CENTRE FOR
ECONOMIC
POLICY
RESEARCH

CEPR PRESS

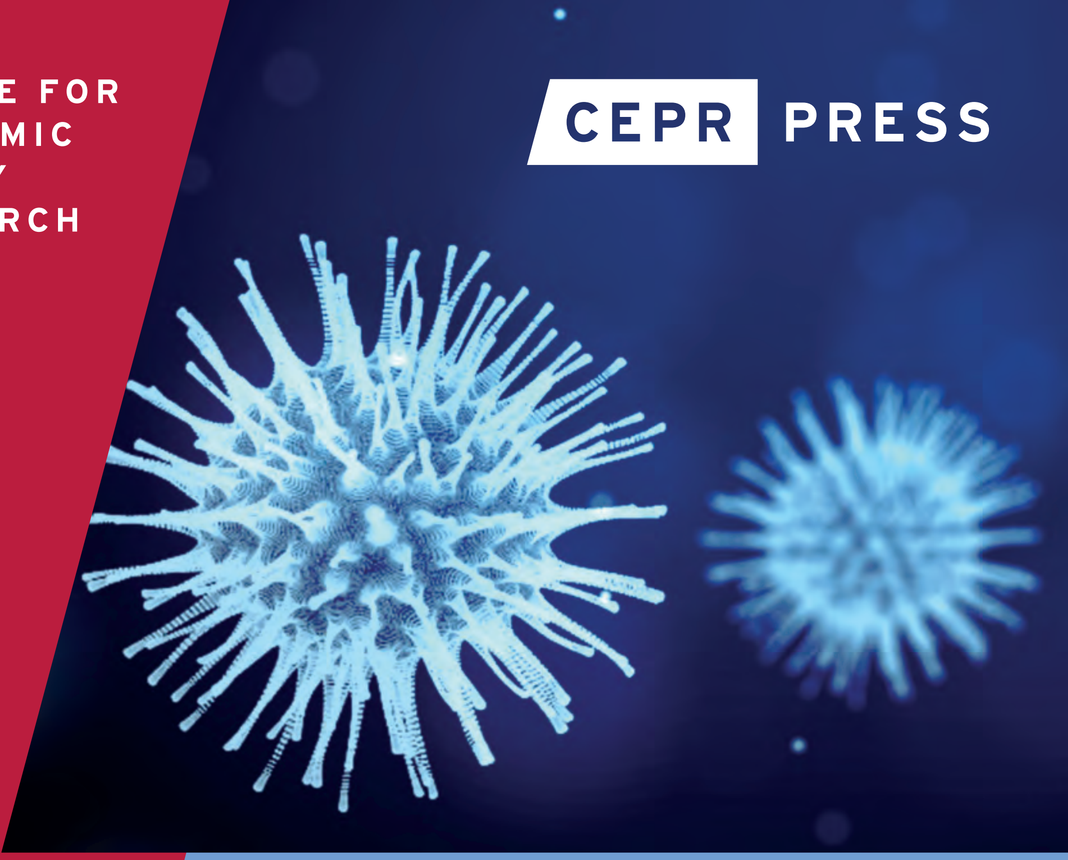

COVID ECONOMICS

VETTED AND REAL-TIME PAPERS

ISSUE 52
15 OCTOBER 2020

Covid Economics
Vetted and Real-Time Papers

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

Founder: Beatrice Weder di Mauro, President of CEPR
Editor: Charles Wyplosz, Graduate Institute Geneva and CEPR

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

The Centre for Economic Policy Research (CEPR)

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

Editorial Board

Ethics

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

Submission to professional journals

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

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

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

Covid Economics
Vetted and Real-Time Papers

Issue 52, 15 October 2020

Contents

CEPR PRESS

Covid Economics **Issue 52, 15 October 2020**

Uncertainty and lockdown in COVID-19: An incomplete information SIR model

Oliver Forsyth[1]

Date submitted: 13 October 2020; Date accepted: 14 October 2020

This paper introduces an information structure into the standard SIR model to investigate the role of targeted lockdown policies in the presence of incomplete information. By allowing for asymptomatic infected agents and symptomatic susceptible agents, such that the presence of a symptom is an imperfect indicator of an agent's health state, we solve for the optimal lockdown policies on symptomatic and non-symptomatic agents. The model is then calibrated to the UK, where we find mitigation measures have reduced the peak of the infection rate from 23.9% to 3.47%, and decreased the number of fatalities by 39.1% if a vaccine is discovered within 18 months. If a uniform lockdown policy is pursued, the costs to the economy are large with GDP falling by 18.3% in the first year. However, through conditioning lockdown policies on the presence of symptoms, these costs may be substantially reduced, with no increase in the number of fatalities.

Covid Economics 52, 15 October 2020: 1-38

1 Graduate Student, Faculty of Economics, University of Cambridge.

CEPR PRESS

1 Introduction

COVID-19 has brought the world to a halt since the first, publicised outbreak in Wuhan, China. The virus spread rapidly over the globe and was declared a pandemic by the WHO on March 11th (Ghebreyesus 2020). To guard against their healthcare services being overwhelmed, countries have engaged in costly mitigation measures, shutting down places of work and schools, suspending international travel and severely restricting social activity with government mandated lockdowns. Following Stock (2020) and Atkeson (2020), who drew attention to the need for both a greater understanding of the effect of these measures on the economy and also for the establishment of the most efficient means to combat the virus, there has been a flurry of responses from economists. Through imposing additional structure on popular disease models, lockdown policies, simple economies and optimised agent behaviour may be readily introduced. Central to the approach is the assumption that the human interactions which facilitate the spreading of the virus are also necessary for production in the economy. This introduces a trade off between output and transmission, making the virus costly to mitigate. For example, the UK, having applied extensive lockdowns across the population saw an immediate 20.4% decline in GDP in the first quarter since mitigation measures started, according to data from the ONS.

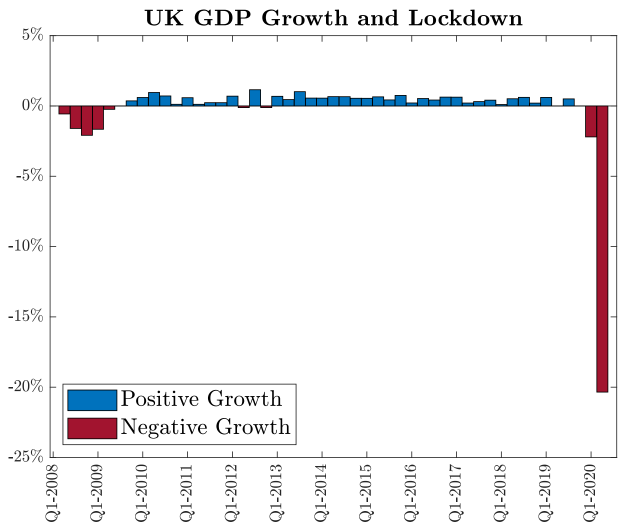

Figure 1

Covid Economics 52, 15 October 2020: 1-38

CEPR PRESS

Covid Economics 52, 15 October 2020: 1-38

A lockdown policy works to limit the number of interactions in the population, thereby limiting the transmission rate. Under the SIR model, since infection may only occur when a susceptible agent meets an infected agent, it is possible to eradicate the virus through the immediate identification and isolation of all infected individuals. This would cause only minor costs to the economy, requiring just a small number of infected agents to stop production. However, there are two sources of information friction which disrupt the identification of infected agents. Firstly, there are infected agents who are asymptomatic, such that even in the absence of symptoms an agent may still be transmitting the disease. Secondly, in the instance an agent does develop symptoms, in the majority of cases, these symptoms are indeterminate and indistinguishable from the common cold or flu. For example, the CDC, the US public health institute, lists over 11 common symptoms for COVID-19, ranging from headaches to fevers. Of these, 8 are also associated with the common cold or flu. These two sources of friction make it impossible to impose lockdown measures that only impact infected agents, and force the consideration of policies that will also lockdown susceptible agents. A complete understanding of the information available to a planner is therefore a key consideration to determine the optimal lockdown policies.

However, the literature has struggled to incorporate these two information frictions into the SIR model. Either it is assumed that there is complete information failure, where the policy maker is unable to distinguish susceptible and infected agents and subsequently applies a uniform lockdown across the entire population (Alvarez et al. 2020). Alternatively, agents are assumed to be fully identified, representing the case of perfect information. To address these problems our model significantly develops the information structure in the SIR model and permits lockdown policies to be conditional on the presence of symptoms. The model allows for the possibility of not only *asymptomatic infected* agents, but also *symptomatic susceptible* agents, such that the presence, or lack of, a symptom only partially reveals an agent's true health state. To achieve this, each period both susceptible and infected agents have a probability λ and $1-\alpha$ respectively of developing a symptom, where α is the asymptomatic rate in the population. To ensure that the signal is informative, the symptom probabilities are restricted such that $\lambda \leq 1-\alpha$, imposing that the probability an agent is infected is non-decreasing in symptom presence. A policy maker must then apply a lockdown policy to each agent type, specifically those with and without symptoms, based on their belief the agent is infected.

To study the impact of the information structure explicitly, we propose that a distinction is drawn between a mitigation path and lockdown allocations. The mitigation path, X, dictates the measures taken by a government to restrict the virus, and completely determines the path of the virus in the economy, including the number of fatalities. The lockdown allocations are the pairs of lockdown policies applied to non-symptomatic and symptomatic agents each period to achieve the mitigation path. From the perspective of a social planner, the problem of the optimal mitigation path is dynamic and depends on a planner's preferences for output and fatalities, and the laws of motion for the virus under the SIR model. However, the problem of optimal lockdown allocation is static, described by the first order derivatives with respect to each lockdown policy. That is, for an exogenous mitigation path, X, there is a unique lockdown allocation for each period that minimises the impact on the economy, subject to maintaining the desired level of infection. This method is used to analyse the UK. The mitigation path is calibrated to estimates of the UK's effective reproduction number, before the model is used to solve for the optimal lockdown allocations each period under different information structures.

This is an alternative approach to solving the full dynamic system given by the social planner problem and offers several advantages. As shown in the robustness section, there is considerable uncertainty around many key epidemic parameters, with many of the virus dynamics, such as peak infection, duration until herd immunity and fatalities, being very sensitive to the values chosen. The focus on the static problem allows for information effects to be found that are robust to epidemic parameters. In addition, this approach also provides greater harmony with the epidemiology literature, where there are models considerably more sophisticated than the SIR model to predict the spread of a virus in the population (Ferguson et al. 2020). This approach allows the mitigation path to be calibrated to replicate the infection spread in these epidemiological models and then the model solved for the most efficient means to achieve this pathway and the resulting impacts to the economy. Finally, the social planner problem is awkward to implement, in that the solution intimately depends on the trade off of output against the value of a fatality. For this, there is no clear calibration target which leaves the solution very exposed to the value chosen by the researcher. Our approach does not need to address this for the number of fatalities is determined exogenously by the mitigation path and is independent of how lockdown is allocated.

