing Households’ Expectations with Simple Economic Policies,” *Working Paper*, 2019.
- D’Acunto, Francesco, Ulrike Malmendier, and Michael Weber, “Gender Roles and the Gender Expectations Gap,” *Working Paper*, 2020.
- D’Acunto, Ulrike Malmendier, Juan Ospina, and Michael Weber, “Exposure to Daily Price Changes and Inflation Expectations,” *Working Paper*, 2019.
- Das, Sreyoshi, Camelia M Kuhnen, and Stefan Nagel, “Socioeconomic Status and Macroeconomic Expectations,” *Working Paper*, 2017.
- de Quidt, Jonathan, Johannes Haushofer, and Christopher Roth, “Measuring and Bounding Experimenter Demand,” *American Economic Review*, 2018, 108 (11), 3266–3302.

- Di Maggio, Marco, Amir Kermani, and Kaveh Majlesi, “Stock Market Returns and Consumption,” *Journal of Finance*, 2019.
- Dietrich, Alexander, Keith Kuester, Gernot Müller, and Raphael Schoenle, “News and Uncertainty about COVID-19: Survey Evidence and Short-Run Economic Impact,” *Working Paper*, 2020.
- Dominitz, Jeff and Charles F Manski, “Expected Equity Returns and Portfolio Choice: Evidence from the Health and Retirement Study,” *Journal of the European Economic Association*, 2007, 5 (2-3), 369–379.
- Fetzer, Thiemo, Lukas Hensel, Johannes Hermle, and Christopher Roth, “Coronavirus Perceptions and Economic Anxiety,” *Working Paper*, 2020.
- Fuster, Andreas, Greg Kaplan, and Basit Zafar, “What Would You Do With \$500? Spending Responses to Gains, Losses, News and Loans,” *Working Paper*, 2020.
- , Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar, “Expectations with Endogenous Information Acquisition: An Experimental Investigation,” *Working Paper*, 2019.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen P Utkus, “Five Facts About Beliefs and Portfolios,” *Working Paper*, 2020.
- , –, –, and –, “Inside the Mind of a Stock Market Crash,” *Working Paper*, 2020.
- Goldfayn-Frank, Olga and Johannes Wohlfart, “Expectation Formation in a New Environment: Evidence from the German Reunification,” *Journal of Monetary Economics*, 2019.
- Gollier, Christian, “Time Diversification, Liquidity Constraints, and Decreasing Aversion to Risk on Wealth,” *Journal of Monetary Economics*, 2002, 49 (7), 1439–1459.
- Gormsen, Niels Joachim and Ralph SJ Koijen, “Coronavirus: Impact on Stock Prices and Growth Expectations,” *Working Paper*, 2020.
- Greenwood, Robin and Andrei Shleifer, “Expectations of Returns and Expected Returns,” *The Review of Financial Studies*, 2014, 27 (3), 714–746.
- Guiso, Luigi and Paolo Sodini, “Household Finance: An Emerging Field,” in “Handbook of the Economics of Finance,” Vol. 2, Elsevier, 2013, pp. 1397–1532.
- , Paola Sapienza, and Luigi Zingales, “Time Varying Risk Aversion,” *Journal of Financial Economics*, 2018, 128 (3), 403–421.
- Haliassos, Michael and Carol C. Bertaut, “Why do so Few Hold Stocks?,” *The Economic Journal*, 1995, 105 (432), 1110–1129.
- Hoynes, Hilary, Douglas L Miller, and Jessamyn Schaller, “Who Suffers During Recessions?,” *The Journal of Economic Perspectives*, 2012, 26 (3), 27–47.
- Hudomiet, Peter, Gábor Kézdi, and Robert J Willis, “Stock Market Crash and Expectations of American Households,” *Journal of Applied Econometrics*, 2011, 26 (3), 393–415.


- Imas, Alex**, “The Realization Effect: Risk-taking After Realized versus Paper Losses,” *American Economic Review*, 2016, 106 (8), 2086–2109.
- Jappelli, Tullio and Mario Padula**, “The Consumption and Wealth Effects of an Unanticipated Change in Lifetime Resources,” *Management Science*, 2015, 62 (5), 1458–1471.
- Kuchler, Theresa and Basit Zafar**, “Personal Experiences and Expectations about Aggregate Outcomes,” *Journal of Finance*, 2019, 74 (5), 2491–2542.
- , **Dominic Russel, and Johannes Stroebel**, “The Geographic Spread of COVID-19 Correlates with Structure of Social Networks as Measured by Facebook,” *Working Paper*, 2020.
- Kuhn, Moritz, Moritz Schularick, and Ulrike Steins**, “Income and Wealth Inequality in America,” *Journal of Political Economy*, 2019.
- Laudenbach, Christine, Benjamin Loos, Jenny Pirschel, and Johannes Wohlfart**, “The Trading Response of Individual Investors to Local Bankruptcies,” *Working Paper*, 2020.
- Maćkowiak, Bartosz and Mirko Wiederholt**, “Lack of Preparation for Rare Events,” *Journal of Monetary Economics*, 2018, 100, 35–47.
- Madrian, Brigitte C and Dennis F Shea**, “The Power of Suggestion: Inertia in 401 (k) Participation and Savings Behavior,” *Quarterly Journal of Economics*, 2001, 116 (4), 1149–1187.
- Malmendier, Ulrike and Stefan Nagel**, “Depression Babies: Do Macroeconomic Experiences Affect Risk-Taking?,” *The Quarterly Journal of Economics*, 2011, 126 (1), 373–416.
- and – , “Learning from Inflation Experiences,” *The Quarterly Journal of Economics*, 2016, 131 (1), 53–87.
- Meeuwis, Maarten, Jonathan A Parker, Antoinette Schoar, and Duncan I Simester**, “Belief Disagreement and Portfolio Choice,” *Working Paper*, 2019.
- Mian, Atif, Amir Sufi, and Nasim Khoshkhoh**, “Partisan Bias, Economic Expectations, and Household Spending,” *Working Paper*, 2018.
- Roth, Christopher and Johannes Wohlfart**, “How Do Expectations About the Macroeconomy Affect Personal Expectations and Behavior?,” *Review of Economics and Statistics*, 2019.
- Vissing-Jørgensen, Annette**, “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 2002, 110 (4), 825–853.
- Vissing-Jørgensen, Annette**, “Perspectives on Behavioral Finance: Does “Irrationality” Disappear with Wealth? Evidence from Expectations and Actions,” *NBER Macroeconomics Annual*, 2003, 18, 139–194.
- Weber, Martin, Elke U Weber, and Alen Nosić**, “Who Takes Risks When and Why: Determinants of Changes in Investor Risk Taking,” *Review of Finance*, 2013, 17 (3), 847–883.

Online Appendix: Exposure to the COVID-19 Stock Market Crash and its Effect on Household Expectations

Tobin Hanspal¹ Annika Weber² Johannes Wohlfart³

A Additional figures

Figure A1: US stock market and number of initial jobless claims around the survey period




Notes: This figure displays the number of initial jobless claims (in thousands, left axis) and the development of the S&P500 stock market index (index points, right axis) over the first 19 weeks in 2020, on a weekly basis. The April 6-13 survey period is highlighted in light red.

¹Tobin Hanspal, Department of Finance, Accounting and Statistics, WU Vienna University of Economics and Business, e-mail: tobin.hanspal@wu.ac.at

²Annika Weber, Department of Finance, Goethe University Frankfurt, e-mail: annika.weber@hof.uni-frankfurt.de


³Johannes Wohlfart, Department of Economics and CEBI, University of Copenhagen, CESifo, e-mail: johannes.wohlfart@econ.ku.dk

Figure A2: Information Treatment *FinCrisisInfo*



Notes: This figure illustrates the information treatment screen, providing an example of the *FinCrisisInfo* treatment arm. The information treatment includes a dynamic figure contrasting the respondent's prior belief (in dark orange, on the right) with the actual number of years it took for the US stock market to recover to its levels before the 2007-2009 Financial Crisis (in yellow, on the left). Recovery durations for the three different information treatments *FinCrisisInfo*, *BlackMondayInfo* and *DotComInfo* are calculated based on monthly time series data of the S&P500.

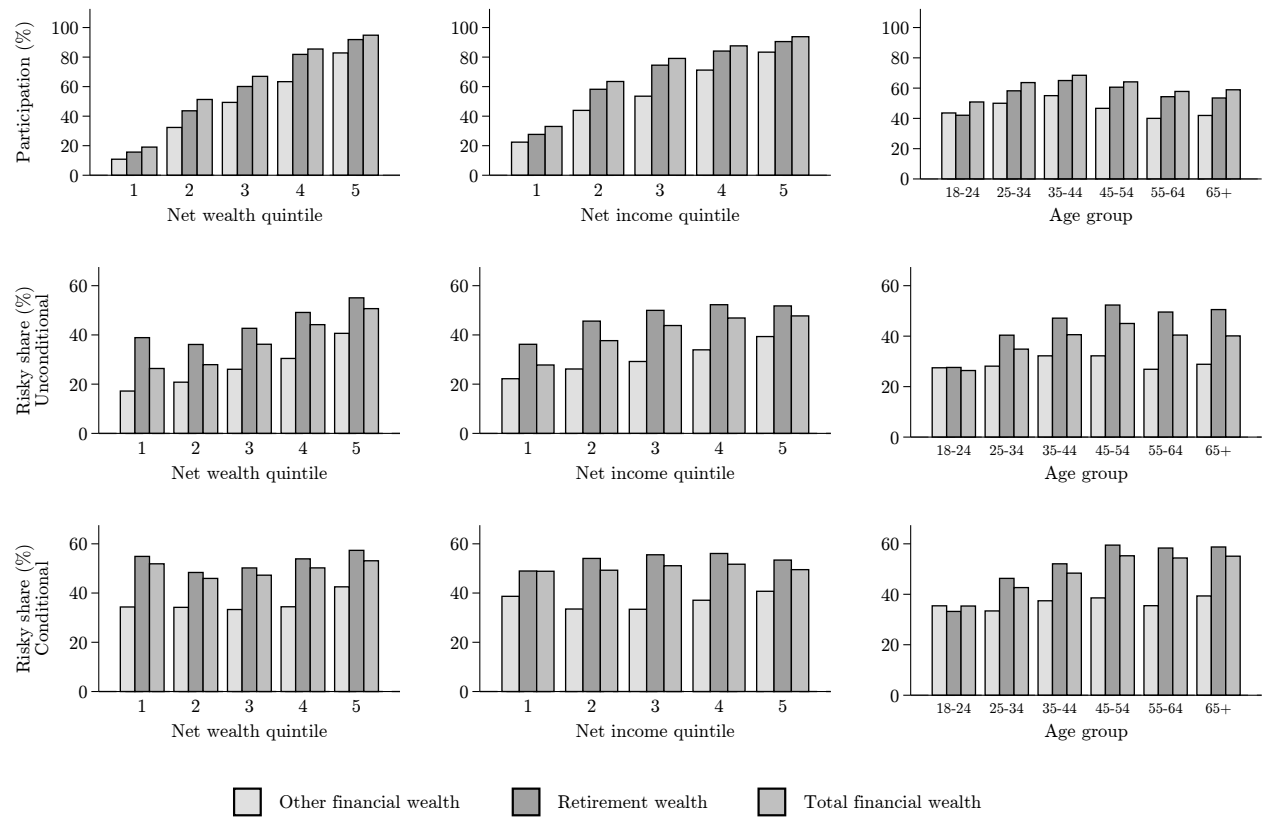
Figure A3: Financial assets and incomes across groups



Notes: This figure displays the average value of financial assets (top row) and gross household income during the first quarter of 2020 (bottom row), by quintile of the pre-crisis net worth distribution (left column), by quintile of the pre-crisis net income distribution (middle column), and by age group (right column). Values of financial assets are displayed separately for financial assets outside of retirement accounts, for financial assets in retirement accounts, and for the combined value of all financial assets. The sample is the full sample without missings in the relevant survey questions.

Covid Economics 23, 28 May 2020: 188-262

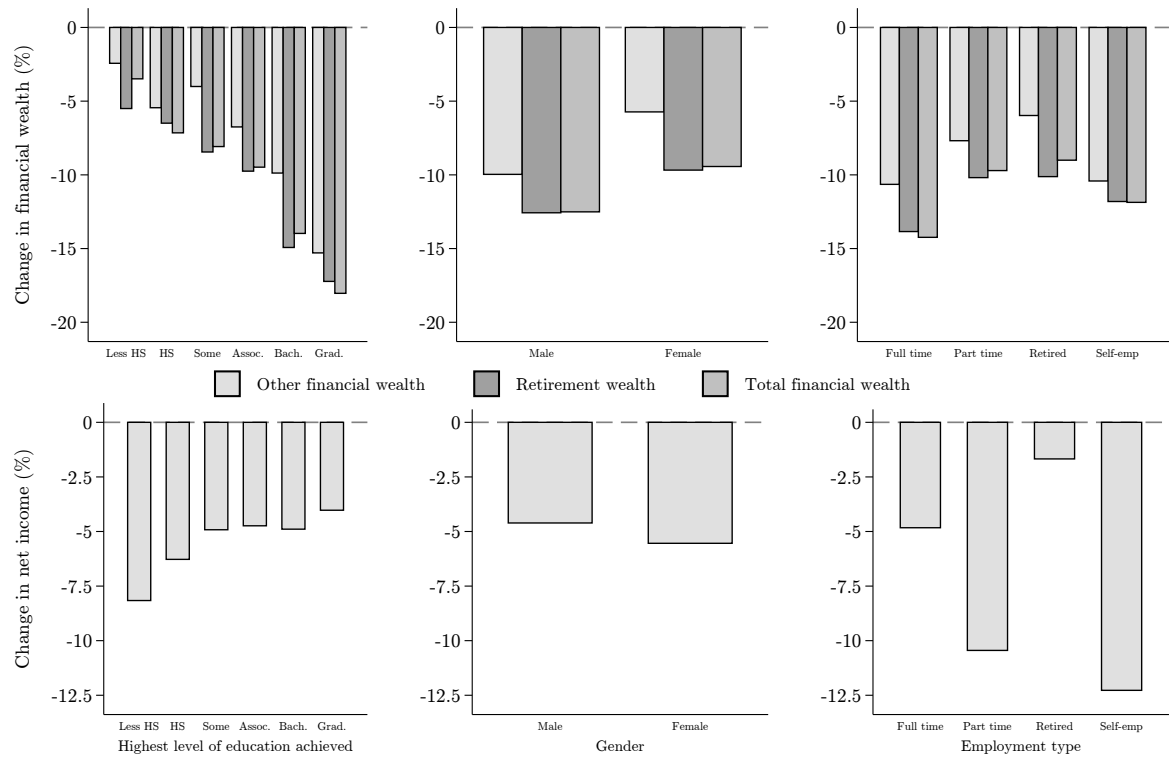
Figure A4: Participation and stock share of financial wealth across groups



Notes: This figure displays participation, and the share of financial wealth held, in stocks or mutual funds by quintiles of pre-crisis net wealth (left column) and of pre-crisis net income, (middle), and by age group (right), respectively. The top row plots the rate of participation in stocks and stock mutual funds in the full sample. The middle row plots the unconditional equity share, and the bottom plots the conditional equity share including only respondents that report positive holdings of stocks or stock mutual funds as of January 2020. Equity shares are displayed separately for financial assets outside retirement accounts, for financial assets in retirement accounts, and for the combined value of financial assets. The sample includes all respondents without missings in the relevant survey questions.

Covid Economics 23, 28 May 2020: 188-262

Figure A5: Income and wealth shocks across groups

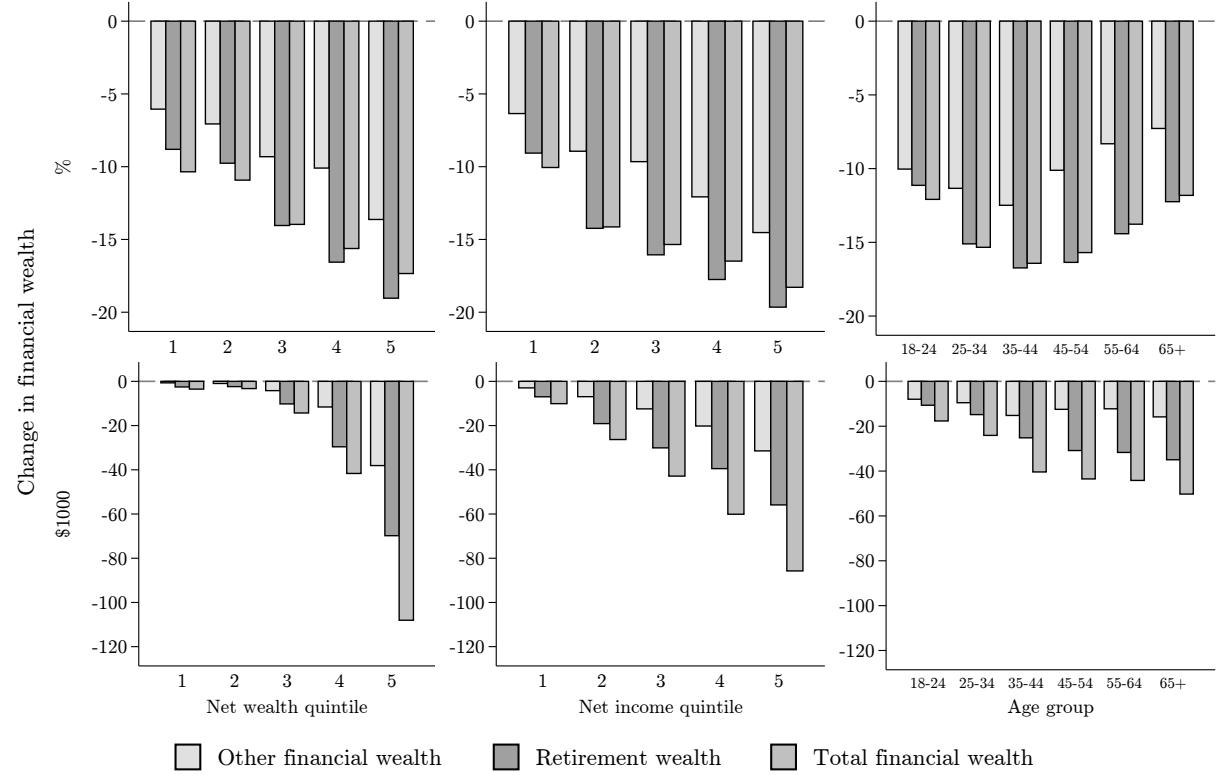


Notes: This figure displays the change in the value of financial assets due to the February/March 2020 stock market drop until the survey date in percentage terms as amounts in USD (top panel) and unexpected changes in net household incomes during the first quarter of 2020 in percentages and USD amounts (bottom panel), by highest level of education achieved (left column), gender (middle column), and pre-crisis employment type (right column). Changes in the value of financial assets are displayed separately for financial assets outside of retirement accounts, for financial assets in retirement accounts, and for the combined value of all financial assets. Changes in value of financial assets are net capital losses for the majority of respondents, and net capital gains for a small fraction of respondents. We trim reported shocks to income and financial wealth at the 2nd and 98th percentiles. The sample is the full sample without missings in the relevant survey questions. Note that the average percent reduction in overall financial wealth can be larger than both average percent reductions for the individual components. This is due to the fact that we coded those with no wealth in a given category as having experienced a shock of zero percent in that category. These cases occur particularly in groups with lower wealth holdings.

Covid Economics 23, 28 May 2020: 188-262

COVID ECONOMICS
 METEOR AND REAL-TIME PAPERS

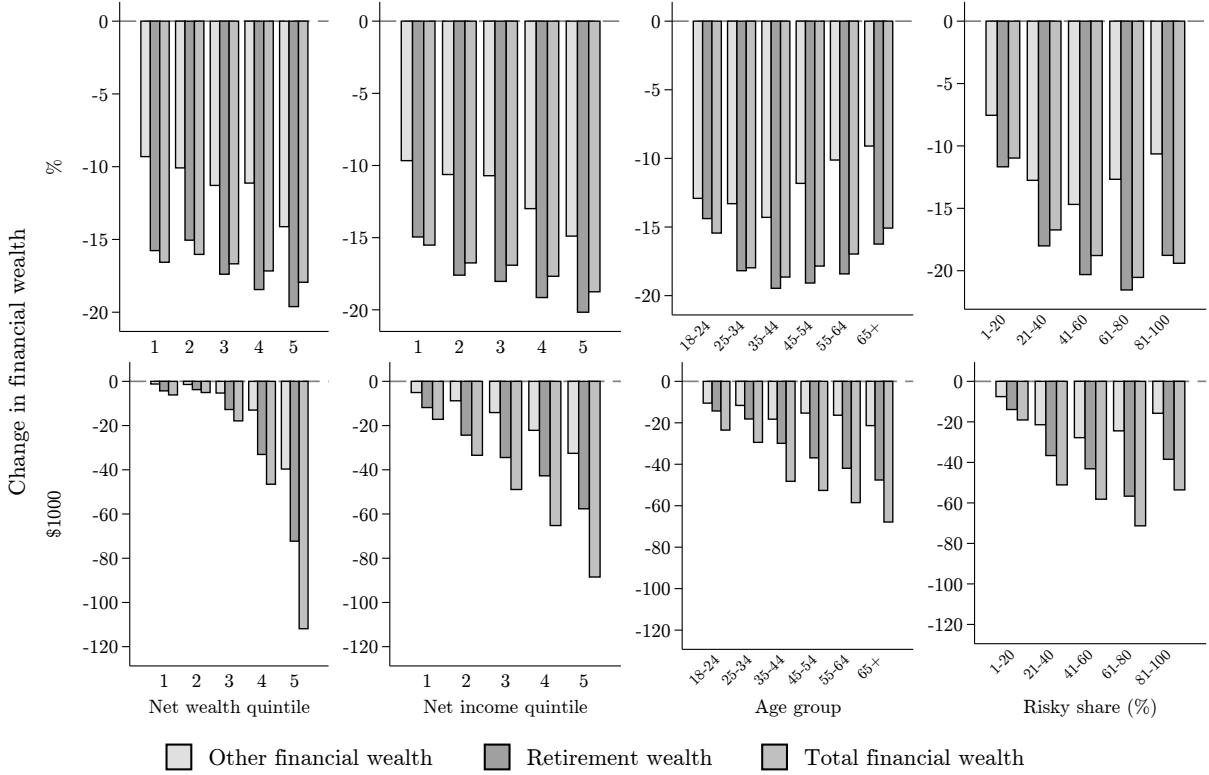
Figure A6: Conditional wealth shocks across groups (Financial wealth > 0)



Notes: This figure displays the change in the value of financial assets due to the February/March 2020 stock market drop until the survey date in percentage terms and as amounts in USD, by quintile of the pre-crisis net worth distribution (left column), by quintile of the pre-crisis net income distribution (middle column) and by age group (right column). Changes in the value of financial assets are displayed separately for financial assets outside of retirement accounts, for financial assets in retirement accounts, and for the combined value of all financial assets. The values are conditional on positive financial wealth in January 2020. Changes in the value of financial assets are net capital losses for the majority of respondents, and net capital gains for a small fraction of respondents. We trim reported shocks to financial wealth at the 2nd and 98th percentiles. The sample is the full sample without missings in the relevant survey questions. Note that the average percent reduction in overall financial wealth can be larger than both average percent reductions for the individual components. This is due to the fact that we coded those with no wealth in a given category as having experienced a shock of zero percent in that category. These cases occur particularly in groups with lower wealth holdings.

Covid Economics 23, 28 May 2020: 188-262

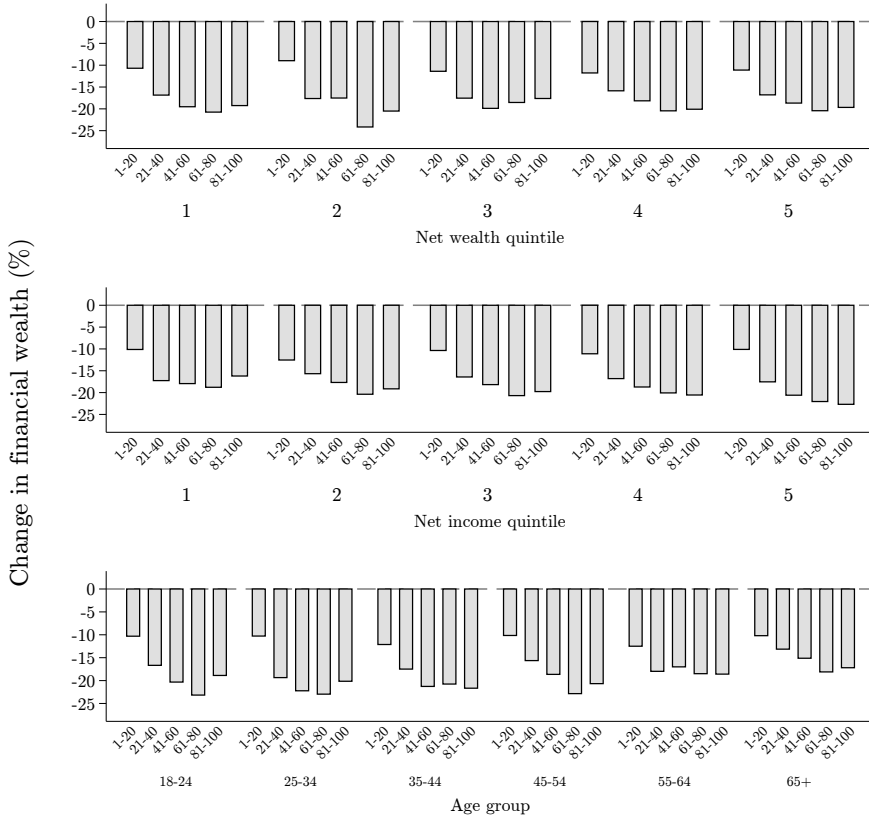
Figure A7: Conditional wealth shocks across groups (Risky share > 0)



Notes: This figure displays the change in the value of financial assets due to the February/March 2020 stock market drop until the survey date in percentage terms, and as amounts in USD, by quintile of the pre-crisis net worth distribution (left column), by age group (third column), and by bin of the equity share in total financial wealth. Changes in the value of financial assets are displayed separately for financial assets outside of retirement accounts, for financial assets in retirement accounts, and for the combined value of financial assets. The values are conditional on positive equity investments in January 2020. Changes in the value of financial assets are net capital losses for the majority of respondents, and net capital gains for a small fraction of respondents. We trim reported shocks to financial wealth at the 2nd and 98th percentiles. The sample is the full sample without missings in the relevant survey questions. Note that the average percent reduction in overall financial wealth can be larger than the average percent reductions for both individual components. This is due to the fact that we coded those with no wealth in a given category as having experienced a shock of zero percent in that category. These cases occur particularly in groups with lower wealth holdings.

Covid Economics 23, 28 May 2020: 188-262


Figure A8: Wealth shocks across groups by risky share



January 2020 risky share allocation

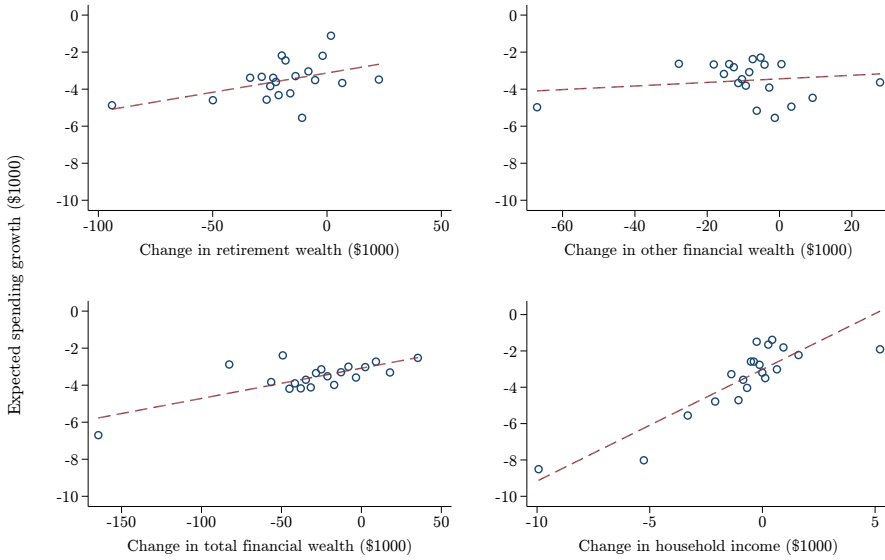
Notes: This figure displays the change in the value of financial assets due to the February/March 2020 stock market drop until the survey date in percentage terms by equity portfolio share bin, separately by quintile of the pre-crisis net worth distribution (top row), by quintile of the pre-crisis net income distribution (middle row), by age group (bottom row). Changes in the value of financial assets are for the combined value of financial assets inside and outside of retirement accounts. The values are conditional on positive equity investments in January 2020. Changes in the value of financial assets are net capital losses for the majority of respondents, and net capital gains for a small fraction of respondents. We trim reported shocks to financial wealth at the 2nd and 98th percentiles. The sample is the full sample without missings in the relevant survey questions.

Figure A9: Job losses across groups



Notes: This figure displays the percentage of respondents who lost their job since January 2020 by pre-crisis net wealth quintile (top left), pre-crisis net income quintile (top right), age (middle left), education (middle right), pre-crisis employment type (bottom left), and gender (bottom right). The question is presented only to respondents who report to have been employed as of January 2020. The sample is the full sample without missings in the relevant survey questions.

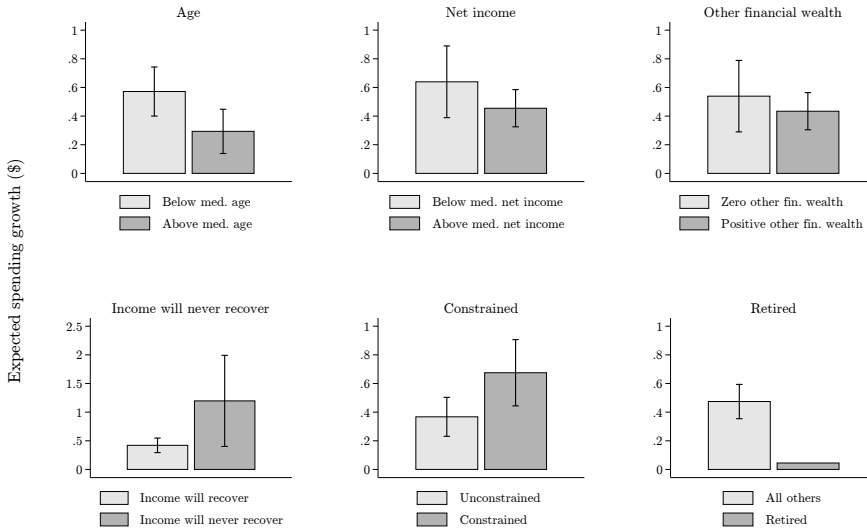
Figure A10: Effects of wealth and income shocks on expected spending



Notes: This figure shows binned scatter plots of the association of shocks to the respondent’s household financial wealth and net income with expected growth of total nominal household spending in 2020 compared to 2019. The plots in the top row and in the bottom right are based on specification 2 shown in Table 3, which jointly includes shocks to income, retirement financial wealth and other financial wealth. The plot in the bottom left is based on a similar specification replacing the shocks to financial wealth in retirement accounts and in non-retirement accounts with the shock to overall financial wealth. The outcome is expected household spending growth in dollars, trimmed at the 2nd and 98th percentiles. Dollar changes are constructed from survey questions for retirement and other financial wealth and for income (assuming that the respondent expected a quarter of her 2019 income for the first quarter of 2020), and for spending from the survey question on percent changes and estimates of the level of spending of different groups in 2019 from the CEX. All specifications are based on respondents in the four control arms, who have not received any information. All specifications control for gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the equity share in total financial assets, investment experience, Census region, survey date, and the survey arm.

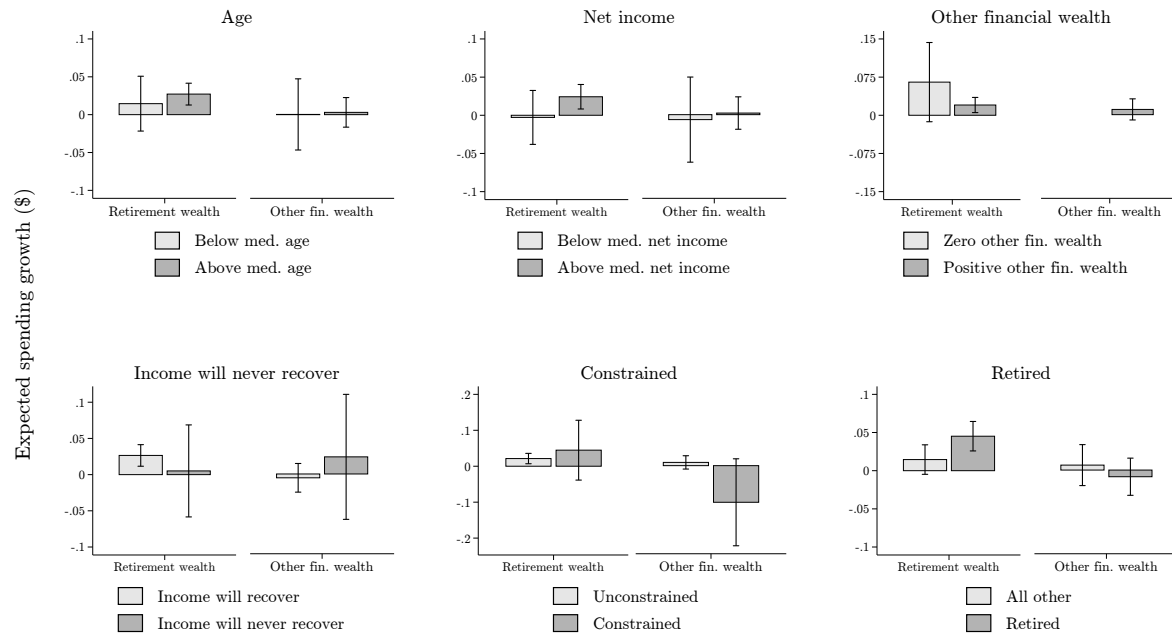
Covid Economics 23, 28 May 2020: 188-262

Figure A11: Effects of income shocks on expected spending: Heterogeneity



Notes: This figure shows heterogeneity in the effect of income shocks on expected growth of total nominal household spending in 2020 compared to 2019 across groups. The plots are based on the 2SLS specification shown in Table 3 column 3, where the respondent’s expected dollar change in household income from 2019 to 2020 is instrumented with the unexpected shock to household income in the first quarter, estimated for different subsamples. The different panels show median splits (below and equal to median vs. strictly above median) according to age (top left) and pre-crisis household net income (top middle), by an indicator (0 vs. 1) for holding other (non-retirement) financial wealth in January 2020 (top right), by an indicator for believing household income will never recover (bottom left), by an indicator for the respondent’s household facing credit constraints (bottom middle), and being retired (bottom right). The outcome is expected household spending growth in dollars, trimmed at the 2nd and 98th percentiles. Dollar changes are constructed from survey questions for retirement and other financial wealth and for income (assuming that the respondent expected a quarter of her 2019 income in the first quarter of 2020), and for spending from the survey question on percent changes and estimates of the level of spending of different groups in 2019 from the CEX. In the bottom right panel we do not include confidence bands as the interval is large and insignificant. For all others we include 90% confidence intervals. All specifications are based on respondents in the four control arms, who have not received any information. All specifications control for gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the equity share in total financial assets, investment experience, Census region, survey date, and the survey arm.

Figure A12: Effects of wealth shocks on expected spending: Heterogeneity



Notes: This figure shows heterogeneity in the association of shocks to the respondent's household retirement financial wealth and non-retirement financial wealth and the expected growth of total nominal household spending in 2020 compared to 2019. The plots for wealth shocks are based on the reduced-form specification controlling for the first quarter-income shock shown in Table 3 column 2, estimated for different subsamples. The outcome is expected household spending growth in dollars, trimmed at the 2nd and 98th percentiles. Dollar changes are constructed from survey questions for retirement and other financial wealth and for spending from the survey question on percent changes and estimates of the level of spending of different groups in 2019 from the CEX. We plot coefficients on changes to a respondent's household retirement financial wealth and other financial wealth by median splits (below and equal to median vs. strictly above median) according to age (top left) and pre-crisis household net income (top middle), by an indicator (0 vs. 1) for holding other (non-retirement) financial wealth in January 2020 (top right), by an indicator for believing individual income will never recover (bottom left), by an indicator for the respondent's household facing credit constraints (bottom middle), and by being retired (bottom right). 90% confidence intervals are displayed in all figures. All specifications are based on respondents in the four control arms, who have not received any information. All specifications control for gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, risky portfolio share, investment experience, Census region, survey date, and the survey arm.

Covid Economics 23, 28 May 2020: 188-262


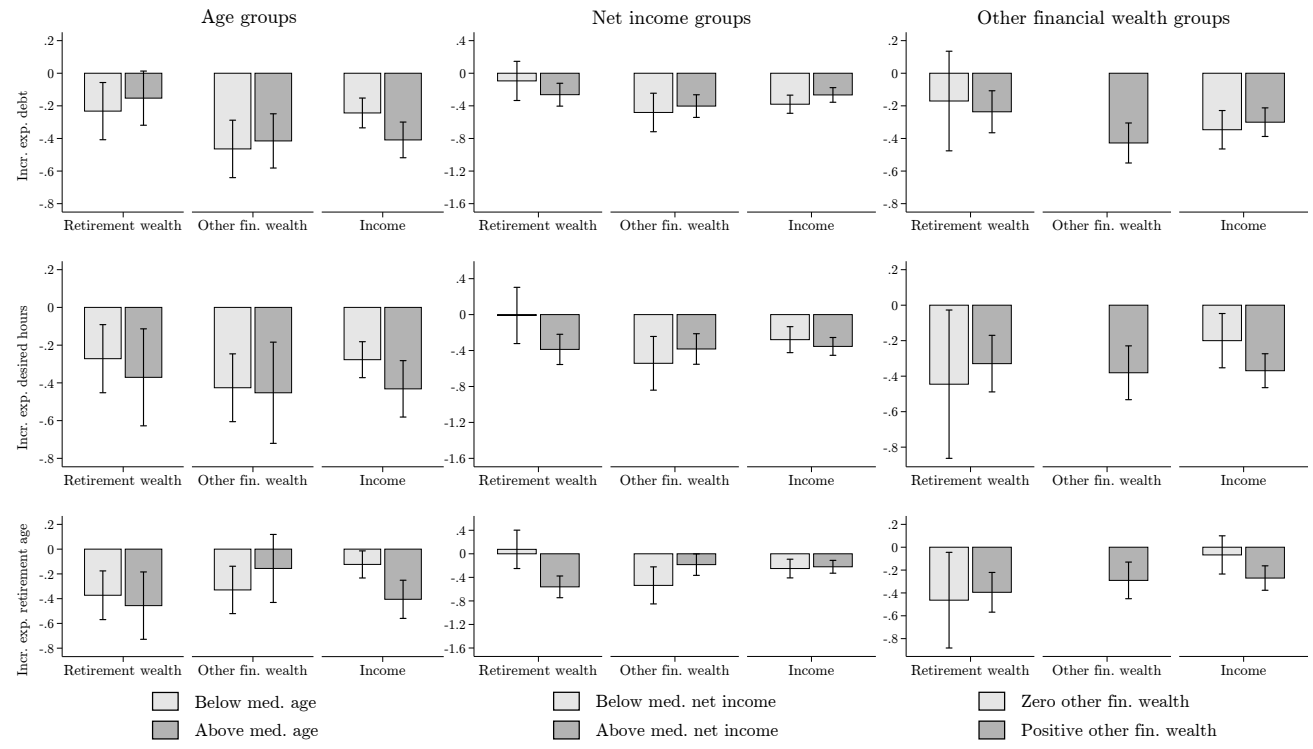


Figure A13: Effects of wealth and income shocks on economic plans: Heterogeneity I

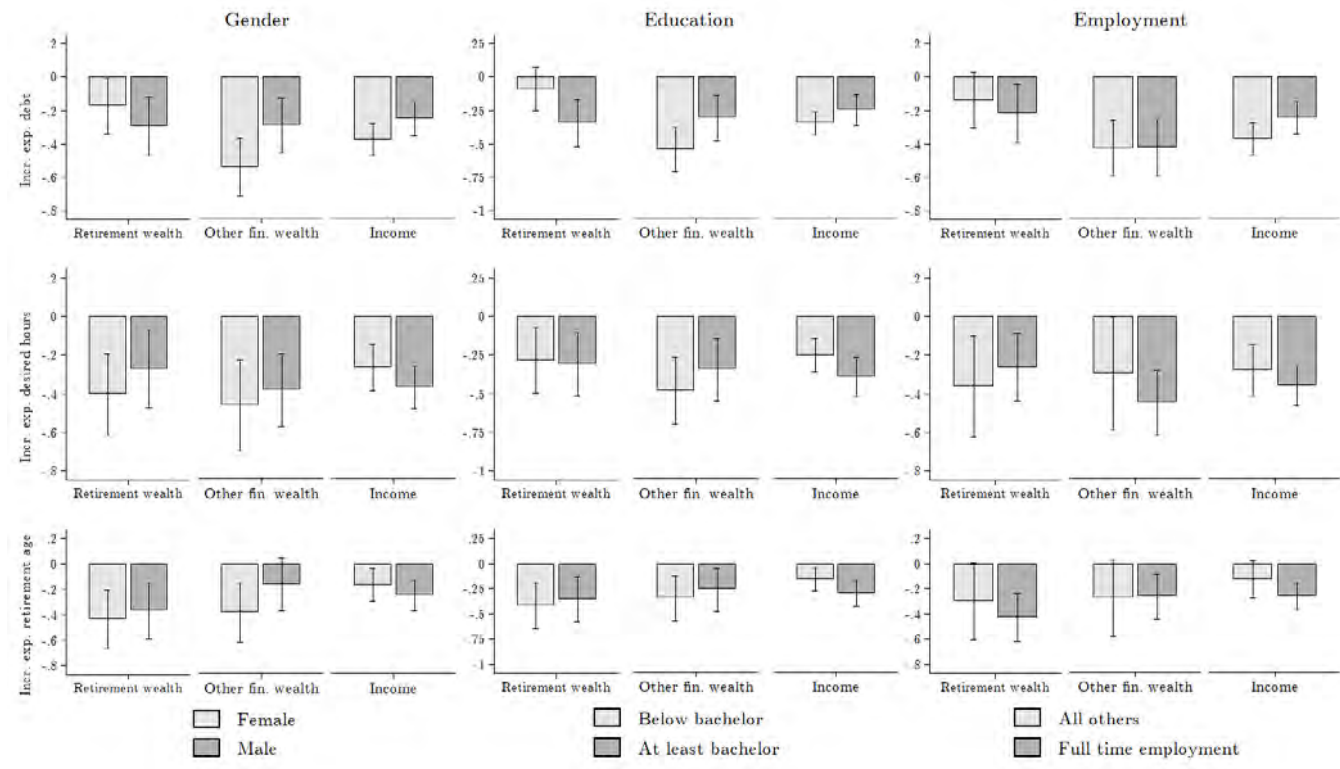


Notes: This figure shows heterogeneity in the effect of shocks to the respondent's household financial wealth and net income on expected economic decisions. We plot the coefficients from specifications regressing indicators for whether the coronavirus crisis increases the respondent's expectations about outstanding household debt by the end of 2020 (top row), expected desired working hours over the next years (middle row) and expected retirement age (bottom row), all coded as 0 or 100, on changes to a household's retirement financial wealth, other financial wealth, and net household income. We plot these coefficients across columns by median splits (below and equal to median vs. strictly above median) according to age (left column) and pre-crisis household net income (middle), and by an indicator (0 vs. 1) for holding other (non-retirement) financial wealth in January 2020 (right). 90% confidence intervals are displayed in all figures. All specifications are based on respondents in the four control arms, who have not received any information. All specifications control for gender, age, employment status, being the household's main earner, being the household's financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the share of equity in total financial assets, investment experience, Census region, survey date, and the survey arm.

Covid Economics 23, 28 May 2020: 188-262

COVID ECONOMICS
 POLICY FEEDBACK AND FEELING PAPERS

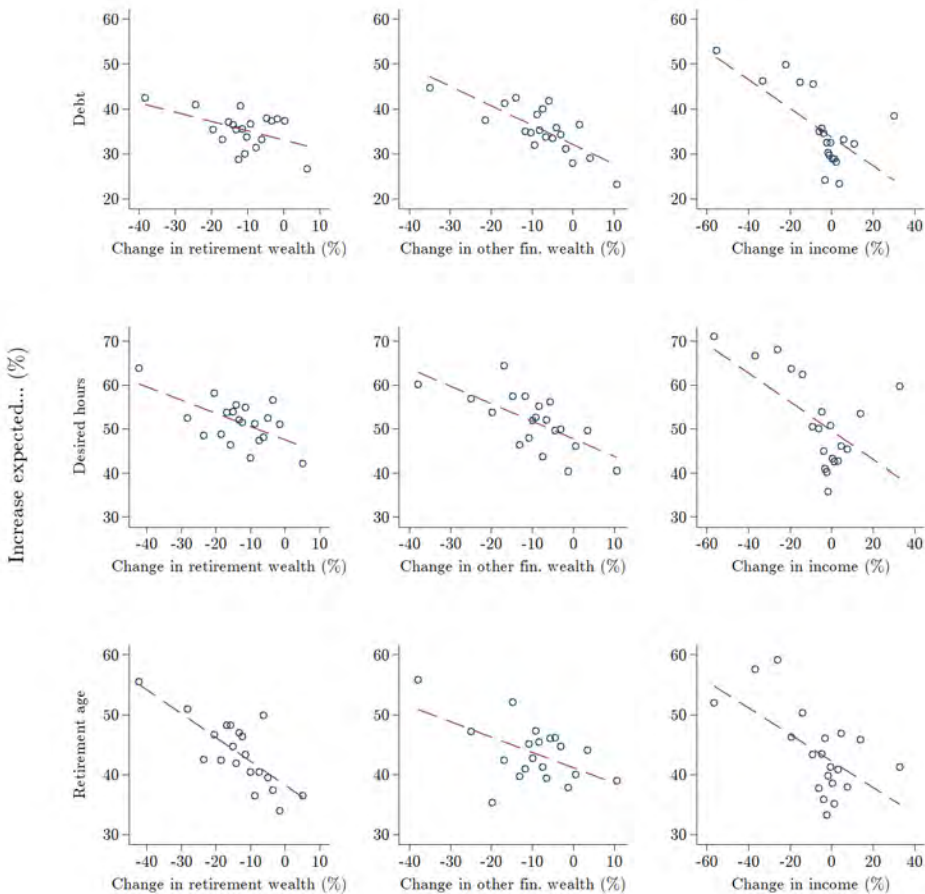
Figure A14: Effects of wealth and income shocks on economic plans: Heterogeneity II



Notes: This figure shows heterogeneity in the effect of shocks to the respondent’s household financial wealth and net income on expected economic decisions. We plot the coefficients from specifications regressing indicators for whether the coronavirus crisis increases the respondent’s expectations about outstanding household debt by the end of 2020 (top row), expected desired working hours over the next years (middle row) and expected retirement age (bottom row), all coded as 0 or 100, on changes to a household’s retirement financial wealth, other financial wealth, and net household income. We plot these coefficients across columns by indicators for gender (left column), education of at least a bachelors degree (middle), and being full-time employed pre-crisis (right). 90% confidence intervals are displayed on all figures. All specifications are based on respondents in the four control arms, who have not received any information. All specifications control for gender, age, employment status, being the household’s main earner, being the household’s financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the equity share in total financial assets, investment experience, Census region, survey date, and the survey arm.

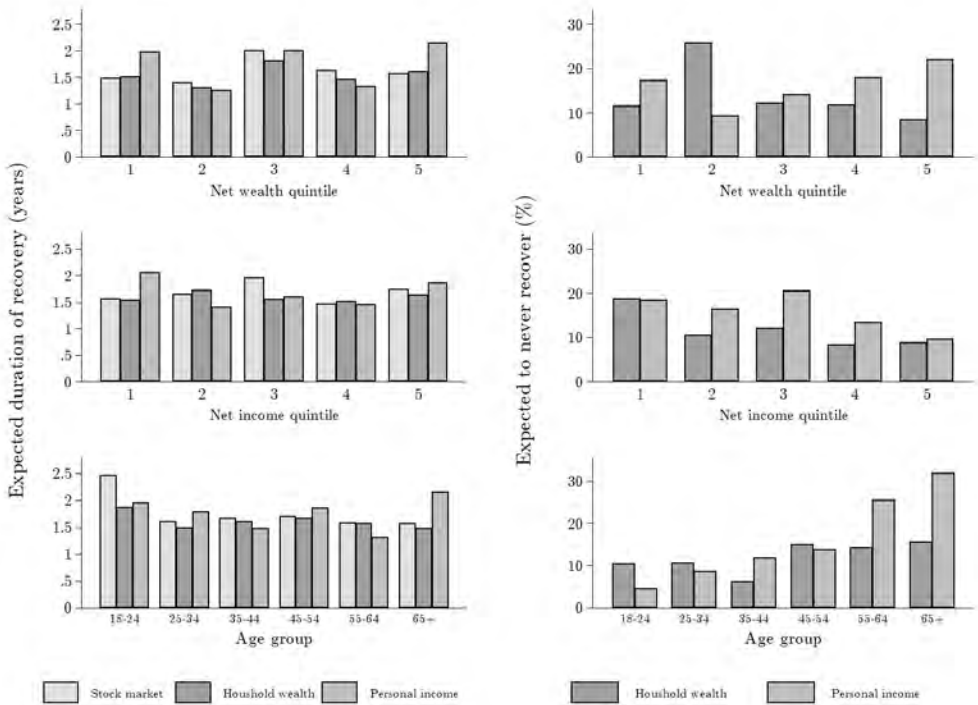
Covid Economics 23, 28 May 2020: 188-262

Figure A15: Effects of wealth and income shocks on economic plans



Notes: This figure shows binned scatter plots of the association of shocks to the respondent's household financial wealth and net income with expected economic decisions. The outcomes are a dummy indicating whether the coronavirus crisis increases the respondent's expectations about outstanding household debt by the end of 2020 (left column), expected desired working hours over the next years (middle column, only if in labor force), or expected retirement age (right column, only if in labor force), all coded as 0 or 100. The underlying regressions are specifications 4, 5, and 6 in Table 3, which jointly include changes to retirement financial wealth, to other financial wealth, and to household net income. For each outcome (debt, desired working hours, and retirement age), we plot coefficients on changes in retirement wealth (top), changes in other financial wealth (middle), and by changes in household income (bottom) in percentage terms, respectively. All specifications control for all other changes to household income and/or wealth, trimmed at the 2nd and 98th percentiles, gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, log of total financial wealth, of real estate wealth, and of debt, share of total financial wealth in retirement accounts, borrowing constraints, stock market participation, stock shares in retirement and other accounts, investment experience, Census region, survey date, and the survey arm. All specifications are based on respondents in the four control arms, who have not received any information.

Figure A16: Expected duration of recovery across groups



Notes: This figure displays respondents' subjective expectations about the duration of the recovery in years for the US stock market, the respondent's pre-crisis household net wealth, and the respondent's pre-crisis household net income (left column) and the fractions of respondents who believe that their household net wealth or income will never recover (right column) by the pre-crisis net household income distribution (top row), by quintile of the pre-crisis net household income distribution (middle row), and by age group (bottom row). The figures on expected recovery duration of the stock market and own wealth condition on those who have made financial losses, while the figures on income recovery duration condition on those who have incurred income losses. The sample consists of respondents in the pure control group, who have not received any questions or information on past crashes before answering to the questions on expected recovery duration.

B Additional tables

Table A1: Treatment and sample details

Treatment	Sample	Information	Respondents
Financial Crisis 2007	Information	5.5 years	1,055
	Control	-	1,050
Dot-com bubble	Information	7 years	1,041
	Control	-	1,063
Black Monday	Information	2 years	1,091
	Control	-	1,058
Pure control	Control	-	1,089
Total			7,447

Notes: The table gives an overview of the various control and treatment arms in the survey. The final number of participants is listed in the column *Respondents*.

Table A2: Integrity of treatment randomization

	FinCrisis Info	FinCrisis Control	DotCom Info	DotCom Control	Black Monday Info	Black Monday Control	Pure Control	P-value	Obs.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Female	0.51	0.50	0.52	0.51	0.51	0.52	0.53	0.815	7,447
Age	48.32	48.65	48.53	48.94	48.21	48.84	48.94	0.907	7,447
Bachelor's degree or higher	0.38	0.38	0.38	0.40	0.39	0.39	0.39	0.981	7,447
HH income (gross, USD)	80,074	80,836	82,281	81,628	81,753	79,662	80,472	0.939	7,417
Republican	0.37	0.37	0.39	0.36	0.36	0.36	0.37	0.817	7,447
Democrat	0.38	0.39	0.38	0.39	0.39	0.40	0.38	0.928	7,447
Region Midwest	0.25	0.24	0.25	0.23	0.25	0.27	0.23	0.470	7,447
- South	0.33	0.30	0.35	0.34	0.33	0.33	0.34	0.190	7,447
- West	0.21	0.25	0.21	0.22	0.20	0.21	0.21	0.273	7,447
Stock investor	0.61	0.64	0.61	0.62	0.61	0.59	0.61	0.565	7,447

Notes: The table shows respondent characteristics across the 7 treatment and control arms. Column 8 shows the p-Value of an F-test that all coefficients are zero when jointly regressing the respective characteristics on all treatment dummies.

Table A3: Average losses across groups and samples

Sort	Sample	Difference in losses between 5th and 1st quintile (%)		
		Other financial wealth	Retirement wealth	Total financial wealth
		(1)	(2)	(3)
Net wealth	Unconditional	-11.33	-15.70	-13.48
	Financial wealth > 0	-7.58	-10.23	-6.99
	Stocks > 0	-4.81	-3.85	-1.37
Net income	Unconditional	-10.50	-13.93	-12.07
	Financial wealth > 0	-8.18	-10.58	-8.23
	Stocks > 0	-5.23	-5.21	-3.23
Age group	Unconditional	0.21	-3.74	-2.67
	Financial wealth > 0	1.72	-3.28	-1.68
	Stocks > 0	2.79	-4.03	-1.52

Notes: The table shows the difference in percentage financial losses between the 5th and 1st quintile of the distributions of household net wealth and household net income, and between the lowest (18-24) and highest (65 and older) age categories. Differences in the percentage losses are displayed separately for non-retirement financial wealth (column 1), other financial wealth (column 2), and total financial wealth (column 3). We display differences between the two extreme quintiles (age categories) unconditionally for the entire sample, for the subsample of individuals with positive financial wealth holdings as of January 2020, and for the subsample of individuals with positive equity investments as of January 2020. Within each (sub-)sample, we include all respondents with nonmissing survey responses.