Our results find that the UK has been following a suppression policy since the onset of the virus, introducing strong mitigation measures to reduce the virus reproduction number to one, where, we assume, it is held until a vaccine arrives or herd immunity is achieved. This policy has been very effective in reducing the severity of the virus with peak infection falling from 23.9% to 3.47%. Cumulative fatalities are reduced by 24.9% if a vaccine arrives within 2 years, saving approximately 24,000 lives per year in a UK adult population of 54 million. This rises to a 39.1% reduction in fatalities if the vaccine arrives within 18 months. The mitigation measures are needed for 19 months in the absence of a vaccine, at which point herd immunity is achieved. If the UK is restricted to following a uniform lockdown policy that does not condition on symptoms, all agents experience a hump-shaped lockdown where activity is restricted by 38.3% at the peak of mitigation. The damages to the economy under such a policy are vast with GDP falling by 18.3% in the first year.

Covid Economics 52, 15 October 2020: 1-38

CEPR PRESS

Covid Economics 52, 15 October 2020: 1-38

We derive several principles to optimally condition lockdown policies on the presence of a symptom, which, if followed, can substantially reduce the costs to the economy with no increase in the number of fatalities. Firstly, it is only optimal to apply a uniform lockdown across the population in the case of complete information failure, where both agent types are equally likely to develop symptoms. In all other instances, since the presence of a symptom may only increase the probability an agent is infected, it is always optimal to impose heavier lockdown measures on symptomatic agents. Finally, as information improves and a symptom better identifies infected agents, the lockdown allocations become increasingly targeted on infected agents. More stringent lockdown measures are imposed on those with symptoms and lighter lockdowns on those without symptoms.

It is the increased targeting of infected agents that leads to substantial reductions in lost GDP as information improves. In the case that $\lambda = 0$ and $\alpha = 0.2$, such that only infected agents display symptoms, an aggressive lockdown of 73.4% is applied to symptomatic agents while non-symptomatic agents only experience a lockdown of 5.4%. This deeply targeted lockdown on symptomatic agents removes a sufficiently high number of infected agents that the majority of the population, who do not display symptoms, need only experience minor restrictions. Consequently, since less of the population is under lockdown, the loss in output is limited to only 2.34% in the first year of mitigation. It should be stressed that this substantial improvement on the uniform policy is achieved with no increase in the fatality rate and is entirely driven by better identification of infected agents and the ensuing ability of policy makers to more precisely extract infected agents.

We also find that lockdown allocations, and therefore the costs to the economy, are particularly sensitive to low values of λ. This is driven by the fact that the number of susceptible agents is substantially larger than the number of infected agents throughout the course of the epidemic. An increase in λ, therefore, introduces a large number of susceptible agents into the symptomatic set and quickly erodes a planner's ability to discern infected agents. This effect is diminishing as λ increases. However, it does imply that accurate assessment of the information structure is important. It also suggests that there should be substantial seasonal variation in lockdown allocations in response to the increased number of symptomatic susceptible agents from the flu and common cold over the winter.

The paper is organised as follows. Section 2 briefly surveys the related literature and Section 3 describes pertinent details of the standard SIR model. Section 4 introduces the model, establishing the information environment, the laws of motion for the health states, S, I and R, as well as detailing the economy. Section 5 solves the social planner problem and provides a discussion of the static problem. In Section 6 the model is applied to the UK with Section 7 discussing parameter choices. Section 8 presents the numerical results, and Section 9 discusses a robustness exercise. Section 10 concludes.

CEPR PRESS

Covid Economics 52, 15 October 2020: 1-38

2 Related Literature

There has been an enormous expansion in the economics literature. The classic disease models, such as the SIR model, have been modified to allow contact rates to depend on interpretable lockdown policies. These lockdown policies in turn influence the economy, creating a framework to analyse the costs from different mitigation paths. Alvarez et al. (2020), Farboodi et al. (2020) and Jones et al. (2020) were early examples of such an approach. Through solving the social planner problem they showed that the optimal lockdown allocation is humped-shaped, and sharply front-loaded to contain the spread of the disease in its infancy. At the core of this approach is the trade off between short-term output and fatalities, introducing the notion that there is an unavoidable frontier by which the planner is constrained.

An alternative approach is to solve for the decentralised equilibrium to investigate the precautionary measures taken by agents in the absence of government policy. Toxvaerd (2020) and Garibaldi et al. (2020) demonstrate that agents do take significant precautionary measures, increasing social distancing efforts as the virus prevalence increases. However, unlike in the planner solution, since the risk of spreading the infection is not considered by private agents, infected agents engage in no social distancing, leading to under provision of lockdown in the population. Eichenbaum et al. (2020) show that this externality may be addressed through a tax on consumption. This is an important area of research as it suggests that using the standard SIR model as a counterfactual for government intervention will overstate the significance of lockdown policies.

As the understanding of the virus has improved, attention has shifted to investigating to what extent targeted lockdown policies, conditioned on heterogeneity in the population, could shift the output-fatality frontier. Acemoglu et al. (2020), Bairoliya & Imrohoroglu (2020), Favero et al. (2020) and Rampini (2020) all consider the impact of age targeted policies, finding that an aggressive quarantine on the elderly could substantially reduce the economic burden of lockdown with no increase in the number of fatalities. In an interesting related branch, Greenstone & Nigam (2020) considers how the costs of lockdown are disproportionately accrued by the younger, working population with Glover et al. (2020) investigating the impact of costly wealth transfers from old to young on the optimal lockdown policy.

There has been substantially less explicit modelling of the information available to a planner, specifically the planner's ability to discern the true health state of an agent. In light of evidence suggesting a high degree of asymptomatic transmission (Nishiura et al. 2020), a common approach has been to use the SEIR model, where the extra E state contains infectious agents who are yet to develop symptoms (Atkeson 2020, Berger et al. 2020, Piguillem et al. 2020). However, this approach captures only part of the information frictions present

Covid Economics 52, 15 October 2020: 1-38

to a planner, maintaining the assumption that symptomatic agents are perfectly identified by a planner and placed under immediate quarantine. Attempts to impair the identification of symptomatic infected agents are limited to Chari et al. (2020) who also introduce an imperfect signal of infection. However, this is to investigate a targeted testing policy on symptomatic agents and does not allow for lockdown policies to be conditioned on the presence of a symptom.

3 Standard SIR model

We begin with a brief introduction to the popular SIR model with constant population. First introduced in Kermack & McKendrick (1927) it is a compartmental model where individuals may be in one of 3 health states: Susceptible, Infected, and Recovered. All recovered agents are assumed to be immune from future infection. Virus transmission is assumed to be proportional to the number of interactions between susceptible and infected agents, modelled by a quadratic matching technology such that the number of new infections each period is given by $\gamma I_t S_t$, where γ may be interpreted as the probability of infection per meeting. There is also a flow from I to R, as δI_t agents recover each period. The epidemic dynamics are:

$$S_{t+1} = S_t - \gamma S_t I_t$$

$$I_{t+1} = I_t(1-\delta) + \gamma S_t I_t$$

$$R_{t+1} = R_t + \delta I_t$$

where $S_0 > 0$, $I_0 > 0$ and $S_t + I_t + R_t = 1 \; \forall t$.

The model while simple captures some important properties of contagion. Firstly, on account of the quadratic matching technology, recovered agents have no impact on the evolution of the virus allowing the system to be reduced to just two state variables. In addition, the virus reproduction number, R_0, defined as the average number of infections generated from a single infection in the economy in a fully susceptible population, is easily characterised as

$$R_0 = \frac{\gamma}{\delta}$$

while, the *effective* reproduction number, defined as the average number of infections from a single infection in a susceptible population S_t, is simply:

$$R_t = \frac{\gamma}{\delta} S_t = R_0 S_t$$

Covid Economics 52, 15 October 2020: 1-38

4 Model

The SIR model, while providing an adequate description of virus dynamics, does not provide the necessary structure to interpret the impact of specific mitigation measures. The model is developed to introduce simple lockdown policies that directly influence the transmission of the virus and output in the economy. This allows for an assessment of the costs incurred for different mitigation plans. In addition, the information structure of the model is significantly enhanced to allow for policy to be conditional on the presence of symptoms in an agent.

4.1 Information Structure

Asymptomatic Infected Agents

The first natural step is to allow for infected agents who do not display visible symptoms and are thus classified as asymptomatic. These agents are indistinguishable from susceptible agents to a policy maker. To introduce asymptomatic agents we allow that each period, there is a $1-\alpha$ probability that an infected agent shows visible symptoms, and an α probability they will be asymptomatic. To keep the model simple we assume that this probability is constant each period. This implies that an infected agent's probability of showing symptoms is independent over time and does not depend on whether they have previously displayed symptoms. Consequently, it is not necessary to track the symptom history of each individual, with the convenient aggregation property that the total number of symptomatic and asymptomatic infected agents each period is given by $(1-\alpha)I_t$ and αI_t respectively.

Symptomatic Susceptible Agents

A key innovation in this paper is to allow for the possibility of susceptible agents who display symptoms. As far as we are aware, this is a new addition and captures a vital, yet often missed, area of uncertainty for policy makers.

Mechanically, this is achieved in the same manner as the introduction of asymptomatic agents. Each period there is a probability, λ, that a susceptible agent develops symptoms, making them indistinguishable from an infected agent to a policy maker. Once again, this probability is constant over time, and independent of an agent's previous symptom history, such that the total number of symptomatic and non-symptomatic agents each period is given by λS_t and $(1-\lambda)S_t$ respectively.

While the presence of symptomatic susceptible agents can be treated simply

CEPR PRESS

Covid Economics 52, 15 October 2020: 1-38

as an instrument to introduce information frictions to the model, one can also consider the interpretation that there is a secondary disease introduced to the population. This secondary disease has a constant probability λ of infecting any susceptible agent each period, regardless of their previous symptom history or current lockdown state. Agents, once infected, will then display symptoms identical to the primary virus, in this case, COVID-19, for one period before recovering. This gives rise to λS_t symptomatic susceptible cases each period. No immunity may be gained against the secondary disease nor does prior or current contraction of the secondary disease grant any immunity to the agent against COVID-19. Agents are therefore equally likely to contract COVID-19 with or without the secondary disease. In addition, in the case that an agent contracts both COVID-19 and the secondary disease, COVID-19 is the dominant virus and an agent transitions to the infected state. This secondary disease is assumed costless to a carrier in the respect that it carries no risk of death or lower productivity.

Partitioned SIR Model

This developed information structure can be seen as simply partitioning the states of the SIR model, with λ controlling the share of S that display symptoms and $1-\alpha$ the share of I displaying symptoms.

For a policy maker, this equates to introducing three information sets an agent may fall into denoted by $J = \{NV, V, R\}$ where NV is the set of agents not displaying symptoms, V is the set of agents displaying symptoms, while R is the set of recovered agents. Figure 2 depicts these information sets. Immediately, it can be seen that, in the case that $\alpha = 0$ and $\lambda = 0$, these sets simplify back to the three states of the standard SIR model.

Since both susceptible and infected agents may develop symptoms, the observable health status of an agent is an imperfect signal as to the agent's true infection state. A policy maker, faced with these information frictions must therefore form beliefs to infer the true state of an agent. These are formed using Bayes rule such that:

$$P(I \mid j = V) = \frac{(1-\alpha)I}{(1-\alpha)I + \lambda S} \qquad\qquad P(I \mid j = NV) = \frac{\alpha I}{\alpha I + (1-\lambda)S}$$

Recovered agents are assumed to have perfect immunity, such that $P(I \mid j = R) = 0$. These probabilities reflect the policy maker's ability to discern an agent's infection state from observing their symptom state.

Covid Economics 52, 15 October 2020: 1-38

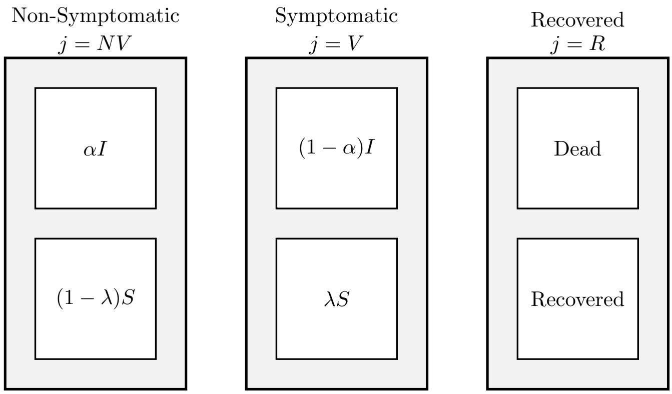

Figure 2: Three Information Sets

To maintain that the presence of a symptom is informative, it is necessary to impose that:

$$\lambda \leq 1 - \alpha \tag{4.1}$$

This restriction states that the incidence rate of symptoms in the susceptible population is never greater than the symptoms incidence rate in the infected population, thereby ensuring that a symptom is an informative signal. The restriction also leads naturally to Lemma 1.

Lemma 1. *If* $0 < \lambda < 1 - \alpha < 1$ *then* $P(I \mid V) > P(I \mid NV)$*. In the case that* $\lambda = 1 - \alpha$ *then* $P(I \mid V) = P(I \mid NV)$*.*

Proof.

$$P(I \mid V) - P(I \mid NV) = \frac{IS(1-\alpha-\lambda)}{((1-\alpha)I + \lambda S)(\alpha I + (1-\lambda)S)} \tag{4.2}$$

Since the denominator is strictly positive, a sufficient condition for (4.2) to be positive is $\lambda < 1 - \alpha$. □

Using the conditional probabilities of infection, the extremes of the information environment can be defined as:

- Perfect Information: $\lambda = \alpha = 0$

  $P(I \mid V) = 1$ and $P(I \mid NV) = 0$

In this environment, there are no asymptomatic infected agents or symptomatic susceptible agents and so a policy maker may fully discern an agent's infection state from their observed health state.

- Complete Information Failure: $\lambda = 1 - \alpha$

  $P(I \mid V) = P(I \mid NV) = \frac{I}{S+I}$

  The conditional probabilities collapse to the population probabilities. There is no additional information available to a planner from observing if an agent displays symptoms.

For intermediate cases, where $0 \leq \lambda < 1 - \alpha$, there is a flexible degree of information failure which is increasing in both λ and α. Figure 3 shows how a policy maker's ability to identify infected agents changes as λ increases from 0 through to 1-α.

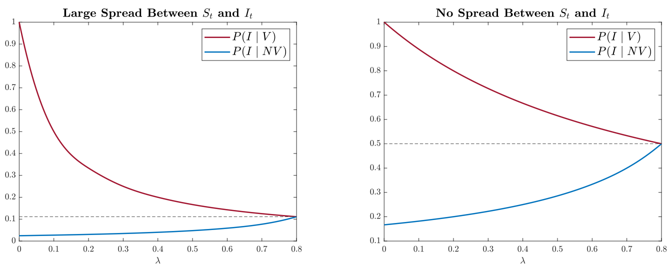

(a) $S = 0.8$, $I = 0.1$ and $\alpha = 0.2$ (b) $S = I = 0.4$ and $\alpha = 0.2$

Figure 3: Probability Convergence as λ Increase

The first observation is that the probabilities move in opposite directions. As λ increases, the set of symptom-showing agents is diluted with susceptible agents, reducing a policy maker's ability to identify infected agents by the presence of symptoms. Conversely, asymptomatic infected agents become a larger share of the non-symptomatic set, such that the lack of symptoms is increasingly insufficient a marker of a susceptible agent. As λ increases further through to $1 - \alpha$, there is convergence in these probabilities, representing the diminishing validity in using symptom status as an indicator of infection. This culminates in the case of complete information failure.

In addition, this rate of convergence depends heavily on the relative number of susceptible and infected agents in the economy. When the level of infection is low relative to the susceptible population, a small increase in λ results in a large increase in the number of susceptible agents displaying symptoms. This

Covid Economics 52, 15 October 2020: 1-38

quickly dilutes the concentration of infected agents, resulting in a steep decline in $P(I \mid V)$. To restate the result simply: provided the level of infection is low in the population, it only takes a small increase in λ to severely impact a planner's ability to discern infected agents by the presence of symptoms. This is a key dynamic in our model. In the case of the UK, I is kept substantially lower than S throughout the epidemic, resulting in lockdown policies that are particularly responsive to even low levels of λ.

4.2 Transmission and Conditional Lockdown Policies

Our second deviation from the standard SIR model is to introduce simple lockdown policies which may be conditioned on the presence of symptoms. Under the SIR model, the transmission of the virus is proportional to the level of interaction in the economy, $I_t S_t$. Each period a planner can implement lockdown policies to control the exposure level of each agent type, where $e_t^V \in [0, 1]$ denotes the exposure level of symptomatic agents and $e_t^{NV} \in [0, 1]$ denotes the exposure of non-symptomatic agents. A lockdown policy is then defined as $1 - e_t^j$ such that $e_t^j = 1$ signifies no lockdown is in place on type j agents. It is assumed that there is perfect adherence to the lockdown policy and that agents have zero interactions when in lockdown.

A lockdown policy can be seen to be restricting the supply of susceptible and infected agents such that the number of interactions each period in the population is reduced to:

$$I_t S_t((\alpha e_t^{NV} + (1-\alpha)e_t^V)((1-\lambda)e_t^{NV} + \lambda e_t^V) = E_t^I E_t^S I_t S_t$$

where

- $E_t^I = \alpha e_t^{NV} + (1-\alpha)e_t^V$ is the average exposure level of an infected agent
- $E_t^S = (1-\lambda)e_t^{NV} + \lambda e_t^V$ is the average exposure level of a susceptible agent.

Unlike in Alvarez et al. (2020) and Farboodi et al. (2020), the level of interaction each period is no longer reduced by the product of the two lockdown policies, but rather by the product of the average exposure levels for susceptible and infected agents, $E_t^I E_t^S$. By expanding this product the interaction level can be decomposed into a weighted sum of 4 types of meetings. Since all agents are equally susceptible and contagious, regardless of their symptom status, these weightings also represent the shares for the different sources of infection. Table 1 shows these sources.

Firstly, it can be seen that λ and α control the relative size of each of the transmission streams. For example, when $\lambda = \alpha = 0$, the model collapses into

Covid Economics 52, 15 October 2020: 1-38

Covid Economics 52, 15 October 2020: 1-38

Table 1: The Four Types of Transmission

		Infected	
		Symptomatic	Non-Symptomatic
Susceptible	Symptomatic	$\lambda(1-\alpha)e^V e^V$	$\alpha\lambda e^{NV} e^V$
	Non-Symptomatic	$(1-\alpha)(1-\lambda)e^{NV} e^V$	$\alpha(1-\lambda)e^{NV} e^{NV}$

the SIR model, with transmission only occurring through infected agents with symptoms and susceptible agents without symptoms. Similarly, if $\lambda = 0$ and $\alpha > 0$, the model becomes comparable to an SEIR model, with transmission occurring solely through the bottom row of Table 1.

Let us now define the product of $E_t^S E_t^I$ equal to

$$X_t = E_t^I E_t^S = (\alpha e_t^{NV} + (1-\alpha)e_t^V)((1-\lambda)e_t^{NV} + \lambda e_t^V) \tag{4.3}$$

where $X_t \in [0,1]$ shall be defined as the mitigation policy taken by a government for that period. The number of new infections each period is then simply $\gamma X_t I_t S_t$.

The epidemic dynamics for the model are therefore given by:

$$S_{t+1} = S_t - \gamma X_t I_t S_t \tag{4.4a}$$
$$I_{t+1} = I_t(1-\delta) + \gamma X_t I_t S_t \tag{4.4b}$$
$$R_{t+1} = R_t + \delta I_t \tag{4.4c}$$

Despite the richer information structure, the introduction of symptomatic agents only partitions the state variables and preserves the structure of the SIR model with time-varying contact rates. This maintains tractability and also conserves on state space, with the system dynamic still depending only on the aggregate variables S_t and I_t.

In addition, the SIR structure allows for easy characterisation of the effective reproduction number as:

$$R_t = \frac{\gamma X_t S_t}{\delta} = R_0 X_t S_t$$

X_t is directly proportional to the effective reproduction number in the population and the interpretation as a mitigation policy is now more apparent. A government, through adjustment of this mitigation policy, can directly control the rate of spread of the epidemic in the country. In addition, given X_t $\forall t$, defined as a mitigation path, X, the course of the epidemic is fully deterministic. Most noticeably, given a mitigation policy, the transmission of the virus is independent of the allocation of lockdown across agents. Figure 5 plots the isoline

Figure 4: Partitioned SIR Model

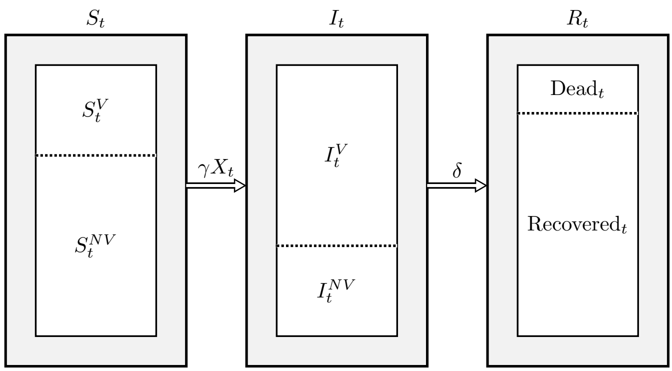

of the mitigation policy, equation (4.2), in policy space for an arbitrary value of X_t, showing the range of lockdown allocations that achieve the same infection dynamics. The line has a negative slope, reflecting that mitigation measures must always move in opposite directions. In addition, as λ increases we see the line pivots clockwise about the 45-degree line to account for the different share of agents under each lockdown policy.