Table A4: Determinants of realized and planned adjustments to stock share: Pure control only

	Changed stock share	Increased stock share	Decreased stock share	Plan change stock share	Plan incr. stock share	Plan decr. stock share
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Net income (%)	-0.203* (0.106)	-0.022 (0.115)	-0.181 (0.113)	-0.267** (0.122)	0.029 (0.142)	-0.296*** (0.110)
Δ Retirement fin. wealth (%)				0.005 (0.181)	-0.203 (0.176)	0.209 (0.155)
Δ Other fin. wealth (%)				-0.480*** (0.175)	-0.348** (0.172)	-0.132 (0.160)
Ln(Total fin. wealth)	2.733 (1.740)	-1.487 (1.712)	4.220** (1.660)	0.960 (1.769)	0.109 (1.845)	0.851 (1.613)
Stock share in ret. wealth	-0.073 (0.069)	-0.020 (0.061)	-0.053 (0.055)	-0.106 (0.072)	-0.034 (0.063)	-0.071 (0.049)
Stock share in ot. fin wealth	-0.044 (0.065)	-0.015 (0.054)	-0.029 (0.058)	-0.036 (0.065)	-0.039 (0.056)	0.003 (0.053)
Share ret. in tot. fin. wealth	-0.074 (0.089)	-0.034 (0.086)	-0.040 (0.086)	-0.092 (0.093)	-0.063 (0.090)	-0.029 (0.081)
Any loss fin. crisis	4.355 (5.250)	0.934 (5.201)	3.421 (4.740)	-0.202 (5.343)	-4.541 (5.330)	4.339 (4.523)
Big loss fin. crisis	10.581* (5.805)	-1.029 (4.675)	11.611** (5.729)	-8.162 (5.398)	-10.969** (5.026)	2.807 (4.661)
Any loss dot-com	-6.620 (5.424)	-7.927 (5.038)	1.307 (4.961)	-2.290 (5.387)	5.642 (5.251)	-7.933** (3.952)
Big loss dot-com	-12.558 (7.925)	7.262 (6.896)	-19.821*** (6.592)	7.637 (8.543)	2.022 (8.043)	5.615 (6.503)
Any loss Black Monday	7.533 (6.558)	2.634 (5.830)	4.899 (5.917)	6.042 (6.147)	3.738 (5.808)	2.303 (4.492)
Big loss Black Monday	-11.810 (10.551)	-14.421** (6.596)	2.611 (9.362)	-12.090 (9.951)	-10.374 (8.497)	-1.716 (6.631)
Male	11.502** (4.970)	7.733* (4.494)	3.769 (4.297)	1.098 (4.944)	4.701 (4.594)	-3.603 (3.779)
At least bachelor	-22.711 (20.880)	-21.361 (22.314)	-1.350 (15.922)	10.944 (21.657)	-11.292 (25.271)	22.236* (12.633)
Republican	-6.972 (4.556)	-7.019 (4.352)	0.047 (4.268)	0.050 (4.445)	-0.710 (4.424)	0.760 (3.816)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	.221	.135	.057	.279	.108	.102
Observations	550	550	550	503	503	503

Notes: This table shows OLS estimates of the determinants of realized and planned adjustments to the share of equities in total financial assets. The outcomes are dummies indicating whether the respondent's household has made any change, has increased or has decreased the equity share in total financial assets since the beginning of the stock market drop (columns 1-3) and dummies indicating plans to change, increase or decrease the equity share in overall financial assets in the weeks after the survey (columns 4-6), all coded as 0 or 100. All specifications are based on the pure control group, which has not received any information and not answered any questions on past crashes, using only respondents who report positive stockholdings as of January 2020. All specifications control for shocks to income, trimmed at the 2nd and 98th percentiles, dummies for having lost any wealth or substantial wealth during past stock market crashes, gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, log of total financial wealth, of real estate wealth, and of debt, share of total financial wealth in retirement accounts, borrowing constraints, stock market participation, stock shares in retirement and other accounts, investment experience, Census region, survey date, and the survey arm. Columns 4-6 additionally control for shocks to retirement and other financial wealth, trimmed at the 2nd and 98th percentiles. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

Table A5: Income and wealth shocks and expected economic decisions: Pure control only

	Exp. spend. growth (%)	Exp. spend. growth (\$)	Exp. spend. growth (\$)	Incr. exp. debt	Incr. exp. desired hours	Incr. exp. retirement age
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Retirement fin. wealth (%)	0.018 (0.068)			-0.148 (0.148)	-0.274 (0.188)	-0.696*** (0.195)
Δ Other fin. wealth (%)	-0.061 (0.070)			-0.443*** (0.140)	-0.371** (0.176)	0.006 (0.179)
Δ Net income (quarterly, %)	0.189*** (0.048)			-0.286*** (0.085)	-0.358*** (0.105)	-0.401*** (0.100)
Δ Retirement fin. wealth (\$)		0.048*** (0.016)	0.039** (0.017)			
Δ Other fin. wealth (\$)		0.011 (0.019)	-0.002 (0.021)			
Δ Net income (quarterly, \$)		0.584*** (0.201)				
Δ Net income (annual, \$)			0.536*** (0.139)			
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	.048	.044	.114	.198	.108	.124
Observations	970	917	906	987	603	603

Notes: This table shows estimates of the association of shocks to the respondent’s household net income and financial wealth with expected economic decisions. The outcomes are expected growth of yearly household spending from 2019 to 2020 in percent, trimmed at the 2nd and 98th percentiles (column 1); expected household spending growth in dollars, trimmed at the 2nd and 98th percentiles (columns 2-3); and dummies indicating whether the coronavirus crisis increases the respondent’s expectations about outstanding household debt by the end of 2020 (column 4), expected desired working hours over the next years (column 5, only if in labor force) or expected retirement age (column 6, only if in labor force), coded as 0 or 100. Dollar changes in columns 2 and 3 are constructed from survey questions for retirement and other financial wealth and for income (assuming that the respondent expected a quarter of her 2019 income in the first quarter of 2020), and for spending from the survey question on percent changes and estimates of the level of spending of different groups in 2019 from the CEX. Columns 1, 2, 4, 5 and 6 show simple OLS estimates. Column 3 shows a 2SLS estimate, where the respondent’s expected dollar change in household income from 2019 to 2020 is instrumented with the unexpected shock to household income in the first quarter. All specifications are based on the pure control group, which has not received any information and not answered any questions on past crashes. All specifications control for shocks to income and financial wealth, trimmed at the 2nd and 98th percentiles, gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, risky portfolio share, investment experience, Census region, survey date, and the survey arm. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

Table A6: Determinants of expectations about the stock market and own wealth: Pure control only

	Stock recovery duration	Stock return: Mean	Stock return: SD	Stock return <-30%	Stock return >30%	Wealth recovery duration	Exp. wealth never to recover	Household financial prospects
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Δ Fin. wealth (%)	-0.010* (0.006)	0.020 (0.058)	-0.024 (0.026)	-0.032 (0.065)	-0.041 (0.086)	-0.022*** (0.005)	-0.229** (0.113)	0.010*** (0.003)
Δ Net income (%)	-0.002 (0.004)	0.087*** (0.033)	-0.005 (0.018)	-0.111** (0.043)	0.087* (0.044)	-0.006** (0.003)	-0.065 (0.067)	0.004** (0.002)
Any loss fin. crisis	-0.138 (0.156)	-1.868 (1.643)	0.234 (0.803)	-1.043 (1.798)	-3.853 (2.417)	0.057 (0.124)	4.437 (3.473)	-0.113 (0.087)
Big loss fin. crisis	0.168 (0.215)	0.693 (2.222)	-0.119 (0.990)	-0.744 (2.414)	3.531 (3.386)	-0.043 (0.170)	2.799 (4.205)	-0.098 (0.111)
Any loss dot-com	0.100 (0.171)	-1.673 (1.776)	0.493 (0.890)	3.121 (2.180)	-0.379 (2.549)	0.229 (0.161)	-3.492 (3.446)	-0.046 (0.097)
Big loss dot-com	-0.537** (0.234)	-0.211 (3.496)	-1.509 (1.497)	-1.645 (3.785)	2.219 (5.506)	-0.302 (0.242)	-3.943 (6.226)	0.340* (0.178)
Any loss Black Monday	0.139 (0.215)	-3.562 (2.268)	0.123 (1.014)	0.860 (2.789)	-6.867** (3.124)	0.138 (0.191)	-3.158 (4.107)	-0.010 (0.117)
Big loss Black Monday	-0.118 (0.254)	0.864 (3.843)	-0.200 (1.496)	0.595 (4.885)	2.592 (4.876)	-0.206 (0.237)	21.798** (8.856)	-0.246 (0.195)
Male	-0.170 (0.169)	-0.267 (1.418)	1.488** (0.682)	-0.556 (1.716)	0.515 (2.033)	-0.107 (0.093)	4.130 (2.803)	0.279*** (0.076)
At least bachelor	-0.258 (0.580)	0.156 (4.319)	-1.612 (1.874)	-9.319 (6.354)	-4.785 (5.214)	-0.169 (0.241)	-0.711 (7.900)	0.194 (0.182)
Republican	-0.194 (0.151)	6.885*** (1.430)	-0.271 (0.665)	-4.814*** (1.733)	7.546*** (2.113)	-0.249** (0.108)	-4.913* (2.780)	0.325*** (0.072)
Stock investor	-0.494* (0.279)	-1.219 (2.512)	2.124 (1.341)	2.196 (3.233)	-1.204 (3.619)	0.613*** (0.210)	-0.679 (5.395)	0.165 (0.144)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R-squared	.045	.049	.049	.017	.055	.234	.075	.091
Observations	1,014	1,014	1,014	1,014	1,014	941	1,014	1,014

Notes: This table shows OLS estimates of the determinants of respondents' expectations about the stock market and their own wealth. The outcomes are the expected duration of the recovery from the current crash in years (column 1); mean and standard deviation as well as probabilities assigned to extreme return realizations based on the respondent's reported probability distribution over the one year-ahead stock market return (columns 2-5); the expected recovery duration of the respondent's household financial wealth (column 6); a dummy indicating whether the respondent thinks her household net wealth will never recover, coded as 0 or 100 (column 7); and a categorical measure of the respondent's subjective household financial prospects, z-scored using the mean and standard deviation in the sample (column 8). All specifications are based on the pure control group, which has not received any information and not answered any questions on past crashes. All specifications control for shocks to income and financial wealth, trimmed at the 2nd and 98th percentiles, dummies for having lost any wealth or substantial wealth during past stock market crashes, gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, risky portfolio share, investment experience, Census region, and survey date. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

Covid Economics 23, 28 May 2020: 188-262

Table A7: Correlational evidence on beliefs about past crashes

	Stock recovery duration	Agree: Recovery many yrs.	Agree: Set back many yrs.	Agree: Recover 2020	Stock return: Mean	Stock return: SD	Stock return <-30%	Stock return >30%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Fin Crisis								
Prior recovery duration Fin. Crisis 2007	0.245*** (0.029)	0.049*** (0.009)	0.048*** (0.009)	-0.017* (0.010)	-0.153 (0.169)	0.101 (0.094)	0.153 (0.231)	-0.159 (0.252)
Observations	1,050	1,050	1,050	1,050	1,050	1,050	1,050	1,050
Panel B: Dot-com								
Prior recovery duration Dot-com 2000	0.225*** (0.025)	0.032*** (0.008)	0.031*** (0.008)	-0.010 (0.008)	0.159 (0.153)	0.135* (0.076)	0.010 (0.196)	0.355 (0.228)
Observations	1,063	1,063	1,063	1,063	1,063	1,063	1,063	1,063
Panel C: Black Monday								
Prior recovery duration Black Monday 1987	0.313*** (0.026)	0.058*** (0.007)	0.051*** (0.008)	-0.034*** (0.008)	-0.409*** (0.153)	0.229*** (0.073)	0.541*** (0.201)	-0.217 (0.228)
Observations	1,058	1,058	1,058	1,058	1,058	1,058	1,058	1,058
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS estimates of the effect of prior beliefs about the duration of a past stock market crash on the respondent's expectations about the stock market. The outcomes are the expected duration of the recovery from the current crash in years (column 1); agreement on 7-point scales to statements describing the severity of the current stock market crash, z-scored using the mean and the standard deviation in the sample (columns 2-4); mean and standard deviation as well as probabilities assigned to extreme return realizations based on the respondent's reported probability distribution over the one year-ahead stock market return (columns 5-8). Panel A is based on the control arm including questions on the Financial Crisis 2007. Panel B is based on the control arm including questions on the burst of the Dot-com bubble 2000. Panel C is based on the control arm including questions on the Black Monday 1987. All specifications control for gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the equity share in total financial assets, investment experience, Census region and survey date. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

Table A8: Effects of information on own outlook and plans: Experimental reduced form

	Wealth recovery duration	Household financial prospects	Plan incr. stock share	Plan decr. stock share	Exp. spend. growth	Incr. exp. debt	Incr. exp. desired hours	Incr. exp. retirement age
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: All								
Info Fin. Crisis 2007	0.484*** (0.078)	-0.100** (0.043)	0.952 (1.931)	3.054* (1.818)	-0.951 (0.914)	0.335 (1.924)	6.429** (2.632)	5.186* (2.662)
Panel B: Underestimators								
Info Fin. Crisis 2007	0.522*** (0.086)	-0.099* (0.052)	2.827 (2.333)	3.215 (2.211)	-1.478 (1.106)	2.978 (2.381)	7.401** (3.210)	8.229** (3.235)
Panel C: Under. & Stocks > 0								
Info Fin. Crisis 2007	0.761*** (0.130)	-0.136** (0.066)	2.827 (2.333)	3.215 (2.211)	-1.281 (1.328)	3.101 (3.002)	9.099** (3.773)	8.995** (3.917)
Panel D: All								
Info Dot-com 2000	0.573*** (0.087)	0.009 (0.043)	0.285 (2.015)	-2.085 (1.843)	-1.255 (0.928)	-1.477 (1.909)	-4.891* (2.623)	-3.710 (2.664)
Panel E: Underestimators								
Info Dot-com 2000	0.683*** (0.109)	0.022 (0.054)	1.149 (2.601)	-2.967 (2.279)	-0.533 (1.141)	-1.019 (2.334)	-4.145 (3.251)	-3.243 (3.321)
Panel F: Under. & Stocks > 0								
Info Dot-com 2000	1.017*** (0.164)	0.051 (0.069)	1.149 (2.601)	-2.967 (2.279)	0.779 (1.374)	0.167 (2.898)	-1.328 (3.822)	-3.998 (3.974)
Panel G: All								
Info Black Monday 1987	-0.157** (0.074)	0.048 (0.042)	3.293* (1.984)	-2.228 (1.710)	-1.027 (0.924)	-5.598*** (1.877)	-0.566 (2.612)	-0.806 (2.611)
Panel H: Overestimators								
Info Black Monday 1987	-0.242*** (0.088)	0.088* (0.047)	4.647** (2.253)	-4.079** (1.981)	-0.066 (1.070)	-6.705*** (2.141)	-1.607 (2.953)	-4.698 (2.962)
Panel I: Over. & Stocks > 0								
Info Black Monday 1987	-0.385*** (0.146)	0.097 (0.062)	4.647** (2.253)	-4.079** (1.981)	1.296 (1.361)	-5.026* (2.709)	-1.170 (3.558)	-3.673 (3.609)
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS estimates of the effect of being shown information on the duration of a past stock market crash on the respondent's expectations about her own financial situation and behavior. The outcomes are the expected recovery duration of the respondent's household financial wealth (column 1); a categorical measure of the respondent's subjective household financial prospects, z-scored using the mean and standard deviation in the sample (column 2); dummies indicating plans to increase or decrease the risky share in overall financial assets in the weeks after the survey (columns 3-4, only for stockholders); expected growth of yearly household spending from 2019 to 2020 in percent, trimmed at the 2nd and 98th percentiles (column 5); dummies indicating whether the coronavirus crisis increases the respondent's expectations about outstanding household debt by the end of 2020 (column 6), expected desired working hours over the next years (column 7, only if in labor force) or expected retirement age (column 8, only if in labor force). All dummy outcomes are coded as 0 or 100. Panels A-C are based on the treatment and control arms including information or questions on the Financial Crisis 2007. Panels D-F are based on the treatment and control arms including information or questions on the burst of the Dot-com bubble 2000. Panels G-I are based on the treatment and control arms including information or questions on the Black Monday 1987. Panels A, D and G are based on the full sample in the corresponding arms. Panels B, E and H are based only on under-estimators (for Financial Crisis and Dot-com bubble) or over-estimators (for Black Monday) of the length of the recovery from the crash. Panels C, F and I are based only on under-estimators or over-estimators who participated in the stock market before the current crisis. All specifications control for the respondent's prior belief about the recovery duration following the corresponding crash as well as for gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the equity share in total financial assets, investment experience, Census region and survey date. The specifications on planned stock trading in columns 3 and 4 also control for realized trading since the onset of the crisis. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

Table A9: Effects of expected stock market recovery on own outlook and plans: Heterogeneity

	Wealth recovery duration	Household financial prospects	Plan incr. stock share	Plan decr. stock share	Exp. spend. growth	Incr. exp. debt	Incr. exp. desired hours	Incr. exp. retirement age
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Age ≤ median								
Expected stock recovery duration (years)	0.244*** (0.078)	-0.047*** (0.012)	-0.809* (0.420)	0.396 (0.415)	-0.317 (0.272)	-0.145 (0.490)	2.085*** (0.525)	1.922*** (0.535)
Panel B: Age > median								
Expected stock recovery duration (years)	0.995*** (0.082)	-0.128*** (0.011)	-0.190 (0.356)	0.750** (0.375)	-0.730*** (0.222)	1.236** (0.496)	2.179** (0.854)	2.136** (0.909)
p-value (A=B)	0.00	0.00	0.26	0.53	0.24	0.05	0.93	0.84
Panel C: Net income ≤ median								
Expected stock recovery duration (years)	0.529*** (0.099)	-0.085*** (0.012)	-0.739* (0.402)	0.506 (0.440)	-0.348 (0.292)	1.234** (0.534)	2.063*** (0.744)	1.407* (0.790)
Panel D: Net income > median								
Expected stock recovery duration (years)	0.698*** (0.070)	-0.090*** (0.012)	-0.554 (0.398)	0.580 (0.356)	-0.762*** (0.226)	0.158 (0.431)	2.184*** (0.556)	2.853*** (0.570)
p-value (C=D)	0.16	0.72	0.74	0.90	0.26	0.12	0.90	0.14
Panel E: Net wealth ≤ median								
Expected stock recovery duration (years)	0.739*** (0.066)	-0.102*** (0.010)	-0.540 (0.336)	0.669** (0.312)	-0.649*** (0.204)	0.695* (0.385)	2.421*** (0.553)	2.661*** (0.573)
Panel F: Net wealth > median								
Expected stock recovery duration (years)	0.348*** (0.116)	-0.057*** (0.015)	-0.709 (0.506)	0.168 (0.552)	-0.332 (0.343)	0.253 (0.657)	1.289 (0.791)	1.107 (0.788)
p-value (E=F)	0.00	0.02	0.78	0.43	0.43	0.56	0.24	0.11
Panel G: Female								
Expected stock recovery duration (years)	0.488*** (0.082)	-0.078*** (0.011)	-0.750** (0.355)	0.570 (0.380)	-0.586** (0.230)	0.339 (0.477)	2.755*** (0.689)	1.871*** (0.682)
Panel H: Male								
Expected stock recovery duration (years)	0.710*** (0.084)	-0.102*** (0.013)	-0.409 (0.456)	0.516 (0.420)	-0.575** (0.277)	0.669 (0.482)	1.508** (0.595)	2.211*** (0.631)
p-value (G=H)	0.06	0.16	0.55	0.92	0.98	0.63	0.17	0.71
Panel I: Below Bachelor								
Expected stock recovery duration (years)	0.507*** (0.085)	-0.083*** (0.011)	-0.596 (0.364)	0.673* (0.364)	-0.359 (0.241)	0.986** (0.461)	2.030*** (0.589)	2.093*** (0.620)
Panel J: At least Bachelor								
Expected stock recovery duration (years)	0.714*** (0.078)	-0.092*** (0.013)	-0.694 (0.436)	0.497 (0.409)	-0.730*** (0.256)	-0.034 (0.496)	2.277*** (0.683)	2.082*** (0.679)
p-value (I=J)	0.07	0.61	0.86	0.75	0.29	0.13	0.78	0.99
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows OLS estimates of the effects of respondents' expected stock market recovery duration on the respondent's expectations about her own situation and behavior for different subgroups. The outcomes are the expected recovery duration of the respondent's household net wealth (column 1); a categorical measure of the respondent's subjective household financial prospects, z-scored using the mean and standard deviation in the sample (column 2); dummies indicating plans to increase or decrease the risky share in overall financial assets in the weeks after the survey (columns 3-4, only for stockholders); expected growth of yearly household spending from 2019 to 2020 in percent, trimmed at the 2nd and 98th percentiles (column 5); dummies indicating whether the coronavirus crisis increases the respondent's expectations about outstanding household debt by the end of 2020 (column 6), expected desired working hours over the next years (column 7, only if in labor force) or expected retirement age (column 8, only if in labor force). All dummy outcomes are coded as 0 or 100. All estimations are based on the four control arms, which have not received any information, and are restricted to those who participated in the stock market before the current crisis. All specifications control for gender, age, employment status, being the main earner, being financial decision-maker, party affiliation, log net household income, logs of retirement wealth, of other financial wealth, of real estate wealth, and of debt, borrowing constraints, stock market participation, the equity share in total financial assets, investment experience, Census region and survey date. The specifications on planned stock trading in columns 3 and 4 also control for realized trading since the onset of the crisis. Robust standard errors are reported in parentheses. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

Table A10: Perceived purpose of the survey

	All	All, excl. pure ctrl	Info treat., all	Ctrl. all, ex pure ctrl.	Difference (4) –(3)
	(1)	(2)	(3)	(4)	(5)
Impact COVID-19 on HH finances	47.25	46.29	44.62	47.97	3.35** [2.68]
Expectations	10.23	10.54	11.39	9.68	-1.71* [-2.22]
Knowledge test	2.73	2.88	3.36	2.40	-0.96* [-2.29]
Comparison of fin. crises	2.32	2.56	3.01	2.11	-0.90* [-2.27]
Experiment	0.07	0.08	0.09	0.06	-0.03 [-0.44]
Do not know	3.32	3.40	3.83	2.96	-0.86 [-1.90]
Other	34.08	34.26	33.70	34.82	1.12 [0.94]
Observations	7,447	6,358	3,187	3,171	

Notes: This table shows the relative frequency (in %) of answers to the question “What do you think was the purpose of the survey?” Answers were given as free text entries. We manually categorize the answers into 8 categories based on meaningful keywords, (1) impact of the corona crisis on household finances, (2) household economic expectations, (3) knowledge test/education, (4) comparison of different financial crises, (6) some form of experiment, (7) do not know, and (8) other. Column 1 is based on the entire sample. Column 2 excludes respondents in the pure control group, who have not received any questions or information on previous crashes. Column 3 is based on respondents in the three information treatment arms *FinCrisisInfo*, *DotComInfo* and *BlackMondayInfo*. Column 4 is based on respondents in the three control treatment arms *FinCrisisCtrl*, *DotComCtrl* and *BlackMondayCtrl*, excluding the pure control group. Column 5 shows the differences across percentages in the treatment and control arms, excluding the pure control group. * denotes significance at the 10 pct., ** at the 5 pct., and *** at the 1 pct. level.

Lockdown strategies, mobility patterns and Covid-19¹

Nikos Askitas,² Konstantinos Tatsiramos³ and Bertrand Verheyden⁴

Date submitted: 25 May 2020; Date accepted: 26 May 2020

We develop a multiple-events model and exploit within and between country variation in the timing, type and level of intensity of various public policies to study their dynamic effects on the daily incidence of COVID-19 and on population mobility patterns across 135 countries. We remove concurrent policy bias by taking into account the contemporaneous presence of multiple interventions. The main result of the paper is that cancelling public events and imposing restrictions on private gatherings followed by school closures have quantitatively the most pronounced effects on reducing the daily incidence of COVID-19. They are followed by workplace as well as stay-at-home requirements, whose statistical significance and levels of effect are not as pronounced. Instead, we find no effects for international travel controls, public transport closures and restrictions on movements across cities and regions. We establish that these findings are mediated by their effect on population mobility patterns in a manner consistent with time-use and epidemiological factors.

- 1 We thank an anonymous reviewer, Jan van Ours and Adrian Nieto Castro for discussions. Missing citations and discussions of related research will be added in future drafts.
- 2 Coordinator of Data and Technology, IZA-Institute of Labor Economics.
- 3 Professor of Labor Economics, University of Luxembourg and Luxembourg Institute of Socio-Economic Research (LISER).
- 4 Senior Researcher, Luxembourg Institute of Socio-Economic Research (LISER).

Copyright: Nikos Askitas, Konstantinos Tatsiramos and Bertrand Verheyden

1 Introduction

In December 2019, the COVID-19 outbreak was registered in Wuhan China. The World Health Organization declared it a “Public Health Emergency of International Concern” on January 30, 2020 and escalated it to a pandemic on March 11, 2020. The disease has been recorded in over 200 countries and territories with several millions of confirmed cases and a case mortality rate of around seven percent.¹ In the early stages of the outbreak, attempts were made to trace every infection back to its origin. Tracing back to the “index” case on an international level soon became impossible and most countries responded by imposing restrictions on international travel. In the later stages of the epidemic, a number of non - pharmaceutical interventions (henceforth referred to as NPIs or public policies) were undertaken, which were of a domestic nature revolving around the idea of “social distancing”. The aim of these interventions was to slow down the pandemic by restricting mobility so that it does not overwhelm health system capacities.

This paper studies how lockdown policies affect the daily incidence of COVID-19 and population mobility patterns across 135 countries based on several data sources.² Understanding the effectiveness of these policies is important as policy makers and the society at large seek to achieve an optimal health outcome in the fight against the pandemic at the lowest economic cost.

We exploit between and within country variation in the type, timing, and level of intensity of lockdown policies in a multiple events study approach, which aims at disentangling the effect of each intervention on COVID-19 incidence and mobility patterns, while controlling for the presence of concurrent policy measures during the event window of the policy of interest,

¹COVID-19 Dashboard by the CSSE at Johns Hopkins University (JHU).

²The data sources include: i) Coded government response data obtained from Hale et al. [2020], ii) prevalence data from European Centre for Disease Prevention and Control (ECDC) and iii) population mobility data from Google Community Mobility Reports. The analysis includes the latest data up to this writing.

as well as for time fixed effects, day of week fixed effects, lagged COVID-19 prevalence, region fixed effects and time-varying country-specific characteristics.

The main contributions of the paper are the following. First, we develop a multiple events model which allows us to identify the dynamic effects of each intervention while taking into account the presence of other concurrent interventions at each event time. Accounting for confounding policies is important because it allows us to avoid attributing the effect of other interventions to the policy of interest, and in addition to establish that it is policies that affect mobility patterns and not that policies ex-post respond to changing mobility patterns in the population.

Second, we consider a wide range of interventions across 135 countries, which vary in their type, intensity, and timing. The policy responses in focus are i) international travel controls, ii) public transport closures, iii) cancelation of public events, iv) restrictions on private gatherings, v) school closures, vi) workplace closures, vii) stay-at-home requirements and viii) internal mobility restrictions (across cities and regions).

Third, we link policy interventions to mobility patterns by studying not only the impact of these policies on the incidence of COVID-19, but also on the time spent in a number of types of places such as i) retail and recreation, ii) grocery and pharmacy, iii) parks, iv) transit stations, v) the workplace and vi) residential areas. Each of these types of places is characterized by different epidemiological features and, therefore, has a different potential for viral transmission. The mobility data can then also be viewed as a measure of compliance to the policies introduced and a mediator between policies and the spread of the disease.

The main result of the paper is that cancelling public events and imposing restrictions on private gatherings followed by school closures have quantitatively the most pronounced effects. They are followed by workplace as well as stay-at-home requirements, whose statistical significance and levels of effect are not as pronounced. Instead, we find no effects for international travel controls, public transport closures and restrictions on movements across

cities and regions. We thus establish i) the order in which public policies help curb the pandemic and ii) that these effects are mediated by the way they change population mobility patterns in a manner consistent with time-use and epidemiological factors.

The rest of the paper is structured as follows. Section 2 contains a literature review, while Section 3 discusses the data and presents summary statistics about NPIs and mobility patterns. Section 4 describes the model and identification issues. The results are presented in Section 5, which contains a discussion linking the evidence on COVID-19 incidence with mobility patterns. Section 6 concludes by summarizing the findings and discussing future research.

2 Literature

Research on infectious diseases focuses on vaccinations and drugs but it also aims at curbing the spread of the diseases by understanding and predicting their spatiotemporal dynamics, especially in the event of a new virus outbreak. Recent epidemiological models have been enriched to incorporate the impact of NPIs on these dynamics, which are at the core of this paper. The canonical model used in epidemiology is the so-called SIR model (Kermack et al. [1927]). It provides a simple and relevant representation of the mechanics of virus propagation with three categories of individuals: i) the people who are susceptible to become infected (the S subpopulation), ii) the infected who can transmit the disease (the I) and iii) those who have recovered and cannot infect anymore (the R). A crucial concept in the SIR model is the R_0 , which is the average number of people that a sick person infects before she recovers. While the R_0 is often considered as a biological characteristic of the virus' transmissibility, it is also affected by environmental, behavioral, and social dimensions, including NPI's.

In its basic form, the SIR compartmental model assumes that the population of interest

is homogeneous in terms of exposure, immunity and chances of recovery. This assumption, however, is not realistic as in practice these factors have proven to be key in guiding policy interventions (Auchincloss et al. [2012]). Relaxing this assumption gave rise to extensions of this model which aim at capturing the multiple dimensions of heterogeneity by partitioning the population into groups based on age or location. Pushed to the extreme, such partitions lead to the individual-based models (Eubank et al. [2004]), which require data on the intensity of contacts between individuals of different age groups to calibrate person-to-person contact rates, for instance via social mixing matrices. Using this approach, Jarvis et al. [2020] find for the UK that lockdown policies reduced the average number of daily contacts by 73 percent, resulting in a drop of the R_0 from 2.6 to 0.62, while Singh and Adhikari [2020] show for India that lockdown policies are unlikely to be effective if applied for 3 weeks or less.