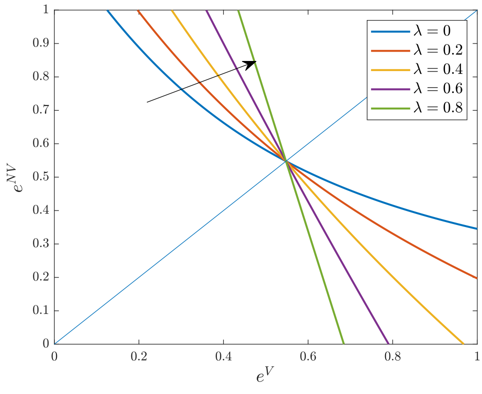

Figure 5: Mitigation Policy as a Function of Lockdown Allocations

Covid Economics 52, 15 October 2020: 1-38

Covid Economics 52, 15 October 2020: 1-38

4.3 The Economy and Welfare

A simple economy is introduced to the model to capture the costs inflicted by the virus and different mitigation paths. A linear production technology is assumed, where each period a type j agent produces e_t^j units of output.[1] In addition, recovered agents are perfectly identified, and having gained immunity to the disease, provide their full endowment of labour. Total output is then equal to the total supply of labour, such that:

$$\text{Output}_t = E_t^S S_t + E_t^I I_t + R_t^a$$

where R_t^a denotes recovered and immune agents. It is assumed there is no saving technology and every agent consumes their entire income each period deriving utility $U(e_t^j)$ from doing so. An egalitarian welfare function is assumed such that total welfare from consumption each period is equal to:

$$TW_t = (\lambda S_t + (1-\alpha)I_t)U(e_t^V) + ((1-\lambda)S_t + \alpha I_t)U(e_t^{NV}) + R_t^a U(e_t^R)$$

It is assumed that U$(e) = ln(e^j) - e^j + 1$ such that utility is maximised at $U(1) = 0$, allowing us to interpret changes in $U(e)$ as welfare losses from lockdown. In addition, since $\lim_{e\to 0} U'(e) = \infty$, if infection levels are non-zero, then interior allocations for (e^V, e^{NV}) are achieved.

4.4 Cost of Fatalities

The high fatality rate of COVID-19, particularly for the elderly, has been a key area of concern. Cumulative fatalities are modelled to evolve according to:

$$D_{t+1} = D_t + \delta(1-\alpha)\kappa I_t$$

where $(1-\alpha)\kappa$ represents the case fatality rate. This assumes that only the proportion of infected agents displaying symptoms are at risk of death each period. The fatality rate is exogenous, with a constant proportion of those infected dying each period, abstracting from explicit modelling of healthcare capacity constraints as has been popular in the literature (Acemoglu et al. 2020). The number of recovered and immune agents, R_t^a is defined as $R_t^a = R_t - D_t$ to preserve the SIR structure of the model.

The costs of a fatality are difficult to tackle. However, it forms a critical value for a planner deciding on an optimal mitigation path to implement. To model this, in addition to the loss in consumption resulting from the death of an agent, we introduce the parameter, Θ, to represent the additional lost value to society from a fatality.

[1] Implicit here is the assumption that an agent produces no output during the periods they are locked down. See Jones et al. (2020) for deeper insight into the impact of work from home options.

Covid Economics 52, 15 October 2020: 1-38

5 The Social Planner Problem

The Social planner problem is formulated as:

$$\max_{\{e_t^V, e_t^{NV}\}_{t=0}^T} \sum_{t=0}^{T} \beta^t \big[(\lambda S_t + (1-\alpha)I_t)U(e_t^V) + ((1-\lambda)S_t + \alpha I_t)U(e_t^{NV}) + R^a U(e_t^R) - (1-\alpha)\delta\kappa I_t \Theta\big]$$

subject to equations (4.4a), (4.4b) and (4.4c).

The sum is finite as it is assumed a vaccine is found in period T, which cures all current infections and grants immunity to all remaining susceptible agents, ending the pandemic in the population.

However, since the recovered agents are assumed to be perfectly identified and immune to further infection, it is optimal to set $e^R = 1$ in all time periods. Since our model has maintained the standard SIR dynamic, the recovered state, R, has no direct impact on transmission. In addition since utility is normalised such that $U(1) = 0$, the R state may be dropped from the problem, reducing the state space from three to two. The Lagrangian can then be written as:

$$\begin{aligned} V(S, I) = \max_{e^V, e^{NV}} \big[&(\lambda S + (1-\alpha)I)U(e^V) + ((1-\lambda)S + \alpha I)U(e^{NV}) \\ &-(1-\alpha)\delta\kappa I\Theta + \beta V(S', I') \\ &+\phi^I(I' - I(1-\delta) - \gamma X(e^V, e^{NV})IS) \\ &+\phi^S(S' - S + \gamma X(e^V, e^{NV})IS)\big] \end{aligned}$$

The solution to this problem can give many insights into the incentives of a social planner in mitigating an epidemic and allocating lockdown measures. The envelope conditions are:

$$V_t^S = \lambda U(e_t^V) + (1-\lambda)U(e_t^{NV}) + \beta\left((1 - \gamma X_t I_t)V_{t+1}^S + \gamma X_t I_t V_{t+1}^I\right)$$

$$V_t^I = (1-\alpha)U(e^V) + \alpha U(e^{NV}) - (1-\alpha)\delta\kappa\Theta + \beta\left((1-\delta)V_{t+1}^I + \gamma X_t S_t (V_{t+1}^I - V_{t+1}^S)\right)$$

The value of an extra susceptible agent is the extra consumption that period, plus the value of remaining susceptible or becoming infected next period. Similarly, an extra infected agent increases consumption, however, also carries the risk of death, $(1-\alpha)\delta\kappa$. With probability $(1-\delta)$ the agent enters the recovered state, and with probability $\gamma S_t X_t$, the agent infects a susceptible agent.

For simplicity, consider the case of complete information failure, where $e_t^{NV} = e_t^V$ $\forall t$, such that both agents have the same consumption. Subtracting the two envelope conditions gives:

$$V_t^S - V_t^I = \delta((1-\alpha)\kappa\Theta + \beta V_{t+1}^S) + \beta(V_{t+1}^S - V_{t+1}^I)(1 - \delta + \gamma S_t X_t - \gamma I_t X_t)$$

Covid Economics 52, 15 October 2020: 1-38

A new infection necessarily entails an agent moving from S_t to I_t. Hence the second term captures the difference in the number of infected agents generated from an extra I and an extra S, respectively. An extra infected agent will lead to $1-\delta+\gamma S_t X_t$ infections next period, while an extra susceptible agent becomes infected with probability $\gamma I_t X_t$. The term $\beta\delta V^S_{t+1}$ simply acts to correct for the fact that the δ flow out of I_t is into the recovered state rather than into S_t.

Through forward substitution, and using that $V^I_{T+1} = V^S_{T+1} = 0$ due to the arrival of a vaccine:

$$V^S_t - V^I_t = \delta((1-\alpha)\kappa\Theta + \beta V^S_{t+1})$$
$$+ \sum_{j=1}^{T-t-1} \left(\beta^j \delta((1-\alpha)\kappa\Theta + \beta V^S_{t+j+1}) \prod_{i=0}^{j-1} (1-\delta+\gamma S_{t+i}X_{t+i} - \gamma I_{t+i}X_{t+i}) \right)$$

This may be approximated by supposing that Θ, the cost of a fatality, is large enough to dominate V^S_{t+1}, the value of remaining susceptible next period, such that:

$$V^S_t - V^I_t \approx (1-\alpha)\delta\kappa\Theta \left(1 + \sum_{j=1}^{T-t-1} \prod_{i=0}^{j-1} \beta^j (1-\delta+\gamma S_{t+i}X_{t+i} - \gamma I_{t+i}X_{t+i}) \right)$$

The planner's motivation for disease mitigation is driven by the fatality risk associated with infected agents. For sufficiently low levels of infection, specifically such that $S_t \geq I_t \ \forall t$, $V^S_t - V^I_t$ must be positive, implying a planner always prefers susceptible agents to infected agents. This partly reflects that an additional infected agent carries a risk of immediate death, $(1-\alpha)\delta\kappa$, that period. However, more interestingly, the planner dislikes infected agents as the stream of future infected agents is higher for an infected agent than a susceptible agent since $1-\delta+\gamma S_t X_t > \gamma I_t X_t$, when $S_t > I_t$. Proposition 1 summarises this.

Proposition 1. *If $S_t > I_t \ \forall t$, and the cost of death, Θ, is sufficiently high, such that $(1-\alpha)\kappa\Theta+\beta V^S_{t+1} > 0$, then $V^S_t - V^I_t > 0$. This implies a planner strictly prefers susceptible agents to infected agents in the economy and will therefore use lockdown measures to target infected agents.*

FOC Analysis

The first order conditions may be written:

$$U'(e^j_t) = \beta(V^I_{t+1} - V^S_{t+1})\gamma((1-P^j_t)E^I_t I_t + P^j_t E^S_t S_t)$$

where $j \in \{V, NV\}$ and $P^j_t = P(I \mid j)$ in period t.

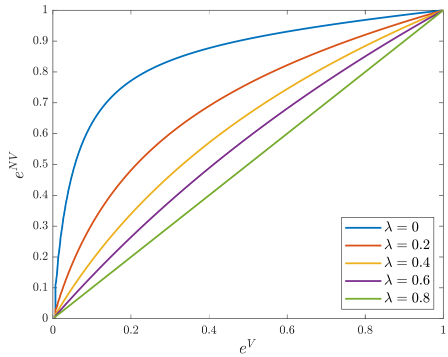

Figure 6: FOC Ratio for Increasing Values of λ

The LHS represents the lost utility from imposing a complete lockdown on an extra, type j agent. The RHS represents the value gained from the reduced rate of infection. If the extra agent locked down is infected, occurring with probability P_t^j, then $\gamma E_t^S S_t$ new infections are avoided. Conversely, if the extra agent is susceptible, occurring with probability $1-P_t^j$, then $\gamma E_t^I I_t$ new infections are avoided.

Combining the two conditions we get a static condition relating the respective lockdown policies on each agent type. To keep notation clean, time subscripts are left out. We shall henceforth refer to this condition as the FOC Ratio.

$$\frac{U'(e^V)}{U'(e^{NV})} = \frac{(1-P^V)E^I I + P^V E^S S}{(1-P^{NV})E^I I + P^{NV} E^S S} \tag{FOC Ratio}$$

The respective probabilities of infection, P^V and P^{NV}, operate like prices, dictating the marginal benefit of increasing each lockdown policy. The higher P^j, the greater the probability that an increase in lockdown policy e^j will remove an infected agent from the economy, and conversely the lower the chance of removing a susceptible agent. Following the intuition in Proposition 1, these infected agents are substantially more dangerous to the population, not from their immediate fatality risk, which is independent of the lockdown decision, but rather owing to the higher stream of future infected agents they produce. As such the higher P^j, the greater the incentive to lockdown type j agents.

Figure 6 plots the FOC ratio in policy space, showing the planner's preference between lockdown policies for given information environments. Immediately, it can be seen that agents with visible symptoms are locked down more aggressively than agents displaying no symptoms, as the indifference curves never lie below the 45-degree line $\forall\ \lambda \leq 1-\alpha$. By Lemma 1, $P^V > P^{NV}$, provided $\lambda <$

Covid Economics 52, 15 October 2020: 1-38

Covid Economics 52, 15 October 2020: 1-38

$1-\alpha$, and so it is always more beneficial for a planner to increase lockdown on symptomatic agents because they are more likely to lockdown an infected agent. Indeed, if agents have linear preferences, it is optimal for a planner to lock down non-symptomatic agents only when all the symptomatic agents are under lockdown. In our model, there are diminishing returns to consumption, giving the indifference curves their concave shape and resulting in the planner allocating lockdown more evenly across agents.

A second interesting feature is that the indifference curves flatten towards the 45-degree line as the quality of information worsens in the economy, represented here by λ increasing. As discussed previously, when λ increases, P^V decreases and P^{NV} increases, reflecting that identification of infected workers becomes harder. As infected agents become increasingly likely to be found in both information sets, the marginal benefits of locking down each agent type converge. Consequently, the planner prefers to spread lockdown allocations more evenly across the population to maximise welfare in the economy. Ultimately, in the instance of complete information failure, a planner is equally likely to lockdown an infected agent in both information sets, and the observation of an agent's symptom state is completely uninformative as to their true infection state. At this point, the indifference curve collapses to the 45-degree line and it is optimal to set uniform lockdown policies that are not conditioned on the presence of symptoms.

It should be noted that Figure 6 displays curves for fixed values of S_t and I_t. The curvature of the indifference curves will increase as the gap between S_t and I_t increases, since the cost of releasing an infected agent is higher when there is a larger population of susceptible agents to infect. For this reason, it becomes more critical for a planner to target infected agents, and as such lockdown is more aggressively placed on agents displaying symptoms than those without. Propositions 2 summarises.

Proposition 2. *Given $0 < \lambda < 1-\alpha$ and $S_t > I_t$, it is optimal to set $e_t^{NV} > e_t^V$. That is to say, provided infection levels are sufficiently low and there is imperfect information, it is optimal for a planner to lockdown those with symptoms more severely. In the case of complete information failure, $\lambda = 1-\alpha$, the optimal lockdown policy sets $e_t^V = e_t^{NV}$.*

Optimal Lockdown Allocation

To solve for the optimal allocation of lockdowns the FOC Ratio condition is combined with the Mitigation Constraint, which was derived previously, and is restated below:

$$X_t = (\alpha e_t^{NV} + (1-\alpha)e_t^V)((1-\lambda)e_t^{NV} + \lambda e_t^V) \qquad \text{(Mitigation Constraint)}$$

This forms a static, non-linear system of two equations and two unknowns, subject to a given mitigation policy. Figure 7a illustrates the two conditions in policy space for $\lambda = 0$ and $\lambda = 0.4$ where the intersection of each set of lines marks the solution to the system. The impact of a change in λ is twofold. Firstly, as discussed above, the increase in λ affects the probabilities of infection, flattening the indifference curve and causing the planner to increase lockdown on non-symptomatic agents and decrease lockdown on symptomatic agents. In addition, the increase in λ also increases the number of agents under the more severe e^V policy, pivoting the Mitigation Constraint clockwise. To maintain the same mitigation level, both lockdown measures must be reduced. Since the first effect dominates the latter, the overall effect is a convergence in lockdown policies.

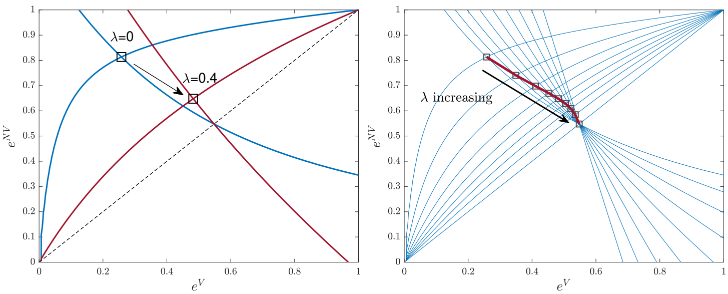

(a) (b)

Figure 7: Illustration of Shifting Lockdown Policies when λ Increases

Using a non-linear solver the system can be solved for $0 \leq \lambda \leq 1-\alpha$ to show how the burden of lockdown shifts between the two agent types as the information in the economy declines. These results are displayed in Figure 7b. The key insight here is that as information deteriorates in the economy, the optimal allocation shifts to the South East, with e^V increasing and e^{NV} decreasing. As infected agents become harder to identify through the presence of symptoms, the planner considers increasingly indiscriminate lockdown policies. Eventually, once $\lambda = 1 - \alpha$, the optimal allocation settles on the 45-degree line, giving a uniform lockdown policy. While these results are shown for arbitrary values of S, I and X, this convergence in policy holds for any values, provided $S > I$.

6 An Application to the UK

The model is now applied to the UK to assess the costs of mitigation on the economy under differing degrees of information failure. First, an approximation of the UK's mitigation path is calibrated from estimates of the UK's effective reproduction number. Under this exogenous mitigation path, the optimal lock-down allocations are solved for each period for different parameterisations of the information environment.

Calibrating the Mitigation Path

To calibrate the mitigation path to the UK, the following relationship is used:

$$R_t = \frac{\gamma X_t}{\delta} S_t = R_0 X_t S_t \tag{6.1}$$

Since S_t is determined in period $t-1$, if provided estimates for $R_t \; \forall t$, as well as initial conditions (S_0, I_0), an iterative process may be set up to deduce a path for X_t.

To simplify analysis, we follow Atkeson (2020) and assume that the effective reproduction number follows the form:

$$R^1_{t+1} = R^1_t - \eta_1(R^1_t - \bar{R}^1_t) \tag{6.2a}$$

$$R^2_{t+1} = R^2_t - \eta_2(R^2_t - \bar{R}^2_t) \tag{6.2b}$$

$$R_t = \frac{R^1_t + R^2_t}{2} \tag{6.2c}$$

where $\bar{R}^1_t$ and $\bar{R}^2_t$ represent the long-run values for R^1 and R^2 respectively and (η_1, η_2) dictate the speed of transition to these states. This is a very flexible specification allowing for differing speeds of mitigation, as well as U-shaped paths for R_t. To inform the calibration of this functional form, we consult government publications and generate estimates for R_t from UK infection data.

Estimating the Effective Reproduction Number in the UK

Estimates for the time path for the UK's reproduction number are found using the R package 'EpiEstim' (Cori et al. 2013). This is a benchmark software, designed to provide quick, real-time estimates of virus reproduction numbers during epidemics.

Covid Economics 52, 15 October 2020: 1-38

The software is popular as it only requires two data inputs. The first is a time series of the number of reported new infections each period. For this, we use the UKGOV *'Coronavirus in the UK'* data set, detailing daily lab-confirmed COVID-19 cases.[2] Since agents are typically tested only when symptoms develop, all points are shifted back 5 days to account for the average incubation period (Lauer et al. 2020). In addition, to account for the reporting delay- the time between the point of testing and reporting of the result- all points are shifted back a further two days (Gutierrez et al. 2020). The second data input is an estimate for the generation interval of COVID-19, defined as the average duration between becoming infected and subsequently infecting another agent. This is difficult to estimate for it is the point at which an agent becomes symptomatic, not the point at which an agent becomes infected, that is typically observed. Instead, the generation interval is approximated by the serial interval, defined as the average duration between symptom onset in the original case and symptom development in the successive infected case (Fine 2003, Flaxman et al. 2020). The serial interval is taken to be gamma distributed with mean 4.7 and standard deviation 2.9 (Nishiura et al. 2020).

Figure 8a displays the estimates starting from March 1st through to July 21st. While the exact timings of these estimates are uncertain, there are 3 main observations. Firstly, the reproduction number appears stable before mid-March, oscillating around $R_t \approx 2.5$. This is taken to be the period before meaningful action was undertaken by the UK government. Secondly, there is then a pronounced fall in R_t, falling to below 1 before the end of April. This coincides with the government introducing social distancing measures from March 16th (Johnson 2020a). These measures were then rapidly escalated over the next week, with most social venues and schools closing on March 20th (Johnson 2020b). By March 23rd, the country was placed under a nationwide lockdown, with people instructed to remain indoors except for absolute necessities (Johnson 2020c). Finally, progressing into May, R_t stabilises, before gradually increasing towards 1. This is interpreted as the government switching to a policy of suppression, applying the necessary mitigation to hold R_t just below 1 to maintain a manageable and low level of infection in the economy (UKGOV 2020).

These 3 observations, corroborated with government documents and announcements, inform the model calibration of R_t. Figure 8b plots this specification. It is assumed there is no mitigation before March 16th and R_t follows a standard SIR model with $R_0 = 2.5$. When the UK starts a suppression policy on March 16th, the reproduction number begins at $R_t = 2.45$, before rapidly falling below one, following the dynamic described in equations (6.2). The remainder of the model must now be calibrated to allow for the mitigation path to be constructed.

[2] Data may be accessed and downloaded at https://coronavirus-staging.data.gov.uk

[3] Parameters for fitting R_t: $\bar{R}_1 = 0.95$, $\bar{R}_2 = 1.05$, $\eta_1 = \frac{1}{22}, \eta_2 = \frac{1}{37}$. In addition, it is assumed $R_0 = 2.5$ and $R_t = 2.45$ on March 16th

Covid Economics 52, 15 October 2020: 1-38

Covid Economics 52, 15 October 2020: 1-38

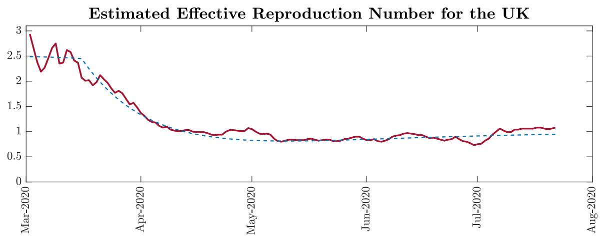

(a) Estimated Effective Reproduction Number for the UK using 'EpiEstim'

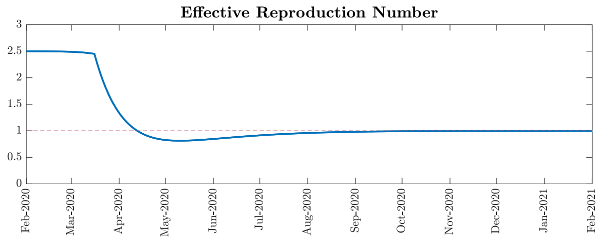

(b) Fitted Suppression Policy for the UK[3]

Figure 8

Covid Economics 52, 15 October 2020: 1-38

7 Calibration

The reproduction number, R_0, is a critical parameter dictating the size and spread of the epidemic. However, despite the fundamental importance of this number, its value is unknown. A benchmark figure frequently cited in the economics literature is $R_0 = 3.1$, based on early infection data from Wuhan (Wang et al. 2020). However, in light of more extensive infection data, this estimate appears too high. Later estimates, also based on Wuhan, suggest substantially lower values ranging from $R_0 = 2.28$ (Zhang et al. 2020) to 2.68 (Wu et al. 2020). Similarly, in a recent survey, the ECDC (2020) found the reproduction number to lie in the interval of 2 to 3 . As seen in the fitting of the R_t time path, we take an intermediate value of $R_0 = 2.5$, coinciding with that considered in Anderson et al. (2020) and Berger et al. (2020). The recovery rate, δ, is less uncertain, with the WHO (2020a) estimating a 14 day recovery period, such that $\delta = 1/14$. This then leaves $\gamma = 0.18$.

We calibrate S_0 and I_0 to achieve a target R_t path before mitigation efforts were enforced in the UK. Assuming a constant contact rate pre-government intervention, for $R_0 = 2.5$ and $R_t = 2.45$ on March 16th, then it must be that $S_0 = 0.999859933$ and $I_0 = 0.00014006665$. This is substantially lower than the values considered in the economics literature, where $I_0 = 0.01$ is common. However, the value is more in line with the medical literature which often considers a seed infection several magnitudes of order smaller (Atkeson 2020). As we shall see, under a suppression policy, the arrival of herd immunity is very sensitive to this calibration. For robustness, we consider alternative values of R_0, paired with appropriate levels of initial infection, to maintain $R_t = 2.45$ when government mitigation begins.