Beyond the partitioning of the population by age groups or communities, the compartmental model has been extended to take into account important specificities of the disease, such as the incubation period, the duration of the acquired immunity, or the challenges it presents given the current state of knowledge. In the context of the COVID-19 pandemic, several variants of the compartmental model have been used. The SEIR model takes into account the group of “exposed” individuals who can be asymptomatic carriers during the incubation period (Karin et al. [2020], Pang [2020]; Roy [2020], Lyra et al. [2020], Lai et al. [2020]). In the SIRS model, recovery only provides a short-lived immunity, so that the R group moves back to the S group after some time (Ng and Gui [2020]). The SIOR model considers that only a fraction of the infected group is detected, or “observed”, by healthcare services (Scala et al. [2020]). While these various models capture different important features of the COVID-19 pandemic and provide predictions on the impact of NPI’s on the spread of the disease, their results are usually simulation-based and rely on structural assumptions.

In the field of economics, recent contributions depart from simple versions of the SIR model and introduce confinement policies as well as economic concepts (e.g. incentives, eco-

conomic cost of lockdown, value of life). These papers highlight through calibrated simulations the tradeoffs between the mortality induced by the excess demand for healthcare services and the economic losses induced by confinement policies. Gonzalez-Eiras and Niepelt [2020] develop a SIR model in which policies take into account, among others, the rate of time preference, the learning of healthcare services and the severity of output losses. Garibaldi et al. [2020] depart from the observation that the SIR model treats transitions from S to I as exogenous. In other words, the SIR model does not take into account individuals' decision to reduce the intensity of their contacts and their exposure to the virus. The authors borrow concepts from the search and matching model (Pissarides [2000]) to introduce a contact function into the SIR model with forward-looking agents. They show that the decentralized epidemic equilibrium is likely to be suboptimal due to the presence of externalities: while individuals care about the private benefits of distancing, they neglect its social benefits and the fact that it reduces the risk of hospital congestion; on the other hand, from a dynamic perspective, they do not take into account the benefits of herd immunity. Greenstone and Nigam [2020] develop a method to quantify the economic benefits of social distancing measures in terms of lives saved. They find that 1.7 million lives could be saved by applying mild social distancing for 3 to 4 months, which they estimate to be worth 8 trillion dollars accruing for 90 percent to the population above 50 years of age. Barro [2020] studies the impact of NPI's in the US during the Great Influenza Pandemic at the end of 1918 finding that even though NPIs reduced deaths peaks, and thereby reduced the stress imposed on healthcare services, they failed to significantly decrease overall mortality, which is likely due to the short application of the NPI's, with an average duration of around one month.

The papers related to this study are Chen and Qiu [2020], Gao et al. [2020], Engle et al. [2020] and Huber and Langen [2020]. Chen and Qiu [2020] focus on the reproduction number, which depends on the timing of NPI's with a parametric time lag effect, and predict for 9 countries the transmission dynamics under various sets of NPI's showing that country

differences lead to different optimal policies with heterogeneous tradeoffs between health and economic costs. By combining geographic information systems and daily mobility patterns in US counties, derived from smartphone location big data, Gao et al. [2020] show that in many counties in which mobility restrictions were only recommended but not imposed, mobility did not decrease. Engle et al. [2020] using US county-level GPS and COVID-19 cases, study the impact of local disease prevalence and confinement orders on mobility solving a utility maximization problem after splitting the utility derived from traveling a unit of distance into costs independent from the epidemic and costs related to perceived risk of contracting the disease. They find substantial effects of local infection rates, while official confinement orders lead to a mobility reduction of less than 8 percent. Huber and Langen [2020] exploit regional variation in Germany and Switzerland to assess the impact of the timing of COVID-19 response measures finding that a relatively later exposure to the measures entails higher cumulative hospitalization and death rates.

We differ from these papers in the following ways: i) we develop a multiple events model exploiting the timing, type and level of intensity of several public policies with the advantage of flexibility in the non-parametric estimation of their dynamic impacts, taking into account the contemporaneous presence of multiple interventions; ii) we consider as outcomes both the COVID-19 cases and various mobility patterns, with the latter capturing how often and how long certain public places or one's residence is frequented; and iii) we analyze a panel dataset of 135 countries.

3 Data and Descriptives

The analysis combines information from multiple data sources on (i) the non-pharmaceutical interventions implemented by governments, (ii) the daily number of infections, (iii) the evolution of population's mobility patterns, and (iv) various country characteristics.

Non-pharmaceutical interventions are collected by the Oxford COVID-19 Government Response Tracker (henceforth OxCGRT) for most countries of the world. The OxCGRT gathers publicly available information on several indicators of public policies aiming at mitigating the propagation of the virus. We focus on the following interventions: i) international travel controls, ii) closure of public transport, iii) cancellation of public events, iv) restrictions on private gatherings, v) closure of schools, vi) closure of workplaces, viii) restrictions on internal movement and viii) stay-at-home requirements. For each of these policies, we exploit information on the dates of introduction as well as qualitative time-varying information on their intensity. Intensity is measured in a scale from 1 to 6, which reflects whether the intervention is (i) recommended, (ii) mandatory with some flexibility, and (iii) mandatory with no flexibility, and whether it is geographically targeted or applied to the entire country. Recommended policies which are targeted obtain a value of 1, while mandatory policies with no flexibility applied to the entire country obtain a value of 6, with values in between referring to combinations of the policy stringency and its geographic scope.³ We use a sample of 135 countries for the estimation of NPIs, which is the set of countries for which we also have information on country characteristics (for a complete list see Appendix B).

Figure 1 presents the distribution of the number of days it took for each policy to be introduced after the first COVID-19 case averaged across countries. The distribution for the international travel controls is bimodal with the first mode well ahead of the first case. All policies have a main mode close to zero, with cancellation of public events and school closures enacted earlier, followed by restrictions on private gatherings and workplace closures, stay-at-home requirements, internal mobility restrictions and public transportation restrictions,

³To fix ideas, when the schooling policy receives an intensity score of 4, it means that it was not made mandatory in all schools or in all education levels, but it was applied to the entire country. A score of 5 means that it was made mandatory to all schools and education levels, but only in some areas of the country. A score of 6 means that it is mandatory for all schools and areas of the country. The average intensity level across countries is 2.9 for international travel controls, 3.8 for public transport closures, 5.4 for school closures, 3.8 for workplace closures, 4.7 for cancelling public events, 4 for restrictions on private gatherings, 3.2 for stay-at-home requirements and 4.2 for internal movement restrictions.

and a secondary mode late into the epidemic.

The number of confirmed cases of COVID-19 infections is extracted from the ECDC, which examines reports from health authorities worldwide in a systematic way in order to produce the number of COVID-19 cases and deaths every day. This provides us with information on the number of new cases each day in each country. The observed variation in the incidence of COVID-19 cases may be influenced in part by variation in reporting. In order to remove such random variation from the data, we use a 3-day moving average of the confirmed new cases and the inverse hyperbolic sine transformation in order to include days with zero reported new cases.⁴

To study how mobility patterns have evolved worldwide, we resort to Google's Community Mobility Reports. The Google mobility data are created with "*aggregated, anonymized sets of data from users who have turned on the Location History setting*" on their phone and show how "*visits and length of stay*" at different types of places change compared to the median value, for the corresponding day of the week, during the 5-week period from January 3, 2020 to February 6, 2020.⁵ Google's ability to accurately locate phones and to correctly categorize places varies both across countries as well as within (urban vs. rural areas). These data contain information on various epidemiologically relevant categories of places such as: i) retail and recreation, ii) grocery and pharmacy, iii) parks, iv) workplaces, v) transit stations and vi) residential areas. Retail and recreation covers visits to restaurants, cafes, shopping centers, theme parks, museums, libraries, and movie theaters. Grocery and pharmacy covers grocery markets, food warehouses, farmers markets, specialty food shops, drug stores, and pharmacies. Parks encompass national parks, public beaches, marinas, dog parks, plazas,

⁴Using the 3-day moving average helps visualization without changing the main findings. We also considered a 7-day moving average, which similarly maintains the main findings while removing reporting idiosyncrasies but additionally flatten features of the data which might be of interest. We opted for the 3-day moving average as the middle ground. We also conduct our analysis without the 3-day moving average as discussed below.

⁵COVID-19 Community Mobility Reports.

and public gardens. Transit stations cover subway, bus and train stations. From the sample of 135 countries we have information on mobility patterns from Google for a subsample of 108 countries (for a complete list see Appendix B).

Figure 2 presents the distribution of these data averaged across countries both before and after the first confirmed COVID-19 case. Before the first case, all distributions are highly concentrated around zero, which suggests no substantial change of movement compared to the baseline period. After the first case, retail and recreation as well as transit stations have a mean of just above -40, suggesting a 40 percentage points drop, while differing in their variance. Grocery and pharmacy as well as workplaces have a mean of around -20 differing in their skewness (grocery and pharmacy is heavy on the left). Parks stand out for having a mean closest to zero and being heavy on the right. Finally, residential areas have a mean of just over 10.⁶ In terms of densities, retail and recreation, grocery and pharmacy, transit stations as well as workplaces are somewhat similar, while parks and residential areas are on their own on opposite sides.

⁶Considering that staying at home is by far the most time intensive activity, according to time-use studies, this value is quite large. We expand on this point in Section 5.3.

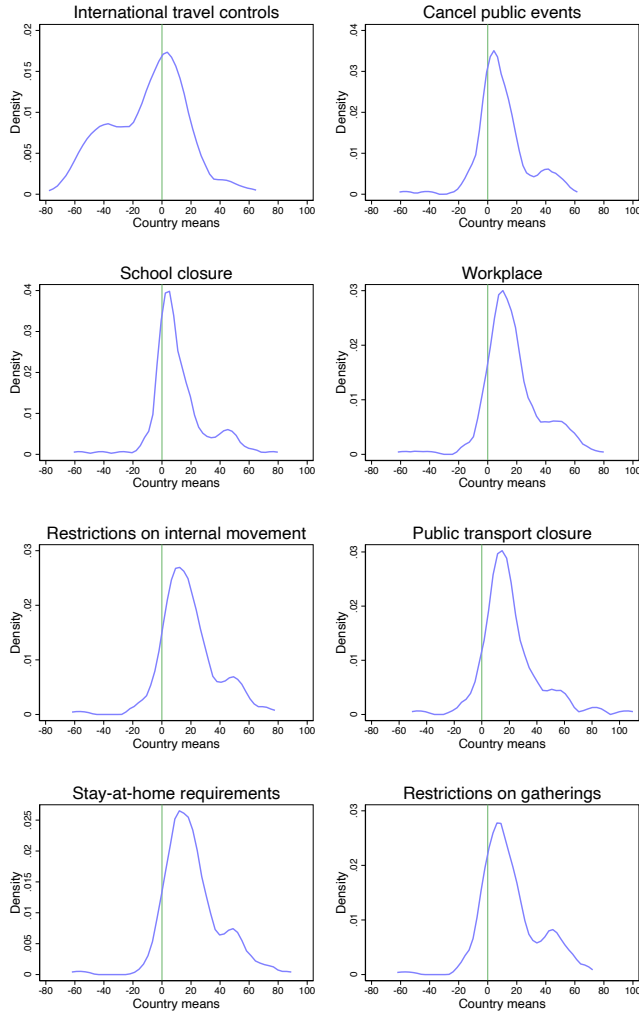


Figure 1: Lockdown Policies - Days after first COVID-19 case each policy was introduced.

Covid Economics 23, 28 May 2020: 263-302

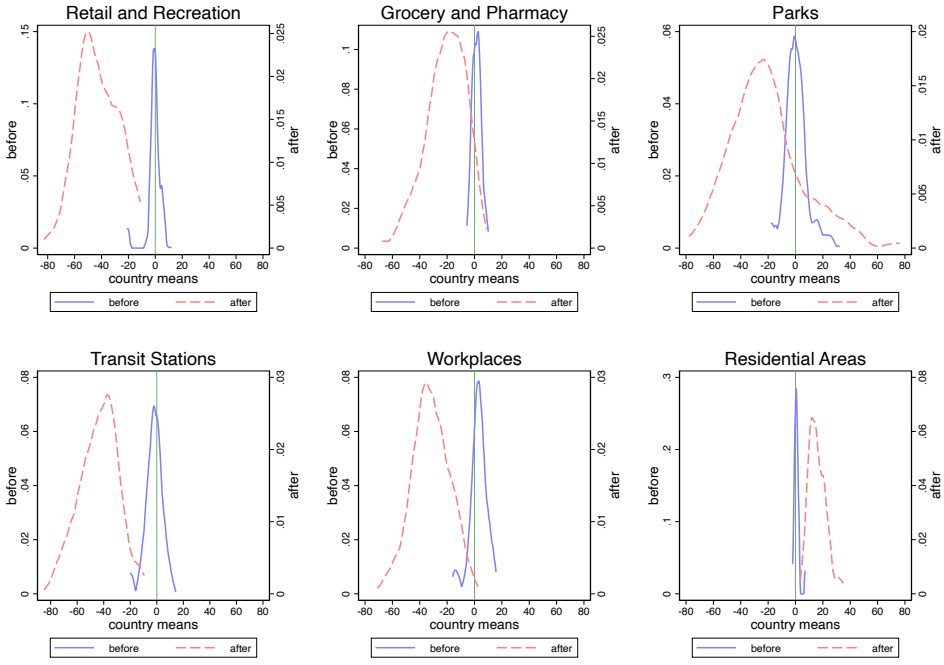


Figure 2: Google Mobility Patterns - Densities before and after first COVID-19 case.

Covid Economics 23, 28 May 2020: 263-302

4 The model

We follow an event-study approach around the time of policy implementation, which we extend to account for multiple events. The single event study (e.g. [Kleven et al., 2019]) can be expressed with the following equation:

$$Y_{c,t} = \sum_{j \neq -20} \alpha_j I[j = t - t_c^\pi] + \sum_l \gamma_l I[l = t] + \sum_d \delta_d I[d_c(t) = d] + \sum_r \rho_r R_r + \phi Z_{c,t-1} + \theta X_c + \epsilon_{c,t}, \quad (1)$$

where $Y_{c,t}$ denotes the outcome in country c at event time t . The first term, on the right hand side, is a set of event time dummies for the intervention of interest π , where t_c^π denotes its implementation day in country c . We consider the outcome in the window starting 20 days before the intervention up to 35 days after its implementation, so the event time runs from -20 to $+35$. We omit the event time dummy at $j = -20$ so that the event time coefficients of interest α_j measure the impact of intervention π at time j relative to the twenty days before the policy was implemented.

The second term, on the right hand side of equation (1), is a set of dummies which control non-parametrically for trends in the time since the first-observed COVID-19 case. Identifying the coefficients of the event time dummies conditional on these time effects is possible because the timing of NPIs differs across countries. The third term, is a set of day-of-week dummies controlling for potential day-specific differences both in terms of reporting of new cases and of mobility patterns (here $d_c(t)$ returns the day of the week for event time t in country c). The fourth term, is a set of dummies for the following regions: Europe, Asia, Middle East, North America, South America, Oceania and Africa. The fifth term, is the log value of the total number of confirmed cases at $t - 1$ in country c . Including this variable allows us to capture the size of the pool of infected people, which is a crucial factor both when

the outcome is the incidence of new cases, in line with the SIR framework, as well as when it is mobility patterns, as populations may react to the perceived threat of contamination.⁷ The sixth term, is a set of country specific variables controlling for differences across countries, such as per capita GDP, population density and the urbanization rate, followed by the error term.⁸

When evaluating the effect of the intervention of interest π , it is important to take into account the presence of other contemporaneous interventions, which can have their own contribution in affecting the outcome, and thus, if ignored can lead to biased estimates. However, identifying the effect of the policy of interest π with multiple events is more challenging than in the single-event case, especially when the multiple events fully overlap during the event window of the policy of interest. Concurrent NPIs, denoted by π' , can be controlled for by introducing in equation (1) a new term, $F^{\pi'}[j = t - t_c^\pi]$, which is a set of dummies - one dummy for each event time of the policy of interest π - which are equal to one if any other interventions π' are in effect at event time j for country c . The multiple events regression equation can then be written as follows:

$$Y_{c,t} = \sum_{j \neq -20} \alpha_j I[j = t - t_c^\pi] + \sum_j \beta_j F^{\pi'}[j = t - t_c^\pi] + \sum_l \gamma_l I[l = t] + \sum_d \delta_d I[d_c(t) = d] + \sum_r \rho_r R_r + \phi Z_{c,t-1} + \theta X_c + \epsilon_{c,t}. \quad (2)$$

The identification problem in the multiple-event case emerges as soon as other policies have been introduced before the start of event window of policy π . This would result in a complete overlap of policies within the event window, making it impossible to separately identify the effect of the event of interest from the other contemporaneous events.⁹

⁷Adjusting the total number of confirmed cases by the number of deaths does not affect our main results.

⁸When the outcome is the incidence of COVID-19, these controls are epidemiologically relevant, whereas, when we consider Google mobility types as our outcome they help control for differences in Google's ability to geo-locate phones and detect types of places.

⁹When other policies are enacted within the event window, then the two set of event dummies are not

To achieve identification in the multiple events case, we use the level of intensity of each policy which varies both within and across policies, as well as across countries. This variation of policy intensity allows us to identify separately the effect of the policy of interest π , while taking into account other concurrent NPIs, π' . The extended multiple events regression equation can be written as follows:

$$Y_{c,t} = \sum_{j \neq -20} \alpha_j S^\pi [j = t - t_c^\pi] + \sum_j \beta_j \bar{S}^{\pi'} [j = t - t_c^\pi] + \sum_l \gamma_l I[l = t] + \sum_d \delta_d I[d_c(t) = d] + \sum_r \rho_r R_r + \phi Z_{c,t-1} + \theta X_c + \epsilon_{c,t}, \quad (3)$$

where the first term, $S^\pi [j = t - t_c^\pi]$, is taking the value of the level of intensity of the policy of interest π in country c at event time j and zero otherwise, while the second term, $\bar{S}^{\pi'} [j = t - t_c^\pi]$, is equal to the average level of intensity of all other contemporaneous policies π' of country c at the event time j , and zero if there are no other policies active at that event time. That is, we extend equation (2) in two ways: 1) we multiply the event dummies for policy π with the intensity level of the policy at event time j - in other words, $I[j = t - t_c^\pi]$ in equation (2) generalizes to $S^\pi [j = t - t_c^\pi]$ in equation (3); and 2) we multiply the dummies controlling for the presence of any other policies π' - at event time j for policy π - with their average intensity at event time j - in other words, $F^{\pi'}$ in equation (2) generalizes to $\bar{S}_{c,t}^{\pi'}$ in equation (3).

Our identification relies on the variation in the timing and intensity of various interventions both within and across countries, conditional on the prevalence of COVID-19, time effects since the first observed case, day effects, country-specific characteristics and continent effects. This variation allows to separately identify the effect of intervention π from that of other concurrent NPIs, π' . The coefficient estimates α_j in equation 3 measure the unit level perfectly collinear so the coefficient estimates α_j and β_j can be separately identified, but at the cost of high variance because of multicollinearity.

intensity effect of policy π at event time j on the outcome.

5 Results

This section contains the results split in three subsections. Subsection 5.1 contains results on the effect of NPIs on the incidence of COVID-19, whereas Subsection 5.2 contains results on the effect of NPIs on population mobility patterns from Google's Community Mobility Reports. Finally, Subsection 5.3 provides a consolidated view on the link between the impact of NPIs on new cases through their effect on mobility.

5.1 Lockdown Policies on COVID-19 Incidence

We start by comparing the estimates for the dynamic effects of each intervention obtained from the two versions of the model: i) ignoring concurrent interventions, i.e. estimating equation (3) without the second term, and ii) controlling for concurrent interventions, i.e. estimating equation (3).¹⁰

Comparing the two sets of estimates, it becomes apparent that ignoring the presence of other interventions leads to biased estimates. Specifically, the results of Figure 3, which report the estimates ignoring concurrent policies, suggest that all policies tend to have a significant impact following a similar pattern. That is, in the days preceding the introduction of the policy, the incidence of COVID-19 increases until it reaches a peak after few a days following its introduction. Then, the number of new cases per day start to decrease, and within a month they become significantly lower than the reference event time (20 days before the policy).

The analysis without controlling for other concurrent policies seems to suggest that all interventions were successful in containing new infections. However, the estimates in Figure

¹⁰We focus on the results where the dependent variable is the 3-day moving average of confirmed new cases. The estimates with the number of confirmed cases are reported in Figures A1 and A2 in Appendix A.

4, which are obtained after controlling for concurrent policies, convey a different message. This is especially true for the two transport related interventions, i.e. *international travel controls* and *public transport closure*, and for *restrictions on internal movement*, which have almost no impact on new cases once other interventions are controlled for.

The two policies with the largest effects, which are robust to confounding by other policies, are *cancelling of public events* and *restrictions on private gatherings*. These are policies which aim to reduce massive contacts.¹¹ For both, we observe a drop in the incidence of COVID-19 starting about one week after implementation, which becomes significantly different than zero within two weeks. Around the end of the event window, a unit increase in the intensity of the policy of interest leads to a 20% decrease in the number of new infections in the case of public events cancelation, and a decrease of about 12% in the case of restrictions on private gatherings, compared to the reference event time.

School and *workplace* closures aim to control contacts in large groups, but unlike public events and private gatherings are easier to monitor and regulate as well as trace whenever infections do occur. We find that new infections start declining a few days after school closures, with the effect becoming negative and significant about 25 days after implementation. Around the end of the event window, a unit increase in the intensity of school closures leads to about a 15% drop of new infections compared to the reference event time. For workplace closures, we find that new infections start declining starting from the second week after implementation and the effect becomes negative and significant only towards the end of the event window, with a unit increase in the policy intensity leading to about a 10% drop of new infections.

Finally, *stay-at-home requirements* aim to impose mobility constraints at the individual level, which is arguably the most extreme of all measures and was generally introduced when infections were reaching alarming growth rates. This is captured in Figure 4 which shows

¹¹See the discussion in Section 5.3.

that, around the date of introduction of the policy, there were on average 20% more new cases every day compared to the reference event time, with the policy reversing that trend immediately. By the end of the window, the coefficient estimates are statistically significantly lower from those around the policy implementation day, although they are not statistically different from zero.

5.2 Lockdown Policies on Google Mobility Patterns

Google mobility patterns are observed as percentage point deviations from a reference calendar period before the onset of COVID-19. As a result, the coefficient estimates of interest - first term of equation (3) - measure the percentage point change in mobility patterns for a unit level of intensity of each intervention compared to the reference point before implementation. Similar to COVID-19 confirmed new cases, we obtain estimates both with, as well as without controls for other ongoing interventions. The estimates with controls for concurrent policies are presented in Figures 5, 6, 7 and 8, while those without controlling for confounding policies can be found in Figures A3, A4, A5 and A6 in Appendix A.

We find that, when we do not control for confounding policies, right after the day of policy implementation there is a general pattern of sharp and large drops in all mobility patterns related to activities undertaken outside residential areas, and an increase in the amount of time spent in the place of residence. However, once we control for other concurrent NPIs, many of these effects are either much smaller, or sometimes not significantly different from zero. For example, the estimates for *international travel controls* without accounting for confounders, shown in panel (a) of Figure A3, suggest a significant decline in movements immediately after the policy implementation across most places (retail and recreation, grocery and pharmacy, parks, transit stations, workplaces) and increases in staying home. However, after controlling for other interventions present around the same time, we find in panel (a) of Figure 5 that restrictions in international travel have a much smaller impact on all types

of movement.

After accounting for multiple events, panel (b) of Figure 5 shows that restrictions on *public transportation* lead to a sharp discontinuity at the day of the intervention with lower movements outside home. Interestingly, the strength of this decrease in mobility tends to weaken with time. The limited - and not very persistent - reductions in mobility patterns observed for international travel controls and closure of public transportation are consistent with the small effects of these policies on the incidence of new cases reported in Section 5.1.

The *cancelation of public events* and *restrictions on private gatherings*, which led to the most important reductions in new infections, as reported in Section 5.1, also exhibit large and persistent negative impacts on retail and recreation, transit stations, workplaces and to a lesser extent grocery and pharmacy (panels (a) and (b) of Figure 6). For both policies, the magnitude of these drops is around 5 percentage points per unit level of policy intensity. These findings are consistent with the fact that attending public events and private gatherings generate spillover effects on various activities outside the homeplace. Conversely, these policies have significantly increased time spent at home.

As reported in Subsection 5.1, the set of interventions with the second strongest reductions on subsequent infections were *school* and *workplace closures*. Figure 7 shows that these policies do change the mobility trends associated with crowded places, such as retail and recreation, transit stations and workplaces. Again, this can be explained by the fact that closing schools and workplaces generate spillover effects on other activities. The sharpness and the magnitude of the mobility decreases is much stronger in workplace than in school closures, with a unit level increase in the intensity of the policy leading to a stable decline of up to about 7-8 and 2-3 percentage points, respectively. This difference is also consistent with the following observations. First, workers generally have access to more mobility patterns and activities than pupils. Second, while pupils staying at home might constrain the mobility of one parent, closing workplaces affects the mobility of all adults working in the

household.

Stay-at-home requirements (panel (a) of Figure 8) result in large drops in all population mobility patterns at the time when they were introduced, a fact which is consistent with the reversal of the increasing trend of new cases reported in Subsection 5.1. Finally, in line with the results obtained for infection cases, *internal mobility restrictions* (panel (b) of Figure 8) have a similar impact on all mobility patterns as home confinement, though neither as sharp, nor as strong. The magnitude of these effects are about half the size compared to stay-at-home requirements.

We conclude this section with three remarks. First, mobility patterns do not exhibit much in the way of anticipation effects once we control for confounding NPIs. This suggests that it is policies affecting mobility patterns and not that policies ex-post responding to de facto changing mobility patterns in the population. It is worth noting that estimates ignoring concurrent NPIs would have led to a completely different conclusion; as shown in Figures A3 to A6 in Appendix A, for several interventions mobility patterns seem to respond before policies are in place. Second, we observe a spike in movements to groceries and pharmacies prior to the introductions of several NPIs, such as public transport and workplace closure, as well as stay-at-home and internal movement restrictions. This is consistent with the widely reported runs on the shelves in anticipation of lockdowns, concerns about imminent shortages, as well as with inadvertent signaling from these interventions about the threat level of the pandemic. Again, we are able to detect these mobility patterns only when we account for confounding policies. In light of the fact that we control for the state of the epidemic by using lags of total confirmed cases, this result is robust and shows the strength of our model. Third, it appears that the decline in mobility patterns is stable towards the last days of analysis, suggesting that compliance does not decline over time, at least within the 35-day window of our study.

5.3 Lockdown policies: a consolidated view

In this section, we expand on how the observed variation in the effects of NPIs on the incidence of COVID-19, reported in Section 5.1, can be understood by the way in which they affect various mobility patterns across places, reported in Section 5.2, which differ in a number of characteristics as they pertain to epidemiology, as well as in their time-use intensity. We thus provide a consistent framework for our results.

First, the degree to which restricting mobility to different places is expected to affect new infections depends on several characteristics of these places, where the most important are the following: i) numerosity, ii) density, iii) social norms, iv) geographical range and v) tracking ability. For example, more numerous and dense places, such as large private gatherings and public events, are more likely to contribute to new infections because the two-meter safe social-distancing rule is more likely to be violated there than say in parks. However, places with similar density can be conducive to different behavior types due to social norms; for example, in a soccer game, where there are large numbers of people densely brought together, there are different norms of accepted behavior compared to the regulated environment of a workplace. Furthermore, places such as schools vs. transit stations, or public events, can have different epidemiological range. For example, an infection at school has a range of perhaps a couple of kilometers (students reside close to their schools), while in the case of a soccer game it might be several kilometers and even cross country borders. Finally, places differ in terms of how easy it is to trace an infection back whenever it occurs, which is important because tracking contains the spread of the virus. For example, an incident at a workplace can be announced immediately to employees and an ad hoc lockdown can be probably enforced at the same time, while an infection which occurs at a transit station is impossible to trace back or treat with a local lockdown.

Second, places differ from a time-use perspective. Based on time-use surveys on how people spend their time in everyday life, for example, European adults in selected countries

between the ages of 20 and 74 years old, spend on average on a daily basis: i) 15 hours at home preparing meals, sleeping, and on household activities, ii) a little less than 3 hours at work, which has mostly a large workplace component, iii) a little more than 1 hour traveling and commuting and iv) about 4 and a half hours on other activities including leisure (recreation, parks, home) and shopping (retail, groceries and pharmacies).¹² These differences in time-use suggest, for example, that observing a 3 percentage points increase in time spent at the place of residence implies an increase of about half an hour, whereas a decrease of 8 percentage points in workplaces amounts to a drop of about 15 minutes.

In light of these differences across places, we find that NPIs tend to reduce activities away from home, while increasing time spent at home to a varying degree depending on their time-use footprint, while their impact on the epidemic depends on the above mentioned epidemiologically relevant characteristics.