Two key variables remaining are the asymptomatic rate, α, and the fatality rate, $(1-\alpha)\kappa$. The asymptomatic rate has been a key area of research since COVID-19 began, however, the evidence is varied and often plagued by small sample sizes. In light of evidence from the Princess Diamond Cruise ship, with a relatively large sample of 3,700 individuals, we set $\alpha = 0.2$ (Mizumoto et al. 2020). The fatality rate of COVID-19 is also an active area of research, however, empirical estimation is difficult due to the under-reporting of mild and asymptomatic cases. In addition there is substantial heterogeneity by age, with fatality rates rising sharply for those over 65 (Ferguson et al. 2020). Consequently, many of the estimates are model-based with mean values ranging from 0.66 (Verity et al. 2020) to 1 (Dorigatti et al. 2020, Famulare 2020). Similarly, the WHO (2020b) reports the fatality rate to be between 0.3 and 1. We consider the lower end of this interval to target a fatality rate of 0.4% to better represent the lower age demographic in the working population of the UK. This sets $\kappa = 0.005$, giving a daily mortality rate of $\delta(1-\alpha)\kappa$. Finally, $\beta = 1 - \frac{0.05}{365}$ to achieve an annual discount rate of 5%.

CEPR PRESS

Covid Economics 52, 15 October 2020: 1-38

Table 2: Parameter Value

Parameters	Value	Definition
γ	0.178571429	Infection contact rate
δ	$1/14$	Daily recovery rate
$(1-\alpha)\delta\kappa$	$(1-\alpha)\delta \times 0.005$	Daily fatality rate
α	0.2	Asymptomatic rate
β	$1-0.05/365$	Daily discount rate
S_0	0.999859933	Initial number of susceptible agents
I_0	0.000140067	Initial number of infected agents

8 Results

We begin with a discussion of the mitigation path followed by the UK, illustrated in Figure 9. We then discuss the epidemic dynamics implied by this mitigation path with Table 3 providing summary statistics for the outcomes with and without mitigation. Table 4 gives summary statistic for lockdown allocations under different values of λ and lost output for alternative vaccine arrival dates.

8.1 Mitigation Path

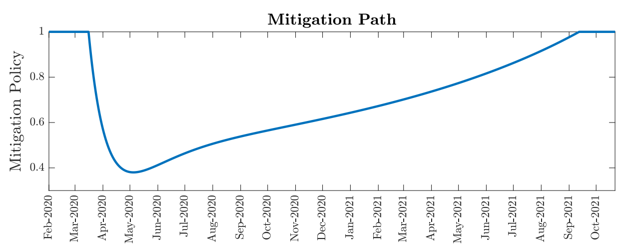

Figure 9: Mitigation Path for the UK

Beginning March 16th, sharp mitigation measures are imposed across the population with contact rates falling by half within 3 weeks. Once R_t settles near 1, mitigation is gradually reduced as the number of susceptible agents in the economy falls. By September 2021, 19 months after mitigation started, herd immunity is reached and the virus prevalence will naturally decline without the need for any further mitigation.

Covid Economics 52, 15 October 2020: 1-38

8.2 Epidemic Dynamics: With and Without Mitigation

$R_0 = 2.5$		No Mitigation	Mitigation	% Change
Months until Herd Immunity		3	19	-
Cumulative Infection		89.70%	68.70%	-
Peak Infection		23.90%	3.47%	-
Fatalities:	12 Months	0.36%	0.16%	-56.06%
	18 Months	0.36%	0.22%	-39.06%
	24 Months	0.36%	0.27%	-24.91%

Table 3: Epidemic Dynamics: Baseline Case

Table 3 and Figure 10 show the significant differences in the epidemic dynamic, with and without mitigation. Without mitigation, the virus tears quickly through the economy, rapidly propagating with peak prevalence reaching 23.9%, only 2.5 months after the first case is observed. However, as quickly as it grows, it soon dissipates as rising immunity levels in the population reduce the force of infection. By the end of the year, the virus is near extinct with only 50 active cases in the whole population, having infected approximately 90% of all agents over this period. In contrast, the government mitigation works to swiftly lower R_t in the early stages of infection, limiting peak infection to just 3.5%. In exchange, the virus persists for significantly longer in the population, held at a low prevalence of 1.2% until either herd immunity is reached in Autumn 2021 or a vaccine arrives. This drawn out epidemic captures the notion of *flattening* the curve, a phrase frequently used in the early days of COVID-19 mitigation in the UK.

These alternative infection paths lead to pronounced differences in fatalities. Over the course of 2 years, in the absence of a vaccine, the suppression policy reduces total fatalities by 24.9%. Furthermore, this number increases the earlier a vaccine arrives, rising to a 39.1% reduction if the vaccine arrives within 18 months. This demonstrates that the value of a suppression policy lies not just in lowering the cumulative fatality rate, but also in delaying the incidence of these fatalities, allowing time for vaccine development.

8.3 Lockdown Allocations

To achieve the mitigation path the planner should optimise the allocation of lockdown across the population, conditioning on an agent's observable symptoms. Below we discuss three key features of the lockdown allocations.

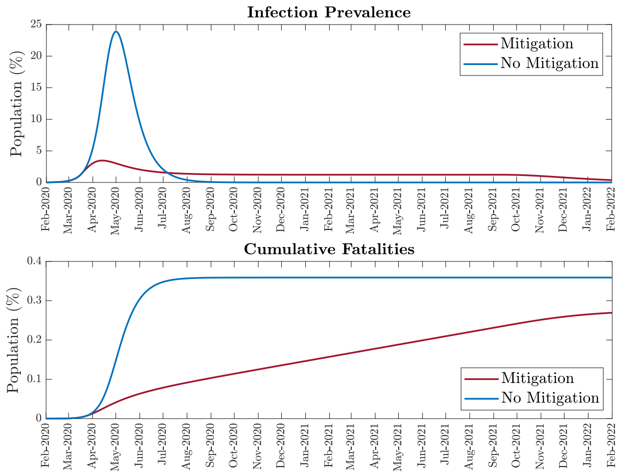

Figure 10

Table 4: Baseline Lockdown Allocations and Lost Output

$R_0 = 2.5$	Peak Lockdown		Lost Output: % of Annual GDP		
λ	$1 - e^{NV}$	$1 - e^{V}$	12 Months	18 Months	24 Months
0	5.4%	73.4%	-2.34%	-2.76%	-2.89%
0.1	16.0%	61.2%	-9.54%	-11.09%	-11.25%
0.2	21.3%	53.8%	-13.09%	-15.22%	-15.40%
0.3	24.7%	48.9%	-15.13%	-17.60%	-17.79%
0.4	27.2%	45.6%	-16.40%	-19.10%	-19.30%
0.5	29.3%	43.1%	-17.25%	-20.11%	-20.31%
0.6	31.5%	41.2%	-17.82%	-20.79%	-20.99%
0.7	34.2%	39.7%	-18.19%	-21.22%	-21.43%
0.8	38.3%	38.3%	-18.33%	-21.39%	-21.60%

Notes: *Output is compared against an economy that has no virus, reflecting the headline reduction in GDP from the virus and the mitigation effort combined. We standardise the lost output under different vaccine arrivals as a proportion of annual GDP. The three rightmost columns indicate the percentage fall in output if a vaccine arrives at 12, 18 and 24 months, respectively.*

Covid Economics 52, 15 October 2020: 1-38

Lockdown allocations become less targeted as λ increases

Figure 11 displays how the spread of lockdown policies across agents depends on the degree of information failure, λ, with the largest difference displayed when λ is low. For $\lambda = 0.1$, the planner implements a strong lockdown of 61.2% on those displaying symptoms, while only restricting activity for non-symptomatic agents by 16.0%. This high level of targeting is driven by the planner's ability to discern infected agents by the presence of symptoms. However, as the degree of information failure increases, this ability is eroded and the planner becomes increasingly likely to find infected agents in both information sets. It is this convergence in $P(I \mid j)$ for $j \in \{NV, V\}$ that drives the convergence in lockdown policies. In the case of complete information failure, $P(I \mid V) = P(I \mid NV)$, and a uniform lockdown, which peaks at 38.3%, is set for all agents.

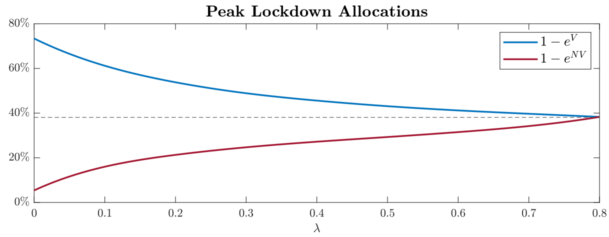

Figure 11: Lockdown Policies Converge as λ Increases

The rate at which lockdown allocations become less targeted is decreasing in λ

As may be seen in Figure 11, the rate at which the lockdown policies converge to a uniform policy is non-linear and decreasing in λ. The increase of λ from 0 to 0.1 induces the largest shift in lockdown allocation, with symptomatic agents seeing a 12.2 percentage point reduction in their peak lockdown. Meanwhile an increase of λ from 0.7 to 0.8 causes a minor reduction of only 1.4 percentage points. As demonstrated previously, for low levels of infection, $P(I \mid V)$ is very sensitive to low values for λ. Consequently, a small increase in λ above zero rapidly dilutes the information carried by the presence of symptoms, causing a large reallocation towards a more balanced lockdown across the population.

Covid Economics 52, 15 October 2020: 1-38

CEPR PRESS

Covid Economics 52, 15 October 2020: 1-38

Figure 12: Lockdown Allocations when $\lambda = 0$ and $\alpha = 0.2$

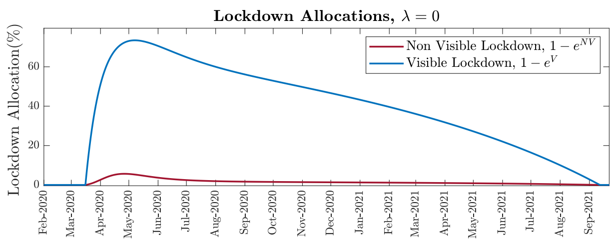

Asymptomatic transmission is costly to reduce

The case where $\lambda = 0$ is noteworthy as the model now has the same information structure as the popular SEIR model. As can be seen in Figure 12, despite the presence of asymptomatic agents in the economy, the planner keeps the burden of lockdown overwhelmingly on those displaying symptoms. Due to the low infection prevalence, the asymptomatic agents comprise only a small proportion of the NV information set. This makes targeting asymptomatic agents very difficult, requiring a planner to impose costly lockdowns on all those without symptoms in order to reduce the rate of asymptomatic transmission. Meanwhile, those displaying symptoms are perfectly identified, allowing the planner to efficiently reduce the number of infected agents in the economy. Furthermore, in general, since the level of infection is typically low, relative to those still susceptible, a change in the asymptomatic rate is less significant than an equivalent change in λ. This result is striking as it shows that the motivation for more uniform lockdown policies is driven largely by how identifiable the virus symptoms are rather than the presence of asymptomatic agents.