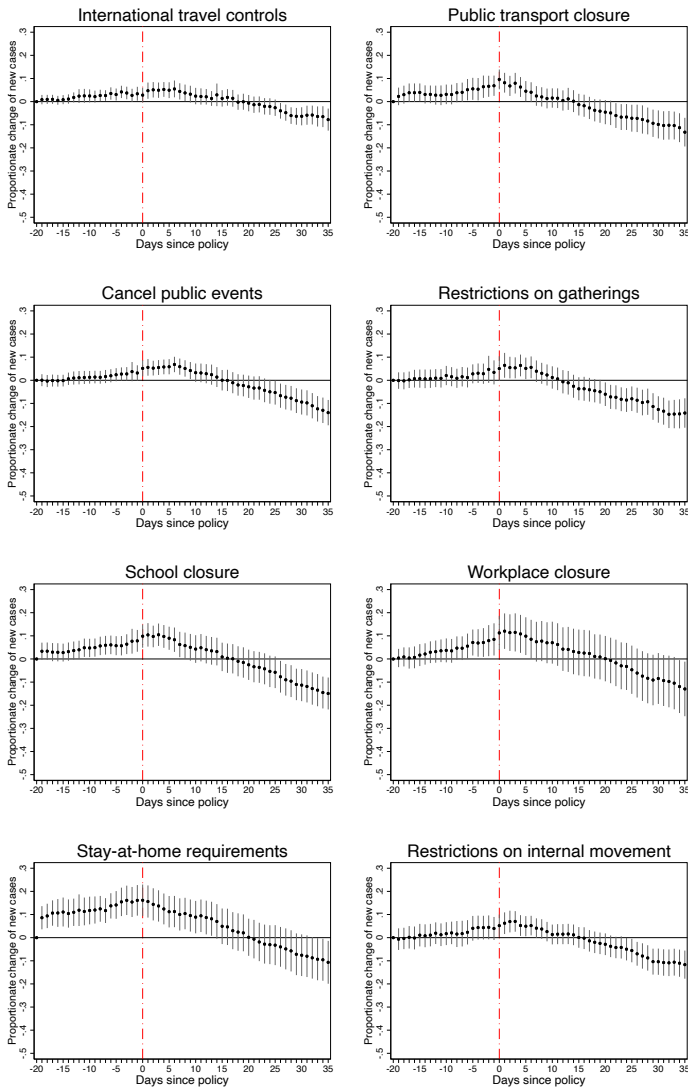
More specifically, *cancellation of public events*, and to a lesser extent *restrictions on private gatherings*, which are seen to lead to a large reduction in new infections, are interventions that reduce exposure to numerous and dense locations, where contact tracing is difficult, and can have a large epidemiological range within and across countries. Similarly, *stay-at-home requirements*, *workplace* and *school closures* reduce activities away from home and lead to significant reductions in the incidence of new infections, which nevertheless are not as large as for public events and private gatherings, possibly because of the differences in numerosity, density and ability to trace new infections in these environments.

On the other hand, although *restrictions on internal movements* reduce mobility across cities and regions, they impact the spread of the disease in a less pronounced way. This is consistent with the fact that these restrictions are not clearly linked to places with high density, and their potential to slow down new infections by restricting geographical mobility is

¹²<https://ec.europa.eu/eurostat/documents/3930297/5953614/KS-58-04-998-EN.PDF/c789a2ce-ed5b-4a0c-bcbf-693e699db7d7?version=1.0>

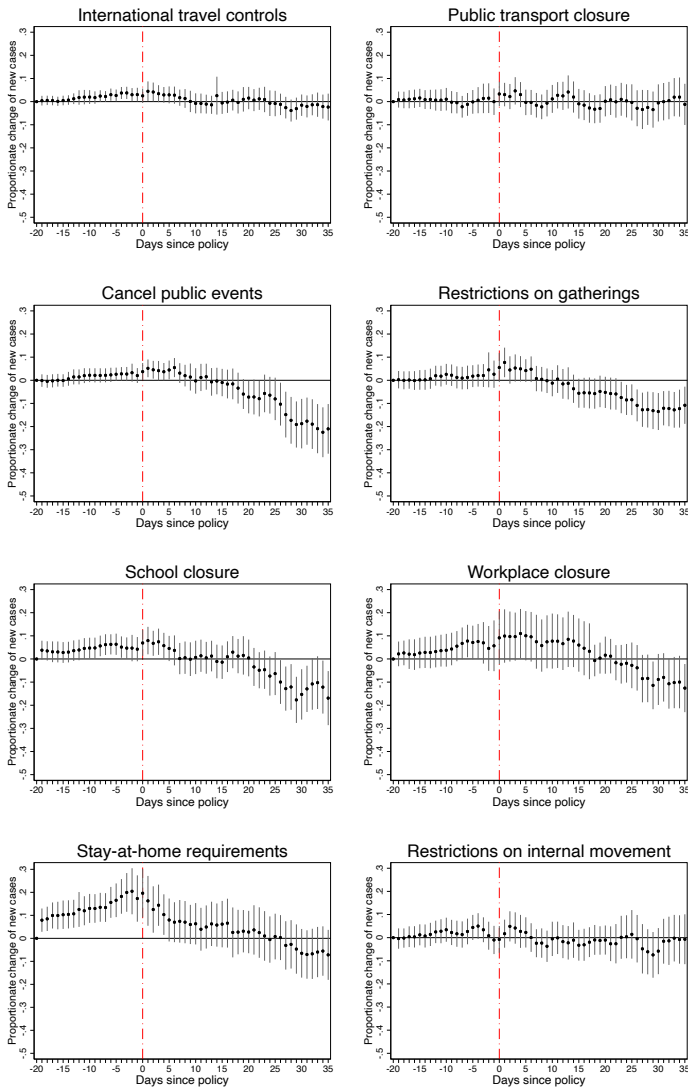
reduced, once other policies such as workplace closures and restrictions on private gatherings are in place. Furthermore, *public transport closures* were introduced on average at a time where demand for traveling and commuting has declined due to other restrictions in place such as workplace closures, which can explain both their limited impact on mobility and on reducing new infected cases.

Finally, the limited impact of *international travel controls*, although they were imposed relatively early by many countries, is likely explained by the lack of stringency of the controls. If countries have banned all international travel soon after the outbreak in China, it would have certainly be an effective measure to seal the country from the virus. However, because most countries did not introduce such bans before the virus has started spreading domestically, or they did introduce some restrictions but not complete bans, those restrictions had a limited impact on mobility and could only reduce new imported infections but not contain the spread of the virus.



Note: Data from Hale et al. (2020), European CDC and own calculations

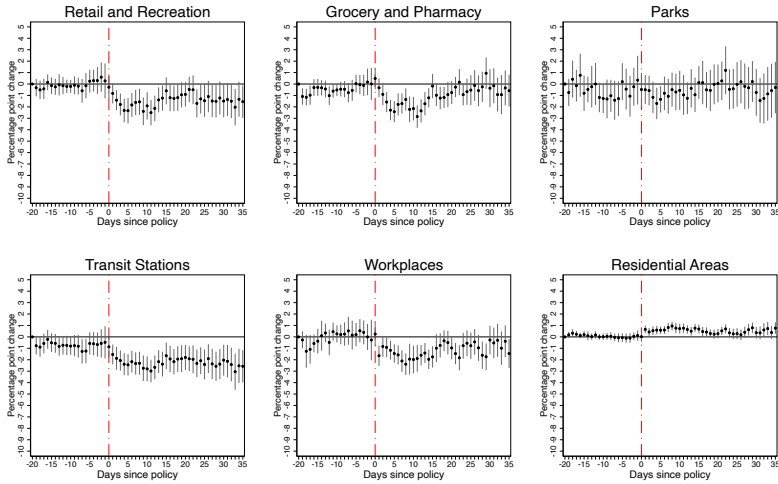
Figure 3: Effects of lockdown policies on COVID-19 confirmed new cases (3-day moving average, in logs) without concurrent policy controls.



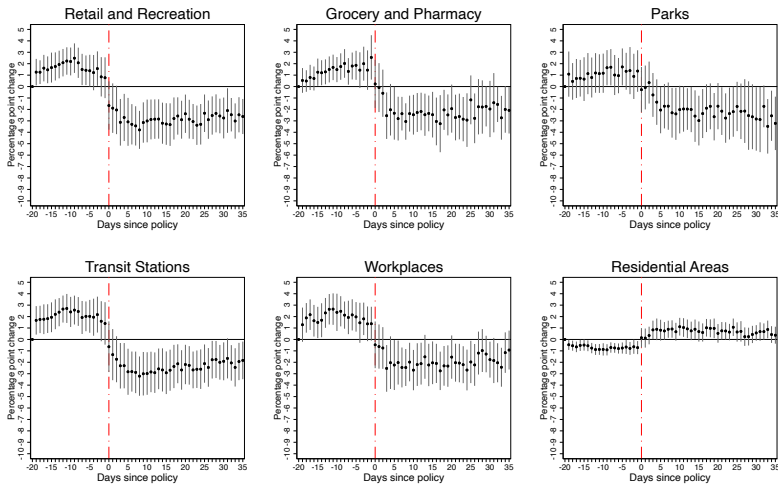
Note: Data from Hale et al. (2020), European CDC and own calculations

Figure 4: Effects of lockdown policies on COVID-19 confirmed new cases (3-day moving average, in logs) with controls for concurrent policies.

(a) International travel controls

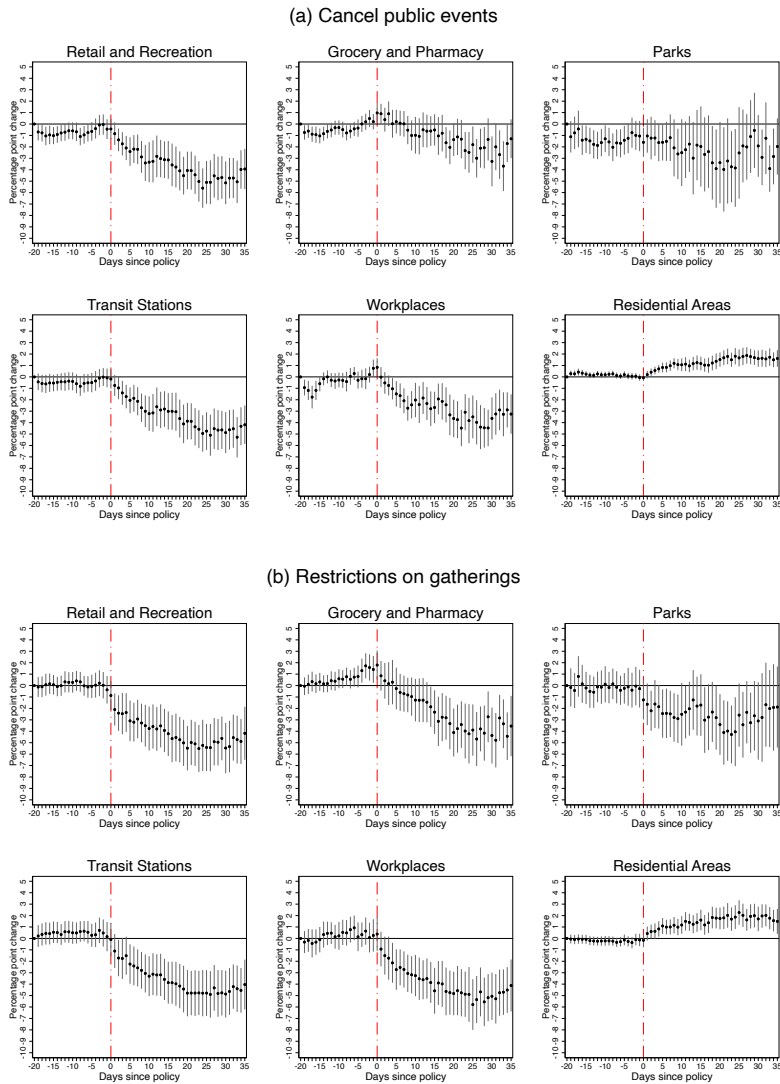


(b) Public transport closure



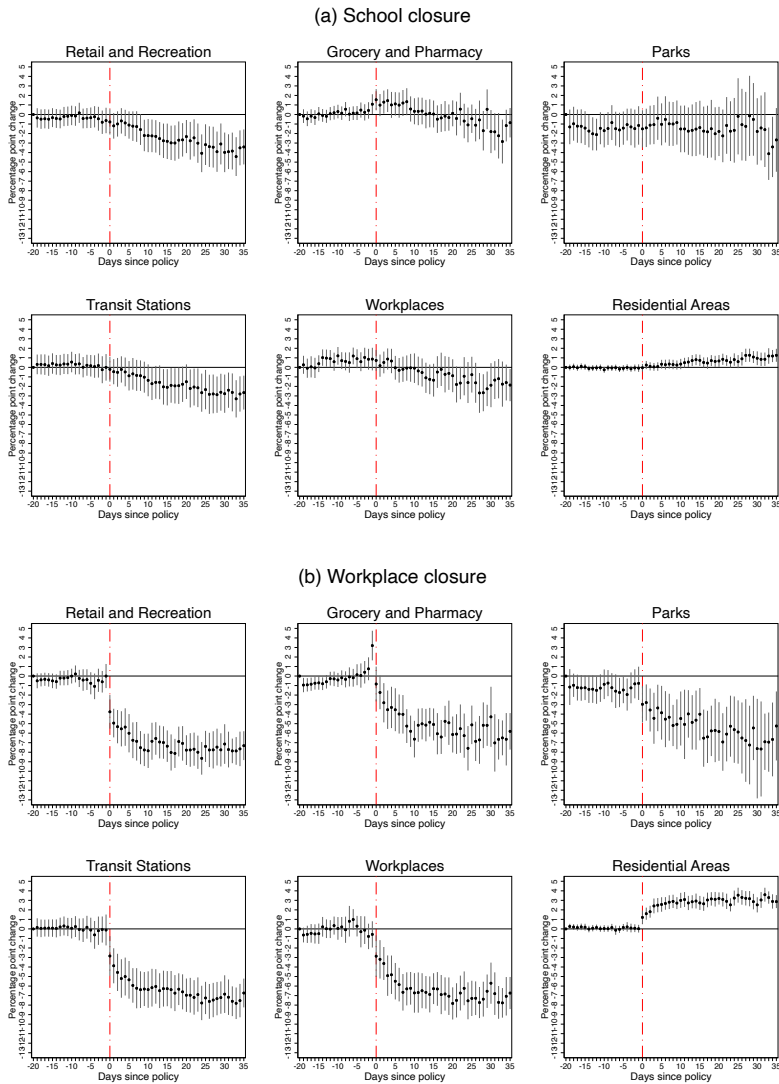
Note: Data from Hale et al. (2020), Google Community Mobility Reports and own calculations

Figure 5: Effects of **international travel controls** (panel a) and **closure of public transportation** (panel b) on Google mobility patterns.



Note: Data from Hale et al. (2020), Google Community Mobility Reports and own calculations

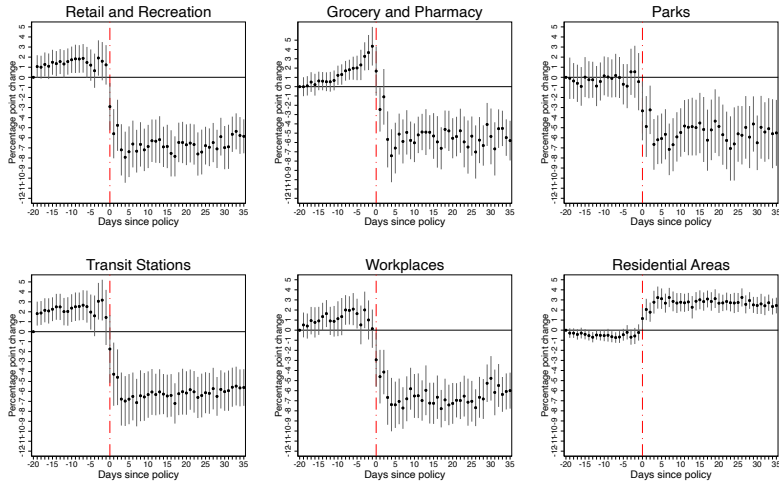
Figure 6: Effects of **public events cancellations** (panel a) and **restrictions on gatherings** (panel b) on Google mobility patterns.



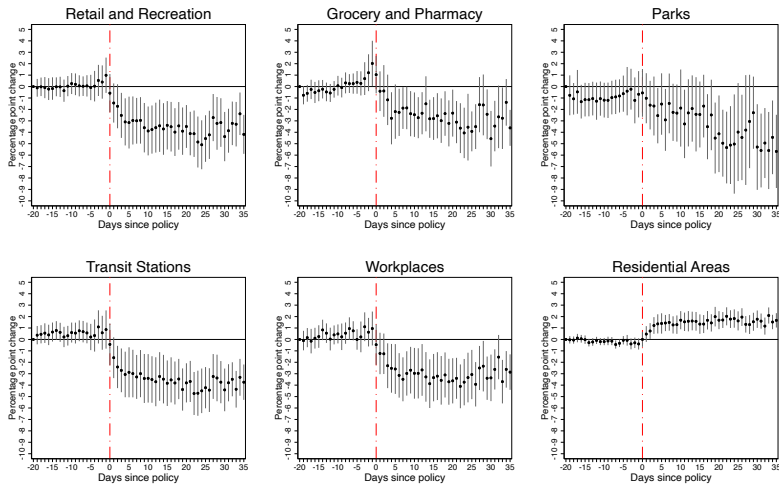
Note: Data from Hale et al. (2020), Google Community Mobility Reports and own calculations

Figure 7: Effects of **school** (panel a) and **workplace** (panel b) closures on Google mobility patterns.

(a) Stay-at-home requirements



(b) Restrictions on internal movement



Note: Data from Hale et al. (2020), Google Community Mobility Reports and own calculations

Figure 8: Effects of **stay-at-home requirements** (panel a) and **restrictions on internal mobility** (panel b) on Google mobility patterns.

6 Conclusion

The COVID-19 pandemic impacts societies and economies in multiple and dramatic ways. The exact extent of this impact in economic and social terms is certainly going to remain a topic of interest in the time ahead. In this paper, we develop a multiple events model to study the effect of lockdown policies on the incidence of new infections and on mobility patterns exploiting variation in the type, timing and intensity of confinement policies across 135 countries. The key contributions of the paper are twofold: i) we model the dynamic effects of each policy on the incidence of new infections accounting for concurrent policies, while in line with the standard SIR model, we specify future infections (incidence) as a function of past cases (prevalence), as well as a number of risk related characteristics, such as GDP per capita, population, population density and urbanization rates, all of which enrich the exposure to risk of infection with heterogeneity within and between countries and ii) we link the effect of NPIs on new infections through their impact on mobility patterns.

Our findings establish that cancelling public events and enforcing restrictions on private gatherings followed by school closures, which reduce mobility patterns in numerous and dense locations, each with their own particular behavioral norms, have the largest effect on curbing the pandemic in terms of statistical significance and levels of effect. They are followed by workplace and stay-at-home requirements, which also reduce activities away from home and lead to significant reductions in the incidence of COVID-19, which nevertheless are not as large as for public events, private gatherings and school closures, possibly because of the differences in numerosity, density and the ability to trace new infections in these environments. Instead, restrictions on internal movement, public transport closures and international travel controls do not lead to a significant reduction of new infections. The limited impact of travel controls, although imposed relatively early in many countries, is likely explained by their lack of stringency allowing the virus to cross borders.

Our econometric framework is suitable for the study of dynamic effects with multiple events, which can be applied in many settings. A natural one is the upcoming exit strategies from the lockdowns, which we will turn to next.

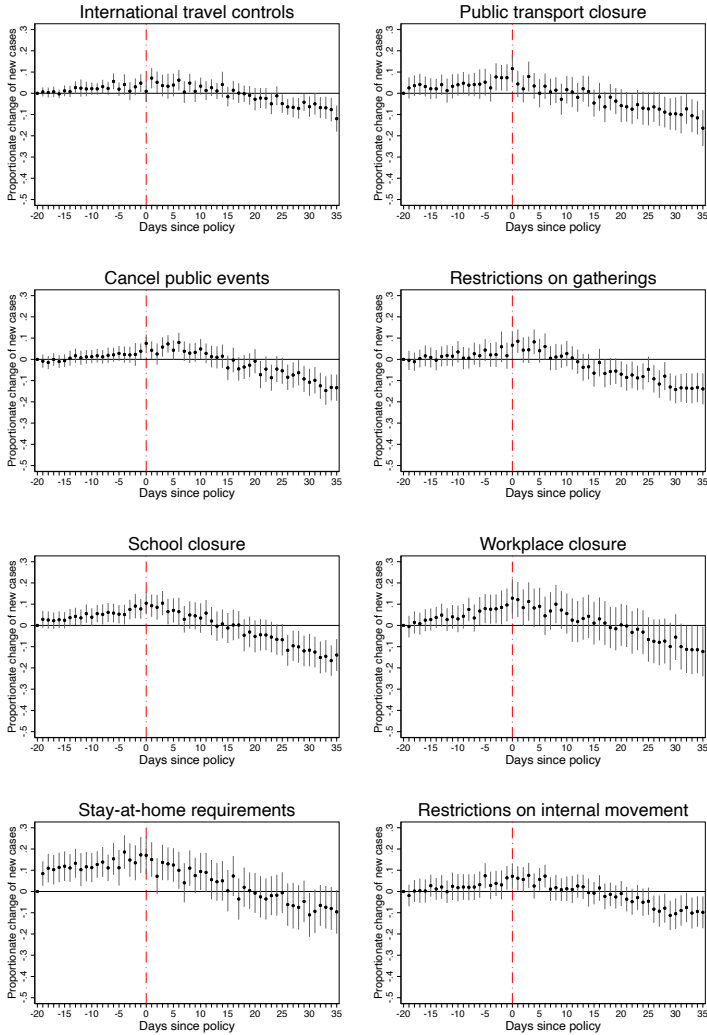
References

- Auchincloss, A. H., Gebreab, S. Y., Mair, C., and Diez Roux, A. V. (2012). A review of spatial methods in epidemiology, 2000–2010. *Annual review of public health*, 33:107–122.
- Barro, R. J. (2020). Non-pharmaceutical interventions and mortality in us cities during the great influenza pandemic, 1918-1919. Technical report, National Bureau of Economic Research.
- Chen, X. and Qiu, Z. (2020). Scenario analysis of non-pharmaceutical interventions on global covid-19 transmissions. *arXiv preprint arXiv:2004.04529*.
- Engle, S., Stromme, J., and Zhou, A. (2020). Staying at home: mobility effects of covid-19. *COVID Economics*, (2020-4).
- Eubank, S., Guclu, H., Kumar, V. A., Marathe, M. V., Srinivasan, A., Toroczka, Z., and Wang, N. (2004). Modelling disease outbreaks in realistic urban social networks. *Nature*, 429(6988):180–184.
- Gao, S., Rao, J., Kang, Y., Liang, Y., and Kruse, J. (2020). Mapping county-level mobility pattern changes in the united states in response to covid-19. *Available at SSRN 3570145*.
- Garibaldi, P., Moen, E. R., and Pissarides, C. A. (2020). Modelling contacts and transitions in the sir epidemics model. *Covid Economics Vetted and Real-Time Papers, CEPR*.
- Gonzalez-Eiras, M. and Niepelt, D. (2020). On the optimal "lockdown" during an epidemic. *Covid Economics Vetted and Real-Time Papers, CEPR*.

- Greenstone, M. and Nigam, V. (2020). Does social distancing matter? *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2020-26).
- Hale, T., Webster, S., Petherick, A., Phillips, T., and Kira, B. (2020). Oxford covid-19 government response tracker. *Blavatnik School of Government*.
- Huber, M. and Langen, H. (2020). The impact of response measures on covid-19-related hospitalization and death rates in germany and switzerland.
- Jarvis, C. I., Van Zandvoort, K., Gimma, A., Prem, K., Klepac, P., Rubin, G. J., Edmunds, W. J., working group, C. C.-., et al. (2020). Quantifying the impact of physical distance measures on the transmission of covid-19 in the uk. *BMC Med*, 18(124).
- Karin, O., Bar-On, Y. M., Milo, T., Katzir, I., Mayo, A., Korem, Y., Dudovich, B., Yashiv, E., Zehavi, A. J., Davidovich, N., et al. (2020). Adaptive cyclic exit strategies from lockdown to suppress covid-19 and allow economic activity. *medRxiv*.
- Kermack, W. O., McKendrick, A. G., and Walker, G. T. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772):700–721.
- Kleven, H., Landais, C., and Sogaard, J. E. (2019). Children and gender inequality: Evidence from denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.
- Lai, S., Ruktanonchai, N. W., Zhou, L., Prosper, O., Luo, W., Floyd, J. R., Wesolowski, A., Zhang, C., Du, X., Yu, H., et al. (2020). Effect of non-pharmaceutical interventions for containing the covid-19 outbreak: an observational and modelling study. *medRxiv*.
- Lyra, W., do Nascimento, J. D., Belkhiria, J., de Almeida, L., Chrispim, P. P., and de Andrade, I. (2020). Covid-19 pandemics modeling with seir (+ caqh), social distancing, and age stratification. the effect of vertical confinement and release in brazil. *medRxiv*.

- Ng, K. Y. and Gui, M. M. (2020). Covid-19: Development of a robust mathematical model and simulation package with consideration for ageing population and time delay for control action and resusceptibility.
- Pang, W. (2020). Public health policy: Covid-19 epidemic and seir model with asymptomatic viral carriers.
- Pissarides, C. A. (2000). *Equilibrium unemployment theory*. MIT press.
- Roy, S. (2020). Covid-19 pandemic: Impact of lockdown, contact and non-contact transmissions on infection dynamics. *medRxiv*.
- Scala, A., Flori, A., Spelta, A., Brugnoli, E., Cinelli, M., Quattrocioni, W., and Pam-moli, F. (2020). Between geography and demography: Key interdependencies and exit mechanisms for covid-19. *Available at SSRN 3572141*.
- Singh, R. and Adhikari, R. (2020). Age-structured impact of social distancing on the covid-19 epidemic in india. *arXiv preprint arXiv:2003.12055*.

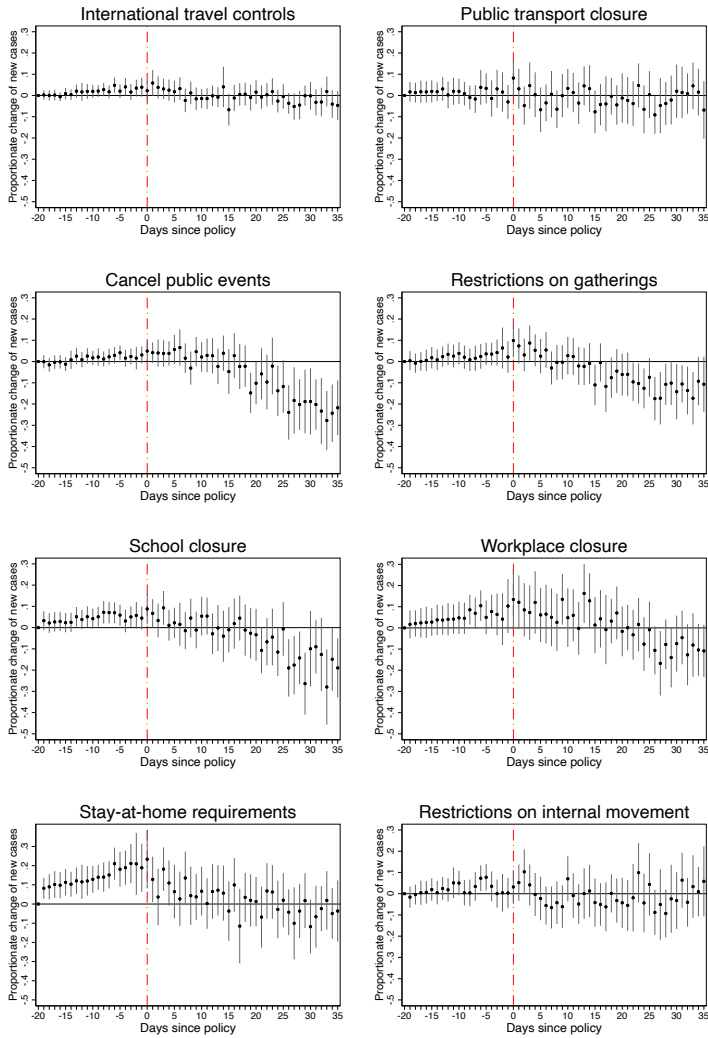
A Appendix - Figures



Note: Data from Hale et al. (2020), European CDC and own calculations

Figure A1: Effects of lockdown policies on COVID-19 confirmed new cases (in logs) without controlling for concurrent policies.