8.4 The Cost of Mitigation on the Economy

Finally, we consider the costs to the economy from the suppression policy, as measured by the total loss in output from the virus and mitigation measures. Since an exogenous mitigation path is implemented, variation in costs are driven entirely by the information environment and the consequent allocations of lockdown. Figure 13 shows the total loss in output as a proportion of annual GDP, with a vaccine arriving in 18 months. There are two key observations here.

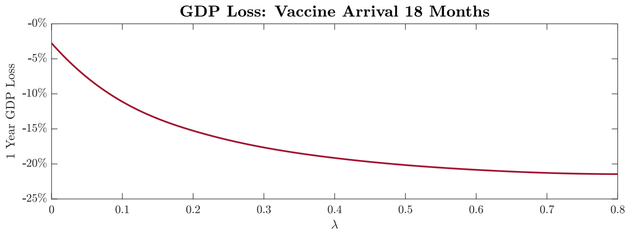

Figure 13: Total Loss in Output as a Proportion of Annual GDP

The loss in output increases as information worsens in the economy

When λ is high and identification of infected agents is difficult, a planner must lockdown a larger share of the population in order to limit the spread of the virus and maintain the mitigation policy. When $\lambda = 0$, the planner is able to implement a very efficient lockdown policy, aggressively locking down the symptom-displaying agents. This allows the majority of the population to continue economic activity unimpeded. As such the contraction in the economy is small, equivalent to a 2.76% reduction in annual GDP, even if a vaccine arrives after only 18 months. As λ increases, and infected agents become harder to identify, the planner must lockdown a larger share of the population to maintain the same level of infection. In the case of complete information failure, mitigation is very costly with a single, mass lockdown applied to all agents, resulting in a contraction in the economy equivalent to a 21.4% fall in GDP that year. It is worth noting that since the virus spreads as before, but the costs of mitigation have increased, a planner may have an incentive to reduce the mitigation level as the information in the economy worsens. This could be shown through the full social planner problem, however, is outside the scope of this paper.

Output is more sensitive to changes in λ when λ is low

This result is the natural conclusion of many of the results shown in this paper. It is built upon the premise that the number of infected agents is significantly lower than the number of susceptible agents, such that only a small increase in λ quickly compromises the effectiveness of diagnosing agents by the presence of a symptom. The increase in λ from 0.1 to 0.2 sees the economy contract approximately 4%, while the increase from 0.2 to 0.3 only causes a further 2% contraction. For $\lambda \geq 0.4$, there is little change in output for a change in λ. At this point, the number of symptomatic susceptible agents is high enough that there is little information contained in a symptom and therefore there is limited value to be gained from conditioning lockdown policies.

Covid Economics 52, 15 October 2020: 1-38

9 Robustness

As a simple robustness exercise, we set $\gamma = 0.196$ such that $R_0 = 2.75$. In order to maintain $R_t = 2.45$ on March 16th this requires $I_0 = 0.000444585$ and $S_0 = 0.999555415$. Despite only a small increase in R_t, this calibration amounts to a virus that is both substantially more virulent and also more established in the population, with three times as many initial cases. Table 5 summarises the new epidemic dynamics, Table 6 details how lockdown is allocated across agent types and Figure 14 illustrates the new infection curve, cumulative fatalities and costs to the economy- these may all be found in the appendix.

The exercise demonstrates that the quantitative predictions of the model are very sensitive to parameter values chosen, and as such, they should be interpreted with caution. Under this calibration, the UK is reacting much later to the virus. By the time that mitigation measures are introduced, the virus is well established in the population with a prevalence of 7.4% in the population. As a result the population has developed more immunity to the virus, such that significantly lighter lockdown measures are needed to control the virus. These measures reduce the peak infection by around a third, however, do little to delay the virus, with herd immunity achieved in mid-April for both cases, well before the arrival of a vaccine. The impact on the economy is mild. The fatality rate, however, is substantially higher at 0.357%, a 32.2% increase on the baseline calibration.

However, we hope to illustrate that the broad principles from Section 8, on how to optimise in the presence of information frictions, are substantially more robust. These principles do not depend on uncertain epidemic parameters, such as R_0, and depend only on the information parameters, α and λ, the state variables, S_t and I_t, and the level of mitigation X_t specified. Consequently, even under this limited, herd immunity based policy, there is substantial targeting of symptomatic agents. In addition, the output costs are still increasing as identification becomes harder, such that a policy maker neglecting to condition properly on the presence of symptoms will experience a larger recession in the economy for no decrease in the number of fatalities.

10 Conclusion

This paper takes steps, through the introduction of asymptomatic infected agents and symptomatic susceptible agents, to parameterise the information failure experienced by policy makers when trying to identify infected agents. This information structure is integrated into the tractable SIR model to study the interaction of targeted lockdown policies, virus transmission and the economy.

Covid Economics 52, 15 October 2020: 1-38

There are a few limitations to the model such that the quantitative predictions should be interpreted with caution. Firstly, under the SIR model, it is assumed that once an agent has recovered they are immune from further infection. This is an important assumption, allowing for immunity to build in the population and the virus to be eradicated in the long run, even in the absence of a vaccine. However, there is still considerable uncertainty around the validity of this assumption in the medical literature. In the absence of rising immunity levels, the hump-shaped mitigation path would be lost, and strong measures would need to stay in place until the arrival of a vaccine. Secondly, to preserve a parsimonious state space, recovered agents are assumed to be perfectly identified in the population and therefore are not subject to any lockdown measures. Future research could relax this assumption and allow for recovered agents to display symptoms, such that they either appear symptomatic or non-symptomatic. This would impair the identification of infected agents further, and lead to higher lockdown levels on both non-symptomatic and symptomatic agents. Finally, as shown in the robustness section, the quantitative predictions from the model are sensitive to some key, uncertain epidemic values.

Despite this, we stress that the qualitative findings of the paper, describing optimality conditions to allocate lockdowns when faced with incomplete information, are robust. This includes the central finding of the paper that the targeting of lockdown policies on the presence of a symptom can reduce the costs to the economy, with no increase in the number of fatalities. This disrupts the notion that there is a fixed frontier between fatalities and output such that increased mitigation measures need not necessarily entail higher costs to the economy.

There are two immediate implications to these results that could warrant further empirical research. It could be the case that countries which achieved a low number of fatalities and only moderate impacts to the economy, may have done so due to better execution of a mitigation path, in the respect that they better targeted infected agents. This could be an important factor to explain the heterogeneity seen in output and fatality rates across countries. Secondly, our results suggest that there is a substantial threat posed by the winter months, where the increased prevalence of the common cold and flu will make the identification of infected agents harder. To ensure the reproduction number stays equal to one, more agents will need to be locked down over winter, with lockdown measures on those not displaying symptoms increasing.

Covid Economics 52, 15 October 2020: 1-38

Covid Economics 52, 15 October 2020: 1-38

Appendix

Robustness

Table 5: Epidemic Dynamics: $R_0 = 2.75$

$R_0 = 2.75$		No Mitigaton	Mitigation
Duration until Herd Immunity		2 Months	2 Months
Cumulative Infection		92.50%	89.35%
Peak Infection		27.55%	19.51%
Fatalities after:	12 Months	0.370%	0.357%
	18 Months	0.370%	0.357%
	24 Months	0.370%	0.357%

Table 6: Lockdown Allocations: $R_0 = 2.75$

$R_0 = 2.75$	Peak Lockdown		Lost Output: % of Annual GDP		
λ	$1 - e^{NV}$	$1 - e^{V}$	12 Months	18 Months	24 Months
0	8.5%	21.9%	-0.62%	-0.79%	-0.95%
0.1	9.5%	19.5%	-0.70%	-0.87%	-1.04%
0.2	10.1%	17.9%	-0.75%	-0.92%	-1.08%
0.3	10.6%	16.8%	-0.78%	-0.95%	-1.12%
0.4	11.1%	16.0%	-0.81%	-0.97%	-1.14%
0.5	11.6%	15.3%	-0.82%	-0.99%	-1.15%
0.6	12.1%	14.8%	-0.83%	-1.00%	-1.16%
0.7	12.9%	14.4%	-0.84%	-1.01%	-1.17%
0.8	14.0%	14.0%	-0.84%	-1.01%	-1.17%

Covid Economics 52, 15 October 2020: 1-38

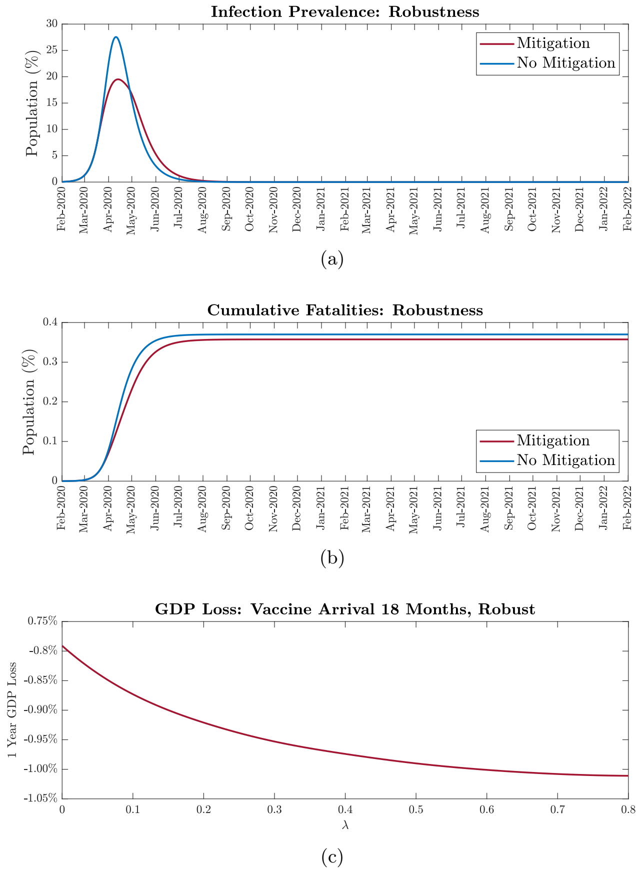

Figure 14: Robustness Exercise, $R_0 = 2.75$

Covid Economics 52, 15 October 2020: 1-38

References

Acemoglu, D., Chernozhukov, V., Werning, I. & Whinston, M. D. (2020), Optimal targeted lockdowns in a multi-group SIR model, Working Paper 27102, National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w27102*

Alvarez, F. E., Argente, D. & Lippi, F. (2020), 'A simple planning problem for COVID-19 lockdown', *Covid Economics* **14**, 1–32.

Anderson, R. M., Heesterbeek, H., Klinkenberg, D. & Hollingsworth, T. D. (2020), 'How will country-based mitigation measures influence the course of the COVID-19 epidemic?', *The Lancet* **395**(10228), 931–934.

Atkeson, A. (2020), What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios, Working Paper 26867, National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w26867*

Bairoliya, N. & Imrohoroglu, A. (2020), 'Macroeconomic consequences of stay-at-home policies during the COVID-19 pandemic', *Covid Economics* **13**, 71–90.