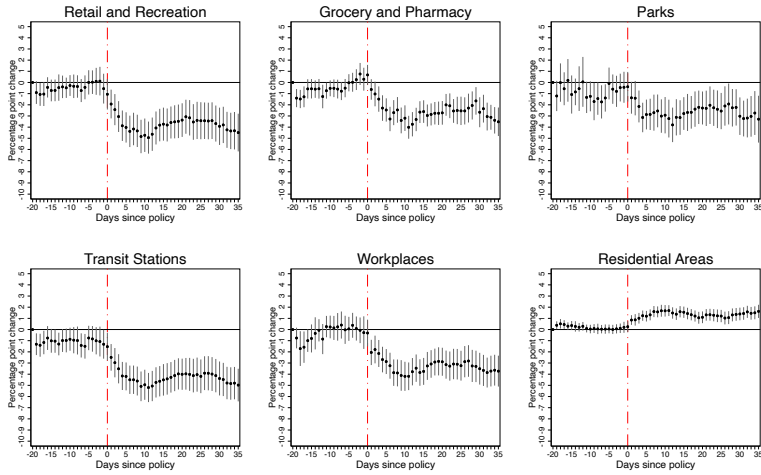
Covid Economics 23, 28 May 2020: 263-302



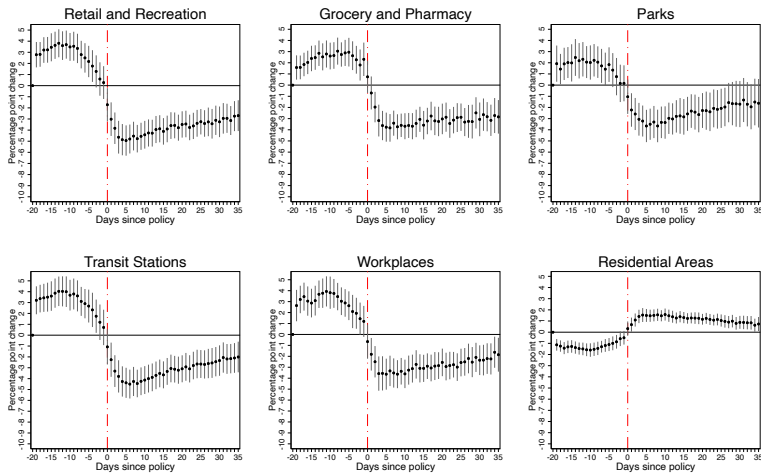
Note: Data from Hale et al. (2020), European CDC and own calculations

Figure A2: Effects of lockdown policies on COVID-19 confirmed new cases (in logs) controlling for concurrent policies.

(a) International travel controls

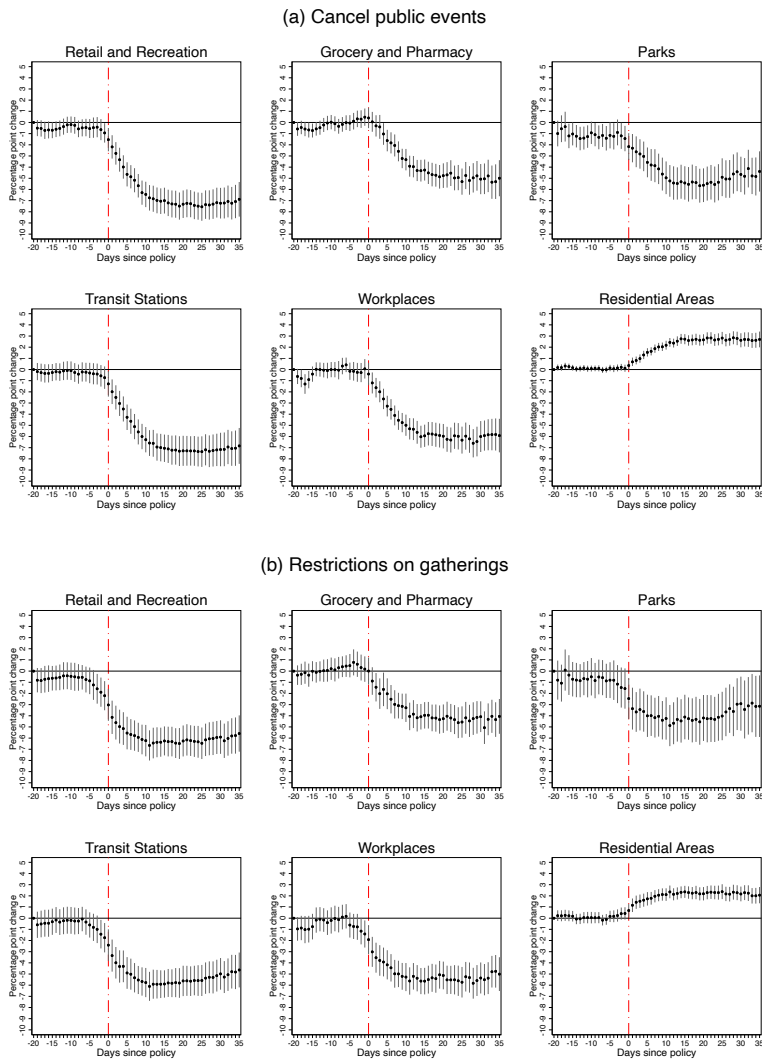


(b) Public transport closure



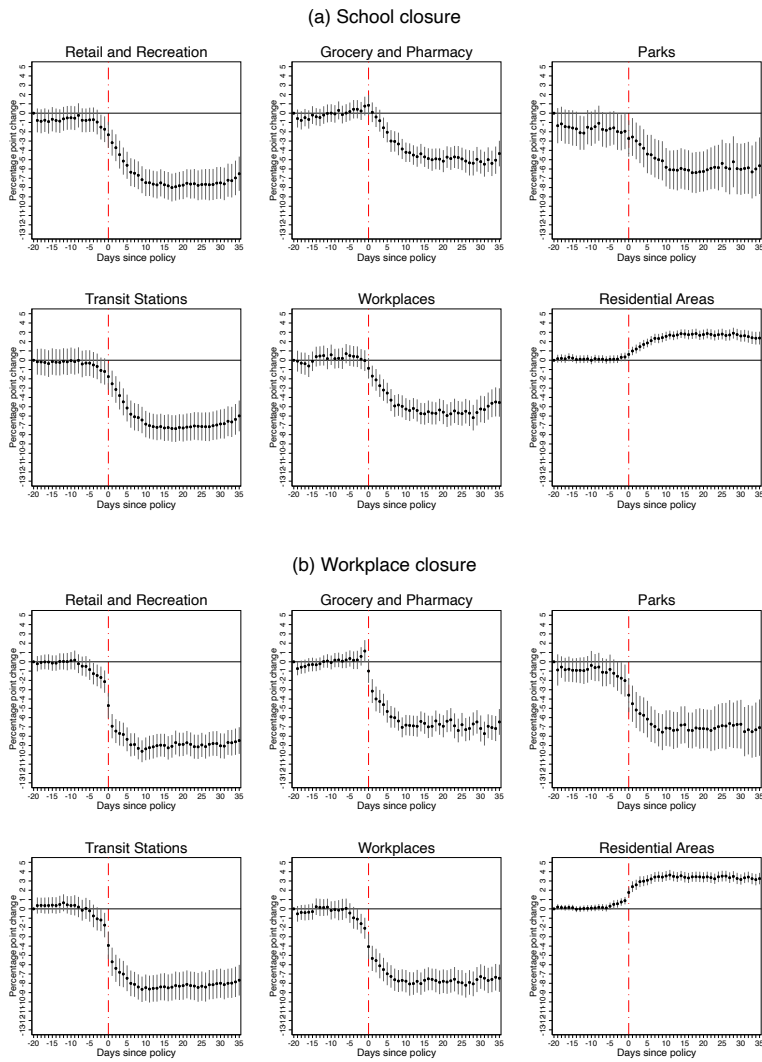
Note: Data from Hale et al. (2020), Google Community Mobility Reports and own calculations

Figure A3: Effects of **international travel controls** (panel a) and **public transportation closure** (panel b) on Google mobility patterns without concurrent policy controls.



Note: Data from Hale et al. (2020), Google Community Mobility Reports and own calculations

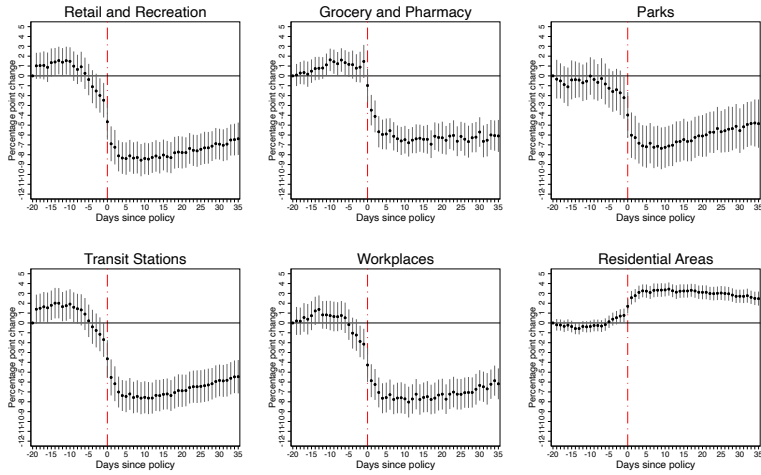
Figure A4: The effect of **public events cancellations** (panel a) and **restrictions on gatherings** (panel b) on Google mobility patterns without concurrent policy controls.



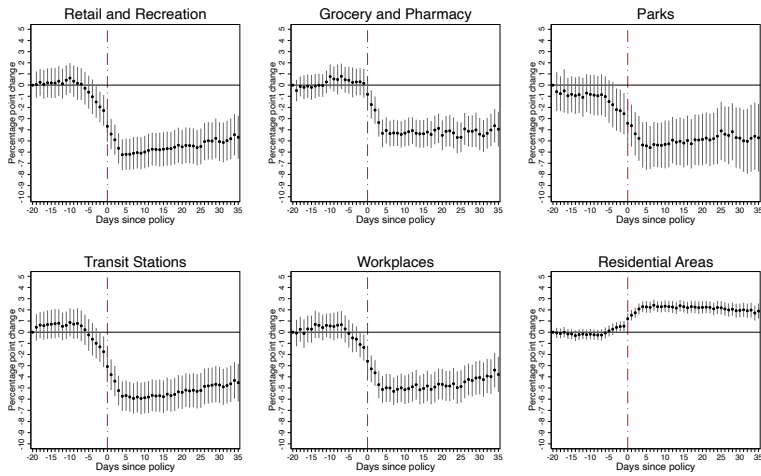
Note: Data from Hale et al. (2020), Google Community Mobility Reports and own calculations

Figure A5: Effects of **school** (panel a) and **workplace** (panel b) closures on Google mobility patterns without concurrent policy controls.

(a) Stay-at-home requirements



(b) Restrictions on internal movement



Note: Data from Hale et al. (2020), Google Community Mobility Reports and own calculations

Figure A6: Effects of **stay-at-home requirements** (panel a) and **restrictions on internal mobility** (panel b) on Google mobility patterns without concurrent policy controls.

B Appendix - Sample of countries

Estimations for COVID-19 cases are based on a sample of 135 countries

- Afghanistan, Angola, Argentina, Australia, Austria, Bahrain, Bangladesh, Barbados, Belgium, Belize, Bolivia, Bosnia and Herzegovina, Botswana, Brazil, Bulgaria, Burkina Faso, Cameroon, Canada, Cape Verde, Chile, Colombia, Costa Rica, Croatia, Czech Republic, Denmark, Dominican Republic, Ecuador, Egypt, El Salvador, Estonia, Finland, France, Gabon, Germany, Ghana, Greece, Guatemala, Honduras, Hong Kong, Hungary, India, Indonesia, Iraq, Ireland, Israel, Italy, Jamaica, Japan, Jordan, Kazakhstan, Kenya, Kuwait, Kyrgyz Republic, Laos, Lebanon, Libya, Luxembourg, Malaysia, Mali, Mauritius, Mexico, Moldova, Mongolia, Mozambique, Myanmar, Namibia, Netherlands, New Zealand, Nicaragua, Niger, Nigeria, Norway, Oman, Pakistan, Panama, Papua New Guinea, Paraguay, Peru, Philippines, Poland, Portugal, Puerto Rico, Qatar, Romania, Rwanda, Saudi Arabia, Serbia, Singapore, Slovak Republic, Slovenia, South Africa, South Korea, Spain, Sri Lanka, Sweden, Switzerland, Tanzania, Thailand, Trinidad and Tobago, Turkey, Uganda, United Arab Emirates, United Kingdom, United States, Uruguay, Vietnam, Zambia, Zimbabwe.
- No school intervention: Burundi, Nicaragua.
- No workplace intervention: Brunei, Burundi, Eswatini, Mozambique, Nicaragua, Niger, Tanzania.
- No events intervention: Burundi, Nicaragua, Sweden.
- No transport intervention: Australia, Brunei, Bulgaria, Burundi, Canada, Chile, Czech Republic, Dominica, Estonia, Germany, Hong Kong, Iceland, Japan, Malawi, Malaysia, Mali, Mauritania, Mozambique, Namibia, Nicaragua, Niger, Panama, South Korea, Sweden, Switzerland, Tanzania, Zambia.
- No mobility intervention: Burundi, Hong Kong, Iceland, Malawi, Mozambique, Nicaragua, Tanzania.
- No travel intervention: Luxembourg, United Kingdom.
- No home intervention: Brunei, Burundi, Cameroon, Iceland, Nicaragua, Norway, Sweden, Tanzania.

Estimations for Google Mobility are based on a sample without these 27 countries

- Albania, Algeria, Azerbaijan, Brunei, Burundi, Chad, China, Cyprus, Democratic Republic of Congo, Dominica, Eswatini, Ethiopia, Gambia, Guyana, Iceland, Lesotho, Madagascar, Malawi, Mauritania, Morocco, Palestine, Russia, Seychelles, Sierra Leone, Tunisia, Ukraine, Uzbekistan.

Impact of the state of emergency declaration for Covid-19 on preventive behaviours and mental conditions in Japan: Difference in difference analysis using panel data¹

Eiji Yamamura² and Yoshiro Tsutsui³

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

During the Covid-19 epidemic in Japan between March and April 2020, Internet surveys were conducted to construct panel data to investigate changes at the individual level regarding preventive behaviors and mental conditions by surveying the same respondents at different times. Specifically, the difference-in-difference (DID) method was used to explore the impact of the Covid-19 state of emergency declared by the government. Key findings were: (1) the declaration led people to stay home, while also generating anger, fear, and anxiety. (2) The effect of the declaration on the promotion of preventive behaviors was larger than the detrimental effect on mental conditions. (3) Overall, the effect on women was larger than that on men.

1 We would like to thank Editage (<http://www.editage.com>) for editing and reviewing this manuscript for English language.

2 Seinan Gakuin University.

3 Kyoto Bunkyo University.

Copyright: Eiji Yamamura and Yoshiro Tsutsui

**Impact of the State of Emergency Declaration for COVID-19 on Preventive Behaviors and
Mental Conditions in Japan: Difference in Difference Analysis using Panel Data**

EIJI YAMAMUR² AND YOSHIRO TSUTSUI

I. Introduction

The COVID-19 epidemic has had a significant impact on social and economic conditions, resulting in drastic changes in lifestyle, even though only a few months have passed since the first infected person was found in China in November 2019. Policymakers have been implementing various measures to mitigate COVID-19 pandemic. On May 16, 2020, the USA's death toll rose to 85,813, making it the highest official figure in the world, which was almost 2.5 times larger than UK, French, Italy, and Spain. At the same time, Japan's death toll was only 687³. A question arises here; was government's policies for COVID-19 control more effective than the USA and other countries? In this note, we examined the question in setting a quasi-natural experiment.

² Corresponding author: Eiji Yamamura, Seinan Gakuin University, 6-2-92 Nishijin Sawaraku Fukuoka, 814-8511, Japan (e-mail: yamaei@seinan-gu.ac.jp). Coauthor: Yoshiro Tsutsui, Kyoto Bunkyo University, Japan (email: tsutsui@econ.osaka-u.ac.jp);. Acknowledgments: We would like to thank Editage [<http://www.editage.com>] for editing and reviewing this manuscript for English language.

³ Johns Hopkins University of Medicine, CORONAVIRUS RESOURCE CENTER. <https://coronavirus.jhu.edu/map.html>. (Accessed on May 16, 2020)

In various countries including China, UK, and the USA, governments have employed a strong measure of lockdown enforcement in response to surging numbers of persons infected by the virus. Countries that have implemented drastic measures, such as lockdown, have seen reductions in the speed of the pandemic spread (Fang, Wang, & Yang, 2020; Tian et al., 2020). On the other hand, a rapid increase in domestic violence has been observed (WHO 2020). However, the negative effects of government policies have not been not sufficiently investigated, as even the closure of schools and non-essential businesses, could increase psychological costs such as the deterioration of mental conditions.

Fetzer et al. (2020) conducted a large-scale web survey between late March and early April covering 58 countries to investigate preventive behaviors and mental conditions within the population. Furthermore, they assessed the changes in the evaluations that people made during this period concerning government policies. However, they did not survey the same respondents to construct panel data and the sample was comprised of countries with different economic and social conditions. Therefore, they could not disentangle the effects of government policies on people's perceptions and behaviors from those other factors. Layard et al. (2020) analyzed the costs and benefits of the lockdown in the UK by considering not only traditional economic indices such as income and unemployment but also mental health⁴.

⁴ Oswald and Pawdthavee (2020) indicated that releasing from lockdown UK citizens aged 20-30 years who did not live with their parents was effective in increasing economic and social benefits and did not lead to higher numbers of COVID-19 victims.

However, no study has evaluated the effectiveness of policy measures to mitigate the COVID-19 epidemic by random assignment, although a few studies have assessed this aspect using simulation⁵. “One possible reason for the lack of systematic testing of mitigation measures in the COVID-19 pandemic is that it is both ethically and practically challenging: standard impact evaluation approaches typically require random assignment of some regions to an intervention and others to a control condition” (Haushofery and Metcalf, 2020, p.3).

On April 7, 2020, the Japanese government declared a state of emergency for COVID-19 in the prefectures where the number of persons infected with the virus was very large. Nine prefectures were clearly affected by the COVID-19 epidemic. However, the declaration was held only for seven of these nine prefectures. In this study, the seven prefectures where the declaration was held were defined as the treatment group. The two remaining prefectures were defined as the control group. The number of individuals infected by COVID-19 in the treatment group was not statistically different from that of the control group, as shown in Figure 1. On the other hand, there was a remarkable and statistically significant difference between the treatment and the control groups. The reason why the prefectures of Hokkaido and Aichi, which were the control group, were excluded from the declaration is ambiguous. This type of setting can be considered as quasi-randomization.

⁵The Susceptible-Infected-Recovered (SIR) model is used to compare the time course of infections in hypothetical control and treatment groups (Alvarez et al., 2020; Haushofery and Metcalf 2020). Atkeson (2020) also used the SIR model to consider the time course by dividing the total population into groups susceptible to the disease, infected by COVID-19, and others.




Figure 1. Comparison of mean number of persons infected by COVID-19 in each group between April 1 and 7, 2020.

Note: The error bar represents the 95% confidence intervals.

In Japan, similarly to the USA and European countries, governments have asked the population to change hygiene and social behaviors to help contain the spread of the disease (e.g., washing hands more carefully and avoiding social gatherings). Later, a request for more strict and costly measures, such as school closure and staying at home, was implemented. Differing from lockdown enforced in other countries, such as Italy, France, Germany, the UK, and the USA, the declaration of a state of emergency by the Japanese Government could not

substantially penalize by law those individuals who did not obey the government's request. In other words, citizens in Japan could decide whether or not to carry out preventive behaviors to mitigate the COVID-19 pandemic. Therefore, it seems plausible that the reactions of Japanese citizens to COVID-19 may vary compared to countries imposing more rigorous lockdowns. The situation in Japan can be considered a natural experiment to examine how the government's policy to appeal to conscience and morale generates changes in citizens' behaviors, which in turn affect their mental conditions under a state of emergency⁶.

The contribution of this study is to assess changes in preventive behaviors and mental conditions after a short period under the COVID-19 pandemic. For this purpose, the treatment and control groups were compared in order to explore how policies requiring preventive behaviors without penalty were effective in setting a quasi-natural experiment.

II. Setting and overview of data

A. Setting

In Japan, the first person infected by COVID-19 was observed on January 16, 2020. The number of infected persons has increased as time has passed, however, the pace of increase is much slower than that of the USA. We carried out Internet surveys to gather data

⁶ Ito et al. (2018) found that moral suasion was useful to persuade Japanese citizens to follow the request of saving electricity in a short period.

concerning citizens' preventive behaviors regarding COVID-19 and their mental conditions, exploring how they behaved and felt in response to the emergent situation of the pandemic.

INTAGE, a company with significant experience in academic research, was commissioned to conduct the Internet survey. The sampling method used was designed to gather a representative sample of the Japanese population regarding gender, age, and prefecture of residence. Our survey selected Japanese citizens between 16 and 79 years old from all regions of Japan. We conducted online surveys three times to assess the same individuals and construct the panel data within a month.

In the first wave, the sample size was 4,359, and its response rate was 54.7 %. In the second and third waves, we surveyed respondents from the first wave, and their response rates were 80.2% and 92.2%, respectively. Totally, observations were 11,867, which included 4,359 individuals. The first wave of the survey was conducted between March 13 and 16. The second and the third waves were carried out between March 27 and 30 and between April 10 and 13, respectively. During the whole period studied, the number of infected individuals in Japan increased from 675 (first wave) to 1,387 (second wave), and then 5,347 (third wave).

On April 7, between the second and third waves, the Government of Japan declared a state of emergency for seven prefectures that had heavily suffered from COVID-19, including Tokyo and Osaka⁷. The declaration requested that people should avoid going out of home

⁷ Besides Tokyo and Osaka, the following prefectures were included in the survey: Kanagawa, Chiba, Saitama, Hyogo, and Fukuoka.

unnecessarily, and also requested closing various public places including schools, museums, theaters, and bars, among others. At the time of the declaration, the request was planned to be valid for one month. Therefore, immediately after the declaration, the third wave of the survey was carried out. We were able to examine the effects of the declaration on citizens' behaviors and mental conditions by comparing them before and after the declaration.

The novel setting for this study was that the state of emergency was not declared in two prefectures, Hokkaido and Aichi, even though the number of infected persons in these prefectures was almost equivalent to that of the other seven prefectures. We divided the 47 prefectures into the treatment group, comprised of the seven prefectures which were the target of the declaration, the control group composed of the prefectures of Hokkaido and Aichi, and other prefectures. Figure 1 compares the mean values of individuals infected with COVID-19 between groups, showing no statistically significant differences between the treatment and control groups, while significant differences were found between the control and treatment groups when compared to other groups. Therefore, citizens from the treatment group and those from the control group experienced almost the same situation.

The population of the nine prefectures covering the treatment and control groups was equivalent to 53.1% of the population of Japan. Specifically, the seven prefectures of the treatment group signified 43.1% of the Japanese population. In the sample used in this

survey, eight prefectures had megacities of over one million people,⁸ with the exception of Chiba prefecture.

B. Features of the data

Observations for the treatment and control groups were 5,492 and 1,269, respectively. Therefore, total observations were 6,761. The survey questionnaire contained basic questions about demographics such as age, gender, educational background, household income, job status, marital status, and number of children. These data were constant in the first, second and third waves because the three waves were conducted within a month. In addition, respondents were asked questions concerning preventive behaviors, which are mentioned as follows:

“Within a week, to what degree have you achieved the following behaviors? Please answer in a scale from 1 (I have not achieved this behavior at all) to 5 (I have completely achieved this behavior).”

- (1) Stay indoors
- (2) Not go to the workplace (or school)
- (3) Wearing a mask
- (4) Washing hands carefully

⁸ Chiba prefecture included Chiba city which has 0.97 million people.

The answers to these questions were proxies for preventive behaviors: *Stay indoors*, *Do not work*, *Wear mask*, and *Washing hands*.

With regards to mental conditions, respondents were asked the following question:


“How much have you felt the emotions of anger, fear, and anxiety? Please answer in a scale from 1 (I have not felt this emotion at all) to 5 (I have felt this emotion strongly).”

(5) Anger


(6) Fear

(7) Anxiety

The answers to these questions were proxies for mental conditions: *Anger*, *Anxiety*, and *Fear*.



Keep indoors



Not going to work




Figure 2. Changes in preventive behaviors.

Note: The solid line indicates the treatment group, while dashed line shows the control group.

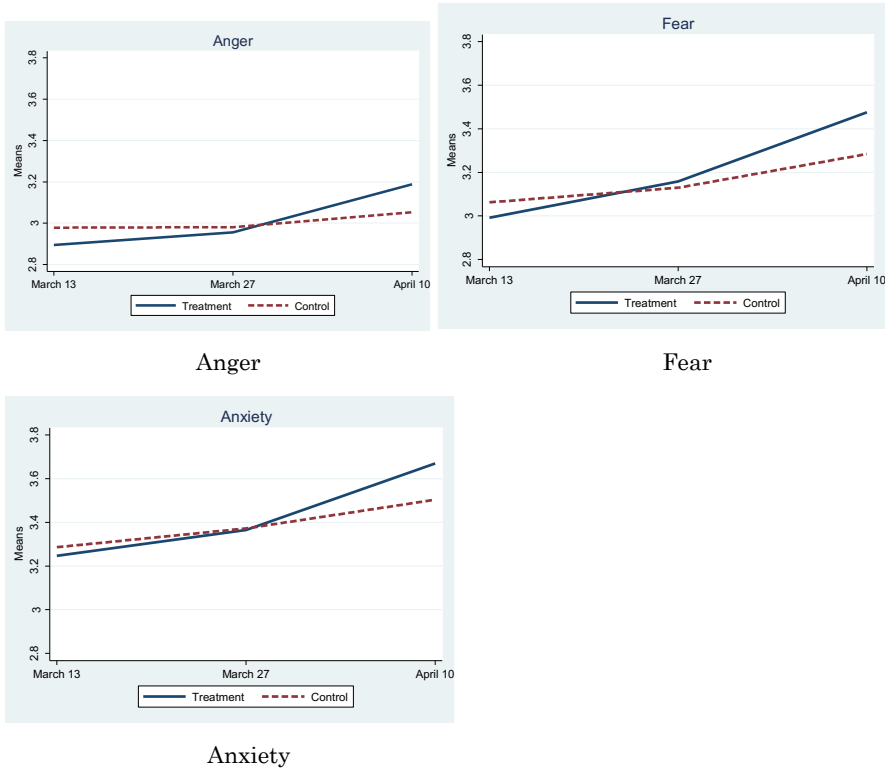


Figure 3. Changes in emotions.

Note: The solid line indicates the treatment group, while the dashed line shows the control group.

Figures 2 and 3 demonstrate the degree of changes in preventive behaviors and mental conditions of respondents between March 13 and April 10 by comparing the treatment and control groups. In the next section, we will explain the difference-in-difference (DID) method to examine the effects of the declaration. Based on this method, there was the key assumption that the trends of variables would be the same for the treatment and control groups before the declaration was announced (Angrist and Pischke, 2009). In Figures 2 and 3, we examined the common trends assumption, confirming that the trends regarding preventive behaviors and mental conditions were almost the same in the control and treatment groups, between March 13 and 27, before the declaration of a state of emergency. Therefore, the common trends assumption was confirmed in this study.

Between March 13 and 27, except for “washing hands,” levels of variables in the treatment group were lower than those of the control group. Later, between March 27 and April 10, the slopes of the treatment group became steeper than those of the control group, leading the mean values of the treatment group to be higher than those of the control group, with the exception of “washing hands.” This suggests that the declaration had a significant effect on citizens’ behaviors and mental conditions.

III. The econometric model

The DID method was used to examine the effects of the state of emergency declaration on preventive behaviors and mental conditions. Data was limited to nine prefectures and divided into the control and treatment groups, which included residents from two and seven prefectures, respectively. As discussed in the previous section, the DID method was considered valid. The estimated function takes the following form:

$$Y_{itp} = \alpha_1 \text{Wave3}_t \times \text{Treatment}_g + \alpha_2 \text{Wave3}_t + \alpha_3 \text{Wave2}_t + \alpha_4 \text{Infected COVID19}_{itg} + k_i + u_{itp},$$

In this formula, Y_{itp} represents the dependent variable for individual i , wave t , and group g . For the estimation of preventive behaviors, Y is preventive behaviors such as *Stay indoors*, *Do not work*, *Wear mask*, and *Washing hands*. Regarding the estimation of mental conditions, Y is *Anger*, *Anxiety*, and *Fear*. The second (*Wave 2*) and third wave (*Wave 3*) dummies were included while their reference group was the first wave. This describes the degree of change in the dependent variables compared to the first wave. *Treatment* is a dummy variable for the treatment group. Key variable was, $\text{Wave3} \times \text{Treatment}$, cross terms for *Wave3* and *Treatment*. When preventive behaviors were dependent variables, their coefficients were expected to be positive if the declaration promoted citizens to achieve preventive behaviors. On the other hand, when mental conditions were dependent variables, their coefficients were expected to be positive if the

declaration deteriorated the mental conditions of the population.

The time-invariant individual-level fixed effects are represented by k_i . Because of short-term panel data, most of the individual level demographic variables were considered as time-invariant features, which were completely described by k_i . The regression parameters are denoted by α . The number of persons infected with COVID-19 increased drastically in the residential areas during the studied period and, thus, it was included as a control variable. The error term was denoted by u .

IV. Results

Results are focused on the key variable $Wave3_t \times Treatment$. The results shown in Table 1 are based on the sub-sample comprised of the treatment and control groups shown in Figure 1. In this table, cross terms display a positive coefficient in all results⁹. Furthermore, we observe a statistical significance at a level of 1 %, with the exception of “*Washing hands*,” which was not significant. The appropriate setting of the control group shows strong evidence that declaring a state of emergency promoted preventive behaviors while at the same time deteriorated mental health. With regards to preventive behaviors, the coefficients for “*Stay*

⁹ The Appendix presents results based on the full sample, where the control group was comprised of respondents from 40 prefectures including the “Control group” and “Others,” indicated in Figure 1. Cross terms show a positive coefficient with the exception of “*Washing hands*.” Furthermore, statistical significance was observed for “*Stay indoors*,” “*Do not work*,” and “*Anger*.”

indoors” and “*Do not work*” were larger than those for “*Wear mask*” and “*Washing hands.*”

Our interpretation of these results is that citizens were requested to wear a mask and wash hands mainly when they went out. Naturally, people may have considered staying indoors as more important than wearing a mask and washing hands.

Table 1. Fixed effects model: Sub-sample of heavily infected areas.

	Preventive behaviors					Mental conditions	
	<i>Keep indoors</i>	<i>Not going to work</i>	<i>Wear mask</i>	<i>Washing hands</i>		<i>Anger</i>	<i>Anxiety</i>
<i>Wave3</i> × <i>Treatment</i>	0.29*** (0.06)	0.32*** (0.07)	0.19*** (0.06)	0.06 (0.04)	0.18*** (0.06)	0.15*** (0.05)	0.15*** (0.05)
<i>Wave3</i>	0.21*** (0.05)	0.07 (0.06)	0.27*** (0.05)	0.11*** (0.03)	0.11** (0.05)	0.28*** (0.05)	0.24*** (0.05)
<i>Wave2</i>	0.10*** (0.03)	0.08** (0.03)	0.07*** (0.02)	0.03 (0.02)	0.05** (0.02)	0.15*** (0.02)	0.11*** (0.02)
<i>Infected COVID_19</i>	0.23 (0.07)	0.23*** (0.07)	0.05 (0.06)	0.03 (0.04)	-0.02 (0.06)	0.07 (0.05)	0.05 (0.05)
Within R- Square Groups Obs.	0.11 2,477 6,761	0.06 2,477 6,761	0.09 2,477 6,761	0.02 2,477 6,761	0.03 2,477 6,761	0.08 2,477 6,761	0.07 2,477 6,761

Note: Numbers within parentheses indicate robust standard errors clustered on individuals. ***, **, * indicate statistical significance at a level of 1%, 5%, and 10%, respectively.

To investigate gender differences concerning the effects of the declaration, we conducted estimations using sub-samples for males and females. In Table 2, we observe a positive coefficient in all cross terms. Differences between males (Panel A) and females (Panel B) are shown below.

Covid Economics 23, 28 May 2020: 303-324

Table 2. Fixed effects model: Sub-sample.

Panel A. Male sample

	Preventive behaviors					Mental conditions	
	<i>Keep indoors</i>	<i>Not going to work</i>	<i>Wear mask</i>	<i>Washing hands</i>		<i>Anger</i>	<i>Anxiety</i>
<i>Wave3</i> × <i>Treatment</i>	0.32*** (0.09)	0.21** (0.10)	0.17* (0.09)	0.002 (0.03)	0.22** (0.09)	0.10 (0.08)	0.14* (0.08)
<i>Wave3</i>	0.12 (0.08)	0.08 (0.09)	0.29*** (0.07)	0.14** (0.06)	0.05 (0.08)	0.29*** (0.07)	0.27*** (0.07)
<i>Wave2</i>	0.05 (0.04)	0.03 (0.04)	0.05 (0.03)	0.01 (0.03)	0.05 (0.04)	0.15*** (0.04)	0.12*** (0.03)
<i>Infected COVID_19</i>	0.21** (0.09)	0.29*** (0.10)	0.05 (0.09)	0.04 (0.06)	0.02 (0.09)	0.11 (0.08)	0.002 (0.08)
Within R-Square Groups Obs.	0.09 1,206 3,296	0.05 1,206 3,296	0.09 1,206 3,296	0.02 1,206 3,296	0.03 1,206 3,296	0.08 1,206 3,296	0.09 1,206 3,296

Panel B. Female sample

	Preventive behaviors					Mental conditions	
	<i>Keep indoors</i>	<i>Not going to work</i>	<i>Wear mask</i>	<i>Washing hands</i>		<i>Anger</i>	<i>Anxiety</i>
<i>Wave3</i> × <i>Treatment</i>	0.25*** (0.09)	0.43*** (0.10)	0.22*** (0.08)	0.11** (0.05)	0.14* (0.08)	0.20** (0.08)	0.16** (0.08)
<i>Wave3</i>	0.30*** (0.07)	0.06 (0.08)	0.25*** (0.06)	0.09* (0.05)	0.17** (0.07)	0.27*** (0.07)	0.21*** (0.07)
<i>Wave2</i>	0.14*** (0.04)	0.12*** (0.04)	0.10*** (0.03)	0.04 (0.03)	0.06* (0.03)	0.15*** (0.03)	0.10*** (0.03)
<i>Infected COVID_19</i>	0.25*** (0.09)	0.18* (0.11)	0.06 (0.08)	0.02 (0.05)	-0.06 (0.07)	0.04 (0.07)	0.09 (0.07)
Within R-Square Groups Obs.	0.14 1,271 3,465	0.07 1,271 3,465	0.10 1,271 3,465	0.03 1,271 3,465	0.03 1,271 3,465	0.09 1,271 3,465	0.08 1,271 3,465

Note: Numbers within parentheses indicate robust standard errors clustered on individuals. ***, **, * indicate statistical significance at a level of 1%, 5%, and 10%, respectively.

First, all columns for females show a statistical significance, whereas statistical significance was not observed regarding two cross terms, “*Wash hands*” and “*Anxiety*.”

Hence, as a whole, the declaration had more significant effects on females than in males.

Washing hands is different from other preventive behaviors in that it is a behavior less likely to be observed by other people. This seems to reduce the incentive to wash hands. In other words, the statistical significance observed in females could mean that women may have the incentive to wash their hands even if other people are not observing their behavior. As Agüero and Belecche (2017) indicate, an exogenous health shock, such as a pandemic, facilitates the adoption of low-cost health behaviors, such as hands washing, which provides long-lasting effects on health outcomes. Therefore, the role of women becomes important to have a long-term effect on the general acceptance of handwashing in a society.

Second, the coefficient value for “*Anger*” in males was two times larger than in females. We may interpret this as suggesting that an increase in anger in husbands could result in an increase in domestic violence against their wives during the state of emergency. On the other hand, females were more likely to feel anxiety and fear than males, which seems to result in mental illness. This could cause social agitation.

Overall, the declaration of a state of emergency not only had positive effects on preventive behaviors addressed to mitigate the pandemic, but also negative effects on mental conditions which may increase domestic violence and social unrest. However, in most of cases, the coefficient values of preventive behaviors were larger than those of mental conditions.

V. Conclusion

The evaluation of government policies should be analyzed by considering their costs and benefits. The purpose of this study was to examine how the declaration of a state of emergency for COVID-19 changed preventive behaviors and mental conditions in Japan. Using individual-level panel data collected through short-term repeated Internet surveys, we conducted DID estimations. After controlling for individual fixed-effects, key findings were: (1) the declaration led people to stay home, while also generating anger, fear, and anxiety. (2) The effect of the declaration on the promotion of preventive behaviors was larger than the detrimental effect on mental conditions. (3) Overall, the effect on women was larger than that on men. In short, we found that the declaration promoted preventive behaviors and at the same time deteriorated mental conditions. More specifically, an increase in anger in husbands is remarkably larger than wives, which could result in an increase in domestic violence against their wives during the state of emergency.

An increase in anger from staying indoors is thought to cause domestic violence. Considering this aspect is important when evaluating the outcomes of the state of emergency declaration in Japan as well as lockdowns in Italy, France, Spain, the United Kingdom, and the United States. Moreover, it is necessary to evaluate government policies through cost-benefit analysis from a long-term viewpoint. Further research should investigate these aspects to scrutinize whether Japanese government's policy is more effective and efficient than

policies adopted by the USA, UK, French, Italy, and Spain.

REFERENCES

Aguero, Jorge, and Trinidad M. Beleche. 2017. "Health Shocks and Their Long-Lasting

Impact on Health Behaviors: Evidence from the 2009 H1N1 Pandemic in Mexico."

Journal of Health Economics, 54: 40-55.

Alvarez, Fernando, Argente, David, and Francesco Lippi. 2020. "A Simple Planning Problem

for Covid-19 Lockdown." *COVID Economics Vetted and Real-Time Papers*, 14(6), 1-32.

<https://cepr.org/sites/default/files/news/CovidEconomics14.pdf>

Angrist, Joshua, and Jorn-Steffen D. Pischke. 2009. *Mostly Harmless Econometrics: An*

Empiricist's Companion. Princeton: Princeton University Press.

Atkeson, Andrew. 2020. "What Will Be the Economic Impact of Covid-19 in the US? Rough

Estimates of Disease Scenarios." NBER Working Paper 26867.

<http://www.nber.org/papers/w26867>.

Oswald, Andrew, and Powdthavee Nattavudh. 2020. "The Case for Releasing the Young from

Lockdown: A Briefing Paper for Policymakers." IZA Discussion Paper 13113.

Fang, H., Wang, L., and Y. Yang. 2020. "Human Mobility Restrictions and the Spread of the

Novel Coronavirus (2019-nCoV) in China.” NBER Working Paper 26906.

<https://doi.org/10.3386/w26906>.

Fetzer, Thiemo, Witte, Marc, Hensel, Lukas, Jachimowicz, Jon M., ... Erez Yoeli. 2020.

“Global Behaviors and Perceptions in the COVID-19 Pandemic.” PsyArXiv, April

16. <https://doi.org/10.31234/osf.io/3kfmh>

Haushofery, Johannes, and Jessica E. Metcalf. 2020. “Evaluation of Non-Pharmaceutical

Interventions is Needed to Mitigate the COVID-19 Pandemic.” Working Paper of

Princeton University.

Ito, Koichiro, Ida, Takanori, and Makoto Tanaka. 2018. “Moral Suasion and Economic

Incentives: Field Experimental Evidence form Energy Demand.” *American*

Economic Journal: Economic Policy, 10(1): 240-67.

Layard, Richard, Clark, Andrew, De Neve, Jan-Emmanuel, Krekel, Christian, Fancourt,

Daisy, Hey, Nancy, and Gus O’Donnell. 2020. “When to Release the Lockdown? A

Wellbeing Framework for Analysing Costs and Benefits.” IZA Discussion Paper

13186.

Li, Sijia, Wang, Yilin, Xue, Jia, Zhao, Nan, and Tingshao Zhu. 2020. “The Impact of COVID-

19 Epidemic Declaration on Psychological Consequences: A Study on Active

Weibo Users.” *International Journal of Environmental Research and Public*

Health, 17(6), 2032. <https://doi.org/10.3390/ijerph17062032>

Tian, H., Liu, Y., Li, Y., Wu, C.-H., Chen, B., Kraemer, M. U. G., . . . C. Dye. 2020. “An Investigation of Transmission Control Measures During the First 50 Days of the COVID-19 Epidemic in China. *Science*. <https://doi:10.1126/science.abb6105>

World Health Organization (WHO). 2020. “COVID-19 and Violence Against Women. What the Health Sector/System Can Do.”

<https://apps.who.int/iris/bitstream/handle/10665/331699/WHO-SRH-20.04-eng.pdf>

Appendix. Results using the full sample.

	Preventive behaviors				Mental conditions		
	<i>Keep indoors</i>	<i>Not going to work</i>	<i>Wear mask</i>	<i>Washing hands</i>	<i>Anger</i>	<i>Anxiety</i>	<i>Fear</i>
<i>Wave3</i> <i>× Treatment</i>	0.11** (0.05)	0.23*** (0.05)	0.05 (0.04)	-0.02 (0.03)	0.08** (0.04)	0.06 (0.04)	0.05 (0.05)
<i>Wave3</i>	0.40*** (0.03)	0.17*** (0.03)	0.41*** (0.03)	0.20*** (0.02)	0.19*** (0.02)	0.36*** (0.02)	0.33*** (0.02)
<i>Wave2</i>	0.13*** (0.02)	0.09*** (0.02)	0.08*** (0.02)	0.04*** (0.01)	0.03 (0.02)	0.13*** (0.02)	0.09*** (0.02)
<i>Infected COVID_19</i>	0.21*** (0.07)	0.23*** (0.07)	0.04 (0.06)	0.02 (0.04)	-0.02 (0.06)	0.07 (0.05)	0.05 (0.04)
Within R- Square Groups Obs.	0.09 4,359 11,867	0.04 4,359 11,867	0.09 4,359 11,867	0.03 4,359 11,867	0.03 4,359 11,867	0.07 4,359 11,867	0.07 4,359 11,867

Covid Economics 23, 28 May 2020: 303-324

Fast and local: How lockdown policies affect the spread and severity of covid-19

Jean-Philippe Bonardi,¹ Quentin Gallea,² Dimtrija Kalanoski³ and Rafael Lalive⁴

Date submitted: 20 May 2020; Date accepted: 22 May 2020

We analyse whether the various types of lockdowns implemented around the world mitigated the surge in infections and reduced mortality related to the Covid-19, and whether their effectiveness differed in developing versus developed countries. Our data cover 184 countries from December 31st 2019 to May 4th 2020, and identifies when lockdowns were adopted, along with confirmed cases and deaths. We find that reducing movements within countries has been effective in developed economies – averting about 650,000 deaths – but not in developing ones, that countries that acted fast fared better, and that closing borders has had no appreciable effect, even after 50 days.

- 1 Professor, University of Lausanne – HEC Lausanne (Faculty of Business and Economics, Enterprise for Society Center (E4S).
- 2 PhD student, University of Lausanne – HEC Lausanne (Faculty of Business and Economics, Enterprise for Society Center (E4S).
- 3 Post-doctoral student, University of Lausanne – HEC Lausanne (Faculty of Business and Economics, Enterprise for Society Center (E4S).
- 4 Professor, University of Lausanne – HEC Lausanne (Faculty of Business and Economics, Enterprise for Society Center (E4S).

Copyright: Jean-Philippe Bonardi, Quentin Gallea,
Dimtrija Kalanoski and Rafael Lalive

1 Introduction

The first death from the coronavirus, on January 11th 2020, was a 61-year-old man in China who had purchased goods from a seafood market. In the middle of May 2020, a few months later, over 300,000 deaths had been registered, and the health and economic effects of the covid-19 have turned out to be massive. There is a lot to learn, however, from looking at how governments around the world have responded to this crisis, and how it impacted the development of the disease. This is the purpose of this paper.

Throughout the first months of the spread of the virus, two alternative strategies have indeed emerged to fight the covid-19 pandemic. First, the so-called 'herd immunity approach', according to which the viral dissemination through the population was critical to develop collective immunity. From this perspective, the only public policy that had to be put in place was one in which patients at risk or infected had to be isolated and taken care of. The second major policy option that emerged from the crisis is the 'lockdown approach', in which most of a country's population had to stay at home to stop the virus dissemination, avoid over-crowding critical care hospital facilities and prevent the deaths of many people. China was the first large country to announce this type of lockdown policy on January 23rd, 2020. Even though the governments of some influential countries (such as the US or the UK) had originally chosen the herd immunity approach, things rapidly evolved and, within a few days in March, most governments had opted for the lockdown approach in a hurry. Now that some time has gone by, it is important to take a closer empirical look at the real impact of these lockdowns on the disease.

Here, we study the effect of lockdown policies, as well as their differences in terms of speed, strength and nature across countries on the increase of new cases and mortality (Flaxman et al. (2020)). Beyond the general question of whether lockdowns are effective or not, several other more subtle aspects can be informative for policy-making. One of them is to know whether there

is a ‘speed premium’ in setting up lockdowns. With a virus growing exponentially, one could expect to observe an advantage for early movers. However, there might also be an option to wait, for instance because countries might learn from what happens to others. So, can we observe a speed premium in the Covid-19 case? Another interesting and politically sensitive issue is related to international border closures. Beyond internal lockdowns, most countries have been closing their borders, something that seem logical to handle a pandemic in a globalized world. But some countries did it right away, while others did it at the last resort. Did closing borders really matter to slow down the spread of the virus? Did the order of the national-international sequence have an impact?

In exploring these questions, endogeneity issues could be major hurdles in order to establish causality, in particular omitted variable bias, reverse causality and measurement errors. We address these issues explicitly in our empirical approach (see Section 3). The panel structure of our dataset, composed of 184 countries, allows us to control for country fixed effects and day fixed effects. Furthermore, we also control for the within-country evolution of the disease both by using a lagged outcome and by controlling for the number of days since the first case was reported in the country.

From an economics perspective, we also explore the underlying mechanisms that can explain why certain types of lockdown measures are more effective than others, and why these might work better in some places than others. The hypothesis driving our empirical investigation is that lockdowns to be effective have to drive down individuals’ opportunity costs of staying home. As long as these opportunity costs are high enough, one could expect that people might not abide by lockdown restrictions, especially since the cost for authorities of monitoring what individuals are doing should typically be quite high. This issue is of particular importance for the effectiveness of lockdown policies in developing countries. Indeed, in these countries where a large number of people earn their living in the underground economy and do not have

social insurance, this opportunity cost approach would predict that lockdown measures will be less effective than in developed countries. This is what our empirical analysis shows. We will return to this issue in the Discussion section.

A few papers have already studied the impact of non-pharmaceutical measures interventions (NPIs) on pandemics and more particularly on the covid-19 (Harris (2020); Hartl et al. (2020); Flaxman et al. (2020)). Chinazzi et al. (2020) and Kraemer et al. (2020) explore to what extent China's travel ban, human mobility, and control measures reduced the spread of the disease, and Maier and Brockmann (2020) finds that measures put in place in China before the lockdown contributed to slow down its viral dissemination. Additionally, Giordano et al. (2020) compare simulation results with real data on the covid-19 epidemic in Italy and show that restrictive social-distancing measures are effective, but their effectiveness could be further enhanced if combined with widespread testing and contact tracing. Hatchett et al. (2007) study cities in the United States and the non-pharmaceutical interventions they adopted to curb the spreading of the Spanish Influenza. Whereas these papers focus on one country, our analysis covers most countries in the world, which allows us to leverage the heterogeneity regarding how lockdowns were implemented. In some cases, in effect, lockdowns were strict and complete, while in others they were partial. In some cases, there was a curfew and in some others not; in some countries, borders were closed right away, whereas in some others bordure closure was the last measure to be taken. As we will see below, these differences matter.

2 Data

We compiled information regarding the lockdown policies undertaken by countries around the world. Using a web-scraping program, we extracted from LexisNexis all news headlines for each country from October 31st, 2019 to April 1st, 2020, and all per country information from US Embassy Covid-19 bulletin. We cross-checked the news headline data against the data

from the US Embassy Covid-19 bulletin to ensure its validity. The final dataset allowed us to generate dates of implementation for several measures designed to stop the spread of the Covid-19, some being internal to the country and oriented towards the outside (See Figure 1). Two measures, *State of Emergency* and *Curfew*, significantly restrain movement of individuals within a country, and thus represent a form of total lockdown within a country. We combined *State of Emergency* and the *Curfew* into one measure, which we call *Total within country lockdown* (see Supplementary Material).

	Measures	Explanation	Example	Severity
INTERNAL MEASURES	<i>Curfew</i>	The effective date when a country announced a limit of the movement of individuals within a given time of the day.	<i>17th of March 2020:</i> Bosnia declares nationwide state of emergency over coronavirus.	[Color scale from red to yellow]
	<i>State of emergency</i>	The effective date when a country announced a state of emergency.	<i>21st of March 2020:</i> President Roch Marc Christian Kaboré closed airports, land borders and imposed a nationwide curfew to curb the spread of the pandemic.	
	<i>Within country regional lockdown</i>	The earliest effective date when a region within a country announced that it will be entering a total lockdown	<i>17th of March 2020:</i> Quebec. Declares State of Emergency to Blunt Pandemic.	
	<i>Partial selective lockdown</i>	The earliest effective date when the earliest effective date when a country announced partial limitation of movement by implementing school closures, limiting the number of people permitted to gather in a group, and/or closure of religious institutions.	<i>16th of March 2020:</i> Cambodia Announces Nationwide School Closures as COVID Response Ramps Up.	
EXTERNAL MEASURES	<i>Selective border close stage 1</i>	The earliest effective date when a country closed borders towards any other significantly affected region or country in the world from Covid-19 (Wuhan, China, Iran, and Italy - individually or as a group).	<i>30th of January 2020:</i> Australia banned the entry of foreign nationals from mainland China.	[Color scale from yellow to red]
	<i>Selective border close stage 2</i>	The earliest effective date after <i>Selective border closure stage 1</i> when a country closed borders towards one or multiple other countries in the world significantly affected from Covid-19.	<i>22nd of February 2020:</i> Fiji extended its travel ban and announced that travelers from Italy, Iran and the South Korean cities of Daegu and Cheongdo would be denied entry.	
	<i>International lockdown</i>	The effective date when a country closed its borders for all flights, rail, and automotive movement internationally.	<i>30th of March 2020:</i> Council of Ministers of Bosnia and Herzegovina issued a decision which bans entrance for all foreigners.	

Figure 1: Lockdown Policies Implemented Around the World

Note: The state of emergency is a situation in which a government is empowered to perform actions or impose policies that it would normally not be permitted to undertake, that is, restriction of movement of individuals and closure of non-essential and essential (if necessary) public and private entities.

We use the John Hopkins University data on the number of cases testing positively for Covid-19 infections (Dong et al. (2020)), as it seems to be the most complete and reliable source regarding reported cases and deaths. We focus here on the number of new infected cases (results on deaths in Supplementary Material), and that for three reasons. First, people who die from the virus got infected first. Hence, controlling the number of contaminated persons in-

Covid Economics 23, 28 May 2020: 325-351

evitably reduces the number of deaths. Second, a major objective for the public management of the pandemic, which is reflected in the “flattening the curve” argument, is to prevent hospitals from being overwhelmed by patients in need of intensive care. Hence, the number of people infected by the virus is a better indicator for the future burden on the healthcare sector than the number of patients who have already passed away. Finally, there is clearly a delay in how a lockdown measure can affect the number of deaths: the patient has to contract the virus, pass the incubation time, experience complications and then eventually die. This process is potentially long and might vary from patient to patient, which might make it harder to observe clear relationships. Our data, of course, represents a lower bound on the total number of people ever infected by the virus; but what is important for us here is to have a measure of the number of people who need medical attention. These people are symptomatic, and possibly quite well represented in our data. Measurement errors will affect our dependent variable, but our estimates should not be greatly affected by them (see Supplementary Material).

It is important to note that the data on the Covid-19 suffer from measurement errors. The dataset contains reported cases which are not equivalent to the total number of cases infected by the virus in the country. To observe reported cases, these have to be reported first. Hence, the person has to be tested, recorded and observed by the John Hopkins University team. However, those three conditions are not met for many individuals. First, the person has to be tested and in most countries, this person requires to have symptoms or even severe symptoms to be tested. When there is no systematic testing (which is the case for an overwhelming majority of countries), asymptomatic people or people contaminated but not experiencing symptoms yet (because of incubation time) are not observed. Second, the new case has to be recorded and transmitted to the authorities or some statistics institute. Some countries are suspected to under-report or modify the data¹. Third, this information has to reach the sources watched by John

¹Can China’s COVID-19 Statistics Be Trusted? (last accessed: 14.04.20) <https://thediplomat.com/2020/03/can->

Hopkins University. Hence, our data represents a lower bound on the total number of people ever infected by the virus. Yet, in our context, we need a measure of the number of people who will need medical attention. These people are symptomatic and possibly quite well represented in our data. Moreover, measurement errors affect our dependent variable, and our estimates should not be greatly affected by them.

A more worrying problem would be the presence of non-classical errors-in-variables. For example, if countries which under-report systematically the number of cases are countries with a lower quality health sector, potentially autocracies. However, as we use country fixed effects in our empirics, these time-invariant unobservables, which might generate measurement errors, are controlled for.

Governments relied on a variety of measures with different levels of strictness to mitigate the effects of Covid-19. On the one hand, many governments focused on what we call "outside measures", i.e. partially or totally restricting international movement from and to a given country for individuals of other countries (International Lockdown of the Country, Selective border closure stage 1 and Selective border closure stage 2). On the other hand, governments took "inside measures", which ranged from closing specific regions within the country (Within country regional lockdown), implementing partial selective lockdown on public and private institutions (Partial selective lockdown) to other stricter measure such as declaring a State of emergency or setting-up Curfews.

Finally, to study the existence of heterogenous effect between developed and developing countries we use the Human Development Index (henceforth HDI) produced by the UN (Programme (2020)). The HDI is a composite index defined as the geometric mean of normalized indices ($\in [0; 1]$) for Life expectancy, Education and GNI. Note that the median in our sample

[chinas-covid-19-statistics-be-trusted/](https://www.economist.com/graphic-detail/2020/04/07/chinas-data-reveal-a-puzzling-link-between-covid-19-cases-and-political-events). China's data, in fact, reveal a puzzling link between covid-19 cases and political events (last accessed: 14.04.20) <https://www.economist.com/graphic-detail/2020/04/07/chinas-data-reveal-a-puzzling-link-between-covid-19-cases-and-political-events>.

is 0.745. We define developing countries as the ones with an index up to 0.699, which refers to Low and Medium human development using the United Nation codebook definition, while above 0.699 will be defined as developed countries (the exact list of countries can be found in the Supplementary Material).

Our final dataset is composed of 184 countries, of which 108 had implemented at least one of the measures at the time we collected the data, observed over 127 days, from the 31st of January 2019 to the 4th of May 2020. We adopt a calendar time definition where the 31st of December 2019 is the starting date, as it is the first day when a country other than China undertook measures to mitigate the Covid-19 dissemination². Figure 2a shows the number of measures taken, and the number of confirmed cases, and deaths, by time since the first measure has been taken. Governments initially adopted “inside” measures, during the period end of January and early February 2020 (20 to 40 days after Taiwan), and moved to outside measures later on. Figure 2b shows measures and confirmed cases by days since the first case is recorded in a country. Countries implement measures during the first three weeks after the first case has been recorded, when the average number of cases is still low.

3 Methods

Our main results are based on models of the growth rate in the total number of confirmed cases in a country (see Supplementary Material for alternative approaches, including the ones about the number of deaths). The growth rate in the number of cases, or new infections, captures whether the lockdown measures reduced the spread of the disease Avery et al. (2020). The underlying mechanism to curb the development of the virus should be the reduction in the number of contacts between people who can be infected and those who are currently infected.

²Taiwan Centers for Disease Control (CDC) implemented inspection measures for inbound flights from Wuhan, China in response to reports of an unidentified outbreak. – 31st of December 2019.

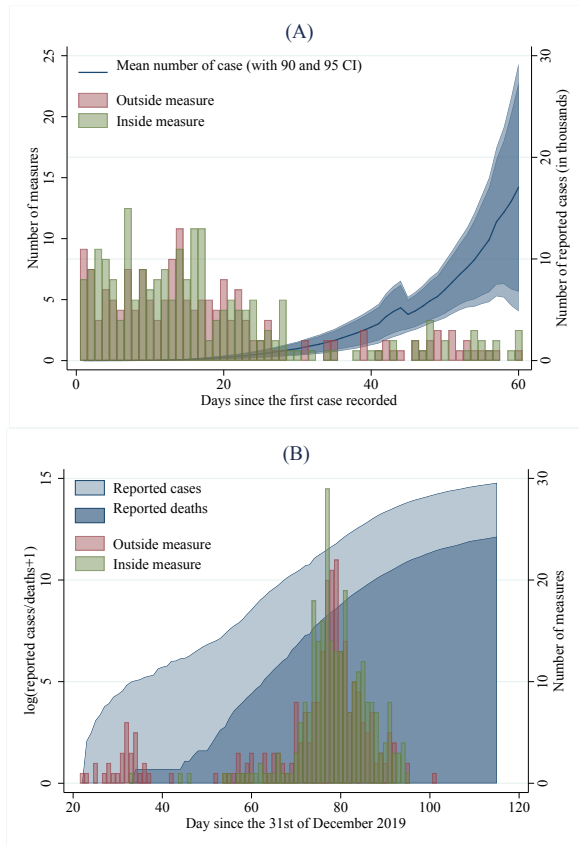


Figure 2: Evolution of Measures, Cases, and Deaths

Note: This Figure displays the distribution of lockdown measures over the beginning of 2020 and the beginning of the outbreak in each country. We exploit this visible variability to quantify the effect of each measure on the growth rate of the virus. "Outside" measures are those that restrict movements out of or into the country, while "Inside" measures are those restricting movements within a country. Both graphs exclude China. (A) Lockdown measures restricting movements within countries or towards the outside take place mostly during the 30 days after the first case is reported in the country, while some measures are taken up to 60 days after the first case. The blue line represents the mean number of reported cases by countries with 90% and 95% confidence intervals. (B) The earliest measures were taken in January with restriction of travel to or from specific locations (outside measure), while most of the measures were taken in March (from day 60 to day 91). The blue areas represent the number of reported death and number of reported cases for the world in log.

Successful lockdown measures are expected to restrict the movements of both the susceptible and the infected Kermack and McKendrick (1927); Maier and Brockmann (2020); Tian et al. (2020).

The panel structure of our data allows us to control quite extensively for the risk of omitted variable bias. First, the countries fixed effects allow to control for unobservables fixed over time at the country level (quality of the healthcare system, age distribution of the population, population density, geographical location, number of neighbouring countries, climate conditions, etc.). Those factors vary over time, but we could expect that they do not significantly vary over the period of interests (a few months). Second, the days fixed effects control for time-varying unobservables affecting the world in the same way (global evolution of the virus (early-stage vs. pandemic), global lockdown, etc.). Finally, the fixed effects also address the measurement errors by controlling for numerous factors that could correlates with the quality of the reporting and the spread of the virus. The countries fixed effects allow to exploit within country variation: if some policies or unobserved country characteristics affect the rate of case reporting (constant bias over time), this does not affect the within-country variation that we exploit.

The second main difficulty to measure the effect of governmental measures on the evolution of the disease comes from reverse causality. Indeed, the spread of the disease in the country influences the timing and the extent of the lock-down measures enforced by the government. To address the timing issue, we either control for the number of days since the first case was reported in country i or we control for the lagged dependent variable (auto-regressive model of order 1). Furthermore, the country fixed effects also serve to address the potential reverse causation of the extent of the measure taken (partial vs. complete lockdown, within vs. outside oriented measures). For example, a country who suffered from several initial “starting points” might require a complete lockdown compared to a country where the initial infections are all geographically concentrated (partial lockdown might be more appropriate).




Figure 3: Empirical Illustration

Notes: This graph reports the average growth rate of confirmed cases in the interval 30 days before, and after an inside measure was taken. The graph also shows a prediction of the growth rate based on fitting a linear model to the data before the measure was introduced.

One crucial empirical challenge is to find an adequate specification to capture the development of the growth rate of cases. Figure 3 reports the average growth rate of confirmed cases in the interval 30 days before, and after an inside measure was taken. Before the measure is introduced, the growth rate of cases is high and this eventually leads to its adoption. There is a sharp decrease in the growth rate after the measure has been implemented. The graph also shows a prediction of the growth rate based on fitting a linear model to the data before the measure was introduced. This is an illustration of how the growth rate of cases might have developed in the absence of the measure.

3.1 Baseline model: Number of days after the measure was taken

This approach allows to assess the change of trend after the country took the measure.

Model (1): First difference:

$$\begin{aligned} \log(\text{cases} + 1)_{it} - \log(\text{cases} + 1)_{i(t-1)} = \\ \beta_0 + \beta_1 \text{Measure}_{it} + \beta_2 \text{DaysAfterMeasure}_{it} + \beta_3 \text{First}_{it} + FE_i + FE_t + \epsilon_{ct} \end{aligned} \quad (1)$$

with i for country and t for the day. cases_{it} is the total number of people who were infected by the virus in country i on or before calendar day t . Measure_{it} is an indicator variable taking the value 1 from the day the measure was taken (onset). $\text{DaysAfterMeasure}_{it}$ is the number of days since the measure was taken. First_{it} records the number of days since the first confirmed case in country i at calendar time t . FE_i and FE_t are countries and days fixed effects. ϵ_{ct} is a error term clustered on the country level.

Model (2): AR(1) (auto-regressive model of order 1):

$$\begin{aligned} \log(\text{cases} + 1)_{it} = \beta_0 + \beta_1 \text{Measure}_{it} + \beta_2 \text{DaysAfterMeasure}_{it} \\ + \beta_3 \log(\text{cases} + 1)_{i(t-1)} + FE_i + FE_t + \epsilon_{ct} \end{aligned} \quad (2)$$

Model (3) is identical as Model (1) but we use an AR(1) instead of a first difference.

3.2 Baseline model: Time trend interaction

This approach allows to assess the global change of trend when a measure is taken

Model (3): First difference:

$$\begin{aligned} \log(\text{cases} + 1)_{it} - \log(\text{cases} + 1)_{i(t-1)} = \\ \beta_0 + \beta_1 \text{Measure}_{it} + \beta_2 \text{Days}_t \times \text{Measure}_{it} + \beta_3 \text{First}_{it} + FE_i + FE_t + \epsilon_{ct} \end{aligned} \quad (3)$$

with i for country and t for the day. Measure_{it} is an indicator variable taking the value 1 from the day the measure was taken (onset). Days_t is the number of days since the 31st of December

2019 (beginning of the sample). $First_{it}$ the number of days since the first case was reported in the country. FE_i and FE_t are country and day fixed effects. ϵ_{ct} is a error term clustered on the country level.

The parameter β_1 estimates the growth rate on calendar day 0, which is 31st of December 2019. The parameter β_2 estimates the change in the growth rate as a function of the number of days since day 0.

Model (4): AR(1) (auto-regressive model of order 1):

$$\begin{aligned} \log(\text{cases} + 1)_{it} = & \beta_0 + \beta_1 \text{Measure}_{it} + \beta_2 \text{Days}_t \times \text{Measure}_{it} \\ & + \beta_3 \log(\text{cases} + 1)_{i(t-1)} + FE_i + FE_t + \epsilon_{ct} \end{aligned} \quad (4)$$

Model (4) is identical as Model (3) but instead we use an AR(1) instead of a first difference.

3.3 Parallel with SIR model

Our estimates can also be interpreted in the context of the Susceptible-Infected-Recovered (SIR) model Kermack and McKendrick (1927). Individuals are either susceptible to the infection, S_τ , or infected, I_τ , so there can be at most $S_\tau \times I_\tau$ potential contacts between infected and susceptible (the SIR model assumes that recovered individuals play no direct role in new infections). The disease is then transmitted at rate β_τ from the infected to the susceptible individuals, so every period τ there are $\beta_\tau S_\tau I_\tau$ new cases reported infected. The total number of cases until day t is $\sum_{\tau=0}^t \beta_\tau S_\tau I_\tau$, and the growth rate of cases is equal to $\beta_{t+1} S_{t+1} I_{t+1}$. Our model provides an estimate of how this growth rate changes as measures are introduced. These changes happen for mainly two reasons. The transmission rate β_τ can decline because the number of actual contacts decreases, and the number of infected individuals decreases thereby creating fewer potential contacts. Our estimates provide the overall effect.

4 Results

4.1 Baseline model: Effectiveness of lockdown measures

We start by presenting how government measures reduce the growth of infections as a function of the time since the measure has been implemented compared to countries which have not implemented any measure yet. Panel (A) and (B) of Figure 4 shows the marginal effects of our baseline model (see Supplementary Material for the equation estimated and the regression tables). Lockdowns are implemented when confirmed cases increase strongly and affect infections only with a delay since the incubation period of the illness is several days. Restrictions within the country are more efficient than measures towards the outside at curbing the spread of the virus (the effect kicks-in quickly and triggers a steeper reduction). Panel (A) and (B) of Figure 4 highlight this results. On average, after 25 days, countries who took internal measures experienced a reduction of the growth rate compared to the other countries. After fifty-days the growth rate is lowered by 7.5%. On the other hand, the aggregation of measures towards the outside does not have a statistically significant effect after fifty-days. We aggregated the measure in two categories to highlight this main results. When we look at the subcategories of governmental measures defined in Figure 1 we obtain a similar split between within country measures and measures towards the outside (See Figure 1 panels (C) to (H) and (III) to (VIII)).

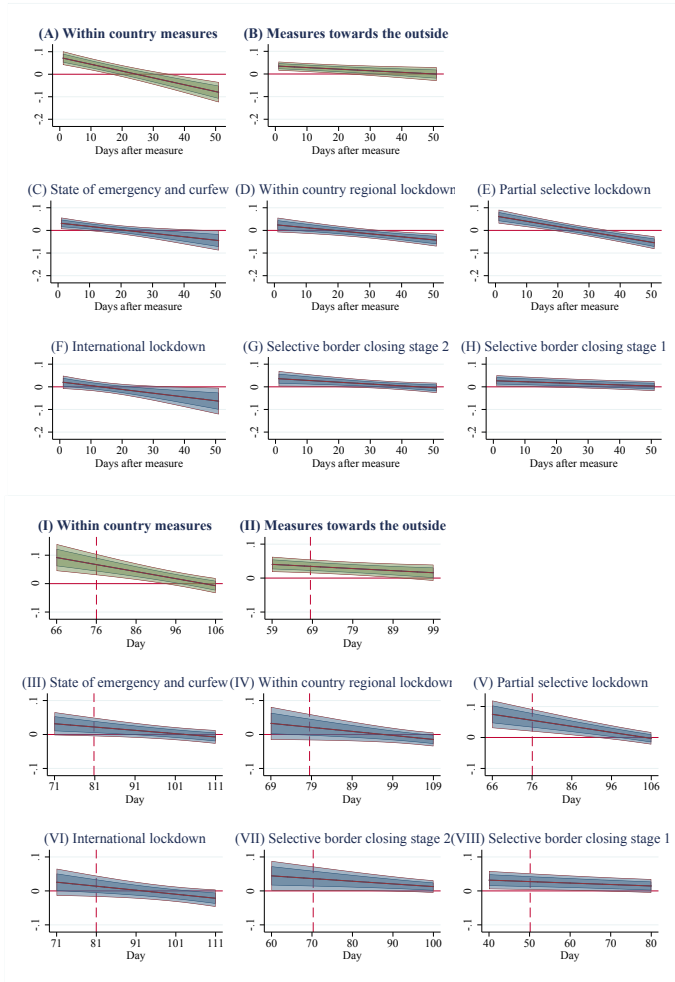


Figure 4: Baseline model (days after the measure): Marginal effects (Cases Confirmed)

Note: Marginal effects computed with our autoregressive model of order 1. 90% and 99% confidence intervals are shown in different shade of blue or green. The vertical dashed line shows the average day where the measure was taken in the sample. The model shows: i) the effectiveness of numerous lockdown measures that governments implemented across countries to mitigate the viral dissemination (statistically significant effect and number of days before the rate of the disease is reduced compared to countries who did not implement the measure), ii) the strength of the effect (steepness of the slope). The corresponding results for deaths are in the Supplementary Material. Panel (A) to (H) show the impact of a measure on the growth rate of infections as a function of time since the measure was implemented. Panel (I) to (VIII) show the impact of a measure on the growth rate of infections as a function of time since the 31st of December 2019 (Day).

Panel (I) and (II) of Figure 4 shows when each type of lockdown measure was adopted on average since December 31st 2019, and when the various measures became effective. Overall, this second approach provides a similar picture: within country measures have a clear impact, while the efficiency of outside measures remain questionable. Restrictions inside countries have been implemented on average on the 16th of March 2020 (76th day) and the average reduction was expected to be observed around the 13th of April (day 103). On the other hand, outside-oriented measures were taken on average on the 9th of March 2020 and their efficiency still had not materialized on the 13th of April (day 100).

Our baseline model thus strongly suggests that lockdown measures focused on blocking relationships among people within a country (inside measures) prevail over measures aimed at blocking international relationships. To explore this point in more depth, we also estimated a model including both measures: inside and outside (c.f: Supplementary Material). With this model, we can observe the effect of one type of measure while taking into account the effect of the other. This model weakens even more the evidence that outside measures had an effect. Results for the fatality growth rate point in the same direction, even though lockdowns measures took more time to have an impact. As discussed earlier, this delay was expectable (See Supplementary Material). We use estimates for deaths to quantify the number of prevented deaths. We find that, world-wide, internal measures have prevented about 650,000 deaths, this is more than three times the actual number of deaths. Internal measures have thus been successful at preventing many pre-mature deaths.

4.2 Quantifying Prevented Deaths

We use model (2) to compare the evolution of the total number of deaths with and without a measure. The model has two parameters which help assess this, β_1 which indicates by how much more the number of deaths grows in a country that has implemented a measure, and β_2

which describes the gradual slowing down of the growth rate in deaths due to the measure.

We base our simulation on countries that have implemented inside measures, as those are shown to be effective. We consider the average time, T , from the day when a measure has been implemented, 0, until the end of our analysis period. For countries that implemented the measure, the increase in the number of deaths between the day they implemented the measure until the end of the observation period is:

$$g_1 = \prod_{t=0}^T \exp(\hat{\beta}_1 + \hat{\beta}_2 \times t)$$

where \prod is the product of its arguments. The counterfactual growth in the number of deaths is

$$g_0 = \prod_{t=0}^T \exp(\hat{\beta}_1) = \exp(\hat{\beta}_1 \times T)$$

The ratio of $(g_0 - g_1)/g_1$ provides information on how many deaths were prevented per actual death that occurred. In our context, this ratio is 3.11 so somewhat more than three deaths were prevented per every death that unfortunately occurred. We then use the average number of cases in countries that implemented the measure, $\bar{d} = 209'799$, to calculate the total number of prevented deaths, which is $\bar{d} * (g_0 - g_1)/g_1 = 652'254$. A total of over 650,000 deaths were prevented, or a bit more than three prevented deaths per actual death.

4.3 Did early lockdown movers fare better?

In this section, we explore whether early reactions by governments influenced the spread of the Covid-19. We define an early reaction with respect to the calendar date when a measure is implemented, and define early to be in the first quartile of the countries implementing the measure (See Supplementary Material.) Figure 5 shows the marginal effects for the impact of moving early, and provide a consistent picture: the growth rate number of days to observe a

reduction of the growth rate is similar for countries that adopted an outside measure and reacted early compared to others, and the slope for early movers is flatter (lower intercept). Note, however, that reaching the zero growth rate at the same moment but at a lower rate implies that the rate was lower to start with, which is in line with the famous idea of “flattening the curve” and thus with the overall objective assigned to lockdown policies. We focus our analysis on inside measure as they proved to be more efficient throughout our analysis. Panels (A) and (B) of Figure 5 show that the countries adopting the inside measure later reached the baseline growth after 36.3 days for the countries which did not react early, while early movers reached the baseline growth rate in 23.4 days. Panel (E) of Figure 5 show that countries who took inside measure late did so on average on the 18th of March 2020 (day 78) and could expect the growth rate to slow down around mid-April. Panel (F) of Figure 5 show that countries who took inside measure late did so on average on the 6th of March 2020 (day 68) and could expect the growth rate to slow down around the end of March.

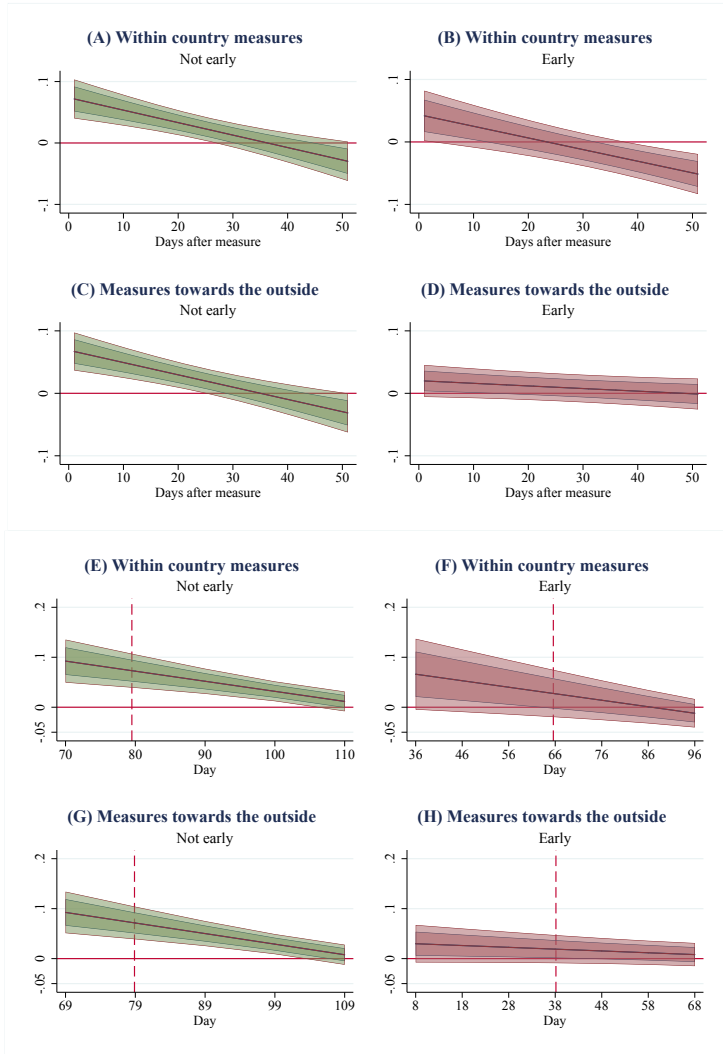


Figure 5: Sequence model: Marginal effects (Cases Confirmed)

Note: “Early” is defined as a measure adopted in the first quartile of the sample. Marginal effects computed with our autoregressive model of order 1. Panel (A) to (B) show the impact of a measure on the growth rate of infections as a function of time since the measure was implemented. Panel (C) to (D) show the impact of a measure on the growth rate of infections as a function of time since the 31st of December 2019 (Day). 90% and 99% confidence intervals are shown in different shade of blue. The vertical dashed line shows the average day where the measure was taken in the sample. The corresponding results for deaths are in the Supplementary Material.

4.4 Developing versus developed countries

This section explores whether the impact of lockdowns is different in developed and developing countries. Figure 6 shows the marginal effect of all the different types of measures for developed and developing countries³. A clear pattern appears in this Figure: lockdown measures have no statistically significant effects in developing economies, while the effects for developed economies are statistically significant. Most of the explanatory variation from our baseline model therefore comes from lockdown imposed in developed countries. Obviously comparing those results to the marginal effects of the baseline model, they are stronger as we are pinning down the group who benefit the most from the lockdown measures. For developed countries, within countries lockdowns have an effect after 20 days on average and the reduction after 50 days is 7.8% on average.

³We define developing countries as the ones with an Human Development index up to 0.699, which refers to Low and Medium human development using the United Nation codebook definition while above 0.699 will be defined as developed countries.




Figure 6: Developing versus developed countries model: Marginal effects (Cases Confirmed)

Note: Developing countries are the one with an Human Development index up to 0.699 which refers to Low and Medium human development using the United Nation codebook definition while above 0.699 will be defined as developed countries. Marginal effects computed with our autoregressive model or order 1. Panel (A) to (B) show the impact of a measure on the growth rate of infections as a function of time since the measure was implemented. Panel (I) to (I) show the impact of a measure on the growth rate of infections as a function of time since the 31st of December 2019 (Day). 90% and 99% confidence intervals are shown in different shade of green/blue. The vertical dashed line shows the average day where the measure was taken in the sample.

5 Discussion

Studying the lockdown measures adopted in the context of the Covid-19 crisis in 184 countries, our paper delivers several important public policy insights for how pandemics should be faced. From these insights, one can also derive ideas about how individuals behave during lockdowns and thus how pandemics can be faced.

The first key insight from our study is that lockdowns are indeed effective measures to stop both the growth of new cases and of the number of deaths. This result is in line with observations from previous pandemics. In his review of the evidence about the 1918 Influenza, Garrett (2008), for instance, compares the cases of Philadelphia, where public officials let a large parade take place during the pandemics, and St. Louis. He wrote: “Officials in St. Louis (a comparable city to Philadelphia at the time), however, responded quickly to the influenza by closing nearly all public places as soon as the influenza had reached the city. As a result, influenza mortality rates were much lower than in Philadelphia” Garrett (2008). With the covid-19 episode so far, lockdown measures have prevented many deaths -our estimates are that about 650,000 deaths have been averted- or more than three deaths were prevented for every death that occurred.

Contrary to common belief, however, our analysis suggests that the most extreme measures such as total lockdowns and immediate border closures are not necessarily the most effective actions to respond to a pandemic, even without considering the economic impact of these lockdowns. Let’s analyse these in turn. First, our empirics show that partial or regional lockdowns are as effective as stricter measures such as those related to declaring a state of emergency or implementing curfews. Since partial measures are likely to be less damaging to the economy than stricter lockdowns, their overall impact can be considered as superior. This analysis should of course be confirmed by a joint study of the economic and health impact of the virus, but the fact that partial internal measures are effective at stopping the spread of the disease and at push-

ing down mortality is an important result by itself. So, why is this the case? One possible explanation is that partial and selective lockdowns are enough to push down the opportunity costs for people of going outside –as schools, stores or local businesses are closed- and taking the risk of being infected. Total lockdowns might thus be superfluous. In a similar manner, one could speculate that partial lockdowns might be strong enough as signals for people not only to stay home but also to quickly adopt sanitary measures or avoid group activities that might spread the disease fast. In other words, our results point to the fact that people might adjust their behaviors quite significantly as partial measures are implemented, which might be enough to stop the spread of the virus at lower economic costs. This questions pure epidemiological models, which typically made projections about the diffusion of the covid-19 without taking into account the adjustments made by rational individuals.

The third striking result of our analysis is that taking inside-country measures matters much more than implementing outside-oriented ones. Blocking borders, in particular, is the least effective policy at curbing the development of the virus, unless it follows effective internal measures. Even in a globalized world, internal policies are the name of the game. This result is in sharp contrast to current political discussions in the US and elsewhere, which often focus on border closure instead of putting the emphasis on within-country lockdowns. Again, why is this? One interpretation, in line with what was discussed above, could be that internal measures are effective at reducing opportunity costs for people of going out during a partial lockdown, whereas outside measures do not have this effect. Here again, what might drive the success of lockdown measures might be their ability to trigger a strong adjustment in individuals' behaviors. Whereas internal measures might have a significant effect, for instance, on the opportunity cost of staying home, it is likely that outside-oriented measures do not change much on that front for many individuals. This reasoning might also explain why outside measures matter only once internal ones have been implemented, an a result we obtained in a post hoc analysis

available from the authors upon request. Outside-oriented measures might thus deliver some added benefits in terms of further limiting interactions, but only when individuals have already adjusted their daily behaviours.

In order to push our idea of opportunity cost further, we split our sample and explored differences between developed and developing countries. Our working hypothesis there was that the opportunity costs of abiding to lockdown rules and staying home are much higher in developing economies in which many people make a living in the informal sector and do not have any safety net. In agreement with our hypothesis, we do find that internal lockdown policies have a significant effect on both the number of cases and on the number of fatalities, whereas this is not the case in developing countries. We cannot firmly conclude from our analysis that lockdowns are not effective in developing countries, as the disease in these countries appeared later and we might thus lack observations and statistical power. However, our results so far indicate that lockdown measures would have to be coupled with other policies, which could push opportunity costs down, to really impact the spread of the disease in developing markets.

Last, our empirical results suggest that there is somewhat of a speed premium for policy-making in the context of a pandemic, especially regarding the objective of 'flattening of curve' to avoid overwhelming intensive care hospital facilities.

In sum, and despite the fact that extreme measures have often been taken by countries in panic situations and for emergency purposes, there are clear learning outcomes from this first large pandemic of modern times: developing organizational structures and decision-making processes favouring fast reaction, agility and targeted lockdowns should be priorities. For similar reasons, these features should help in case we enter into a 'lockdown-release-lockdown' era, a hypothesis that cannot be ruled out in early May 2020 with the apparently low prevalence rate of the coronavirus across countries. One obvious caveat of our study, in that respect, is that the long-term efficiency of lockdown measures will only be known when these lockdowns have

been lifted and when we have had time to observe whether the coronavirus has not surged again Bonardi et al.. If we are right that one key aspect of internal lockdown measures is to have pushed individuals to adjust their daily behaviors, there might hope in that regards nonetheless.

References

- Avery, C., Bossert, W., Clark, A., Ellison, G., and Ellison, S. F. (2020). Policy implications of models of the spread of coronavirus: Perspectives and opportunities for economists. Working Paper 27007, National Bureau of Economic Research.
- Bonardi, J.-P., Bris, A., Brühlhart, M., Danthine, J.-P., Jondeau, E., Rohner, D., and Thoenig, M. The case for reopening economies by sector. Section: Economy.
- Chinazzi, M., Davis, J. T., Ajelli, M., Gioannini, C., Litvinova, M., Merler, S., Pastore y Piontti, A., Mu, K., Rossi, L., Sun, K., Viboud, C., Xiong, X., Yu, H., Halloran, M. E., Longini, I. M., and Vespignani, A. (2020). The effect of travel restrictions on the spread of the 2019 novel coronavirus (covid-19) outbreak. *Science*.
- Dong, E., Du, H., and Gardner, L. (2020). An interactive web-based dashboard to track covid-19 in real time. *The Lancet infectious diseases*.
- Flaxman, S., Mishra, S., Gandy, A., Unwin, H., Coupland, H., Mellan, T., Zhu, H., Berah, T., Eaton, J., Perez Guzman, P., et al. (2020). Report 13: Estimating the number of infections and the impact of non-pharmaceutical interventions on covid-19 in 11 european countries.
- Garrett, T. A. (2008). Pandemic economics: The 1918 influenza and its modern-day implications. *Federal Reserve Bank of St. Louis Review*, 90(March/April 2008).
- Giordano, G., Blanchini, F., Bruno, R., Colaneri, P., Di Filippo, A., Di Matteo, A., and Colaneri, M. (2020). Modelling the covid-19 epidemic and implementation of population-wide interventions in italy. *Nature Medicine*, pages 1–6.
- Harris, J. E. (2020). The coronavirus epidemic curve is already flattening in new york city. Technical report, National Bureau of Economic Research.
- Hartl, T., Wälde, K., and Weber, E. (2020). Measuring the impact of the german public shutdown on the spread of covid19. *Covid economics, Vetted and real-time papers, CEPR press*, 1:25–32.
- Hatchett, R. J., Mecher, C. E., and Lipsitch, M. (2007). Public health interventions and epidemic intensity during the 1918 influenza pandemic. *Proceedings of the National Academy of Sciences*, 104(18):7582–7587.
- Kermack, W. O. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721.
- Kraemer, M. U., Yang, C.-H., Gutierrez, B., Wu, C.-H., Klein, B., Pigott, D. M., du Plessis, L., Faria, N. R., Li, R., Hanage, W. P., et al. (2020). The effect of human mobility and control measures on the covid-19 epidemic in china. *Science*.
- Maier, B. F. and Brockmann, D. (2020). Effective containment explains subexponential growth in recent confirmed covid-19 cases in china. *Science*.
- Programme, U. N. D. (2020). <http://hdr.undp.org>. *Human Development Index (HDI)*. Available On-Line: <http://hdr.undp.org/en/content/human-development-index-hdi>.
- Tian, H., Liu, Y., Li, Y., Wu, C.-H., Chen, B., Kraemer, M. U., Li, B., Cai, J., Xu, B., Yang, Q., et al. (2020). An investigation of transmission control measures during the first 50 days of the covid-19 epidemic in china. *Science*.

Supplementary materials

Data and main variables

Empirical strategy

Additional figures S1 to S9

Regression tables S1 to S38

References (1 - 2)