Berger, D. W., Herkenhoff, K. F. & Mongey, S. (2020), 'An SEIR infectious disease model with testing and conditional quarantine', *Covid Economics* **13**, 1–30.

Chari, V. V., Kirpalani, R. & Phelan, C. (2020), The hammer and the scalpel: On the economics of indiscriminate versus targeted isolation policies during pandemics, Technical report, National Bureau of Economic Research.

Cori, A., Ferguson, N. M., Fraser, C. & Cauchemez, S. (2013), 'A new framework and software to estimate time-varying reproduction numbers during epidemics', *American journal of epidemiology* **178**(9), 1505–1512.

Dorigatti, I., Okell, L., Cori, A., Imai, N., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunubá, Z., Cuomo-Dannenburg, G., FitzJohn, R. et al. (2020), Report 4: severity of 2019-novel coronavirus (nCoV), Technical report, Imperial College London.

ECDC (2020), Novel Coronavirus disease 2019 pandemic: Increased transmission in the EU/EEA and the UK- sixth update, Technical report, European Centre for Disease Prevention and Control.

Eichenbaum, M. S., Rebelo, S. & Trabandt, M. (2020), The macroeconomics of epidemics, Working Paper 26882, National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w26882*

Famulare, M. (2020), '2019-nCoV: preliminary estimates of the confirmed-case-fatality-ratio and infection-fatality-ratio, and initial pandemic risk assessment', *Institute for Disease Modeling* **19**.

Farboodi, M., Jarosch, G. & Shimer, R. (2020), 'Internal and external effects of social distancing in a pandemic', *Covid Economics* **9**, 22–58.

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 COVID-19 mortality and healthcare demand, Technical report, Imperial College London.

Fine, P. E. (2003), 'The interval between successive cases of an infectious disease', *American journal of epidemiology* **158**(11), 1039–1047.

Flaxman, S., Mishr, S., Gandt, A., T.Ulwin H., J., Mellan, T. A., Coupland, H., Whittaker, C., Zhu, H., Berah, T., Eaton, J. W., Monad, M., Chani, A. C., Donnelly, C. A., Riley, S. M., Vollmer, M. A. C., Ferguson, N. M., Okell, L. C. & Bhatt, S. (2020), 'Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe'.
URL: *https://doi.org/10.1038/s41586-020-2405-7*

Garibaldi, P., Moen, E. R. & Pissarides, C. A. (2020), 'Modelling contacts and transitions in the SIR epidemics model', *Covid Economics Vetted and Real-Time Papers, CEPR* .

Ghebreyesus, T. A. (2020), 'WHO Director-General's opening remarks at the media briefing on COVID-19- 11 March 2020'.
URL: *https://www.who.int/dg/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19—11-march-2020*

Glover, A., Heathcote, J., Krueger, D. & Ríos-Rull, J.-V. (2020), 'Health versus wealth: On the distributional effects of controlling a pandemic', *Covid Economics* **6**, 22–64.

Greenstone, M. & Nigam, V. (2020), 'Does social distancing matter?', *University of Chicago, Becker Friedman Institute for Economics Working Paper* **2020-26**.

Gutierrez, E., Rubli, A. & Tavares, T. (2020), 'Delays in death reports and their implications for tracking the evolution of COVID-19', *Covid Economics* **34**, 116–144.

Johnson, B. (2020a), 'Prime minister's statement on coronavirus (COVID-19): 16 march 2020'.
URL: *https://www.gov.uk/government/speeches/controlling-the-spread-of-covid-19-health-secretarys-statement-to-parliament*

Covid Economics 52, 15 October 2020: 1-38

Johnson, B. (2020b), 'Prime minister's statement on coronavirus (COVID-19): 20 march 2020'.
URL: *https://www.gov.uk/government/speeches/pm-statement-on-coronavirus-20-march-2020*

Johnson, B. (2020c), 'Prime minister's statement on coronavirus (COVID-19): 23 march 2020'.
URL: *https://www.gov.uk/government/speeches/pm-address-to-the-nation-on-coronavirus-23-march-2020*

Jones, C. J., Philippon, T. & Venkateswaran, V. (2020), Optimal mitigation policies in a pandemic: Social distancing and working from home, Working Paper 26984, National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w26984*

Kermack, W. O. & 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.

Lauer, S. A., Grantz, K. H., Bi, Q., Jones, F. K., Zheng, Q., Meredith, H. R., Azman, A. S., Reich, N. G. & Lessler, J. (2020), 'The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application', *Annals of internal medicine* **172**(9), 577–582.

Mizumoto, K., Kagaya, K., Zarebski, A. & Chowell, G. (2020), 'Estimating the asymptomatic proportion of coronavirus disease 2019 (COVID-19) cases on board the Diamond Princess cruise ship, Yokohama, Japan, 2020', *Eurosurveillance* **25**(10), 2000180.

Nishiura, H., Linton, N. M. & Akhmetzhanov, A. R. (2020), 'Serial interval of novel coronavirus (COVID-19) infections', *International journal of infectious diseases* .

Piguillem, F., Shi, L. et al. (2020), The optimal COVID-19 quarantine and testing policies, Technical report, Einaudi Institute for Economics and Finance (EIEF).

Rampini, A. A. (2020), Sequential lifting of COVID-19 interventions with population heterogeneity, Working Paper 27063, National Bureau of Economic Research.
URL: *http://www.nber.org/papers/w27063*

Stock, J. H. (2020), 'Data gaps and the policy response to the novel coronavirus', *Covid Economics* **3**, 1–11.

Toxvaerd, F. (2020), 'Equilibrium social distancing', *Covid Economics* **15**, 110–133.

Covid Economics 52, 15 October 2020: 1-38

CEPR PRESS

UKGOV (2020), 'The next chapter in our plan to rebuild: The UK government's COVID-19 recovery strategy'.
URL: *https://www.gov.uk/government/publications/our-plan-to-rebuild-the-uk-governments-covid-19-recovery-strategy*

Verity, R., Okell, L. C., Dorigatti, I., Winskill, P., Whittaker, C., Imai, N., Cuomo-Dannenburg, G., Thompson, H., Walker, P. G., Fu, H. et al. (2020), 'Estimates of the severity of coronavirus disease 2019: a model-based analysis', *The Lancet infectious diseases* .

Wang, H., Wang, Z., Dong, Y., Chang, R., Xu, C., Yu, X., Zhang, S., Tsamlag, L., Shang, M., Huang, J. et al. (2020), 'Phase-adjusted estimation of the number of coronavirus disease 2019 cases in Wuhan, China', *Cell discovery* **6**(1), 1–8.

WHO (2020a), Report of the WHO-China joint mission on coronavirus disease 2019, Technical report, World Health Organisation.
URL: *https://www.who.int/publications/i/item/report-of-the-who-china-joint-mission-on-coronavirus-disease-2019-(covid-19)*

WHO (2020b), Coronavirus disease 2019 (COVID-19): situation report, 30, Technical report, World Health Organization and others.
URL: *https://www.who.int/docs/default-source/coronaviruse/situation-reports/20200219-sitrep-30-covid-19.pdf?sfvrsn=3346b04f$_2$*

Wu, J. T., Leung, K. & Leung, G. M. (2020), 'Nowcasting and forecasting the potential domestic and international spread of the 2019-nCoV outbreak originating in Wuhan, China: a modelling study', *The Lancet* **395**(10225), 689–697.

Zhang, S., Diao, M., Yu, W., Pei, L., Lin, Z. & Chen, D. (2020), 'Estimation of the reproductive number of novel coronavirus (COVID-19) and the probable outbreak size on the Diamond Princess cruise ship: A data-driven analysis', *International journal of infectious diseases* **93**, 201–204.

Covid Economics 52, 15 October 2020: 1-38

Covid Economics **Issue 52, 15 October 2020**

"Heard" immunity: Messages emphasizing the safety of others increase intended uptake of a COVID-19 vaccine in some groups[1]

Nicolas Duquette[2]

Date submitted: 5 October 2020; Date accepted: 6 October 2020

A survey experiment exposes treatment groups to four messages supporting future vaccination against COVID-19. These treatments emphasize either the risks of the virus or the safety of vaccination, to the respondent personally or to others. For a nationally representative sample, self-reported intent to vaccinate is not significantly different from the control for any message. However, there is a substantial divergence between white non-Hispanic respondents, whose response to all four treatments is close to zero, and non-white or Hispanic respondents, whose intention to vaccinate is over 50% higher in response to a message emphasizing pro-sociality and the safety of others.

Covid Economics 52, 15 October 2020: 39-67

1 I am grateful for advice on the design of this experiment from John Romley, Anya Samek, Neeraj Sood, and Joanna Duquette, and from the amazing team at the Understanding America Survey. This research was funded using my discretionary faculty research account, for which I am very grateful to the USC Price School of Public Policy. All mistakes are the author's.

2 Associate Professor, Sol Price School of Public Policy, University of Southern California.

CEPR PRESS

Covid Economics 52, 15 October 2020: 39-67

1 Vaccine hesitancy could make this pandemic longer and worse

The COVID-19 pandemic will only end permanently when the population reaches herd immunity. If this is to happen without hundreds of millions of deadly and dangerous infections, a vaccine must be developed and widely used, quickly. Each person who successfully acquires immunity from a safe and effective vaccine not only avoids their own infection but also slows community spread of the virus.

Even if a safe and effective vaccine is discovered and tested quickly, however, a large minority of Americans may not take it, and policies to coerce vaccination may backfire. The decentralized structure and libertarian ethos of US institutions have historically caused the US to lag peer countries in containment of epidemic disease (Troesken 2015). Anti-vaccine activists exploit these institutions to resist immunization mandates (Reich 2018). Furthermore, even people who are ordinarily supportive of vaccination may not want the COVID-19 vaccine, if they do not trust that its safety has been adequately tested. For example, distrust of the American health system in general is widespread among Black Americans as a result of historical patterns of medical exploitation, causing reduced uptake of life-saving preventive health measures (Alsan and Wanamaker 2017).[1] Furthermore, because children and young adults are at lower risk from COVID-19, the coercive policy most widely and successfully used in the US—school vaccination requirements—would not directly protect the most endangered groups.

Since coercive policies are unlikely to be effective, uptake of the COVID-19 vaccine in the US may depend on finding a message that will successfully persuade the vaccine-hesitant to get their shot. The survey experiment described in this paper tests four messages for effects on respondents' self-reported intention to vaccinate when a vaccine becomes available.

The effect of the messages varies starkly by race and ethnicity. For the complete, nationally representative sample, no message has a statistically significant effect relative to a control group. This is because non-Hispanic white respondents are unresponsive to all messages, with point estimates consistently close to zero. But for non-white and Hispanic respondents, the messages have sizable estimated treatment effects: the most effective message, one emphasizing vaccine safety and protection of others, increases this group's intention to vaccinate by 50.4%.

This heterogeneous treatment effect is explained by varying reasons for vaccine hesitancy. Non-Hispanic white respondents are more likely than others to express lack of fear of the COVID-19 illness, and thus less likely to see any need for a vaccine. Non-white and Hispanic/Latino respondents are more likely than non-Hispanic whites to hesitate because of fears that the COVID-19 vaccine itself will be unsafe or because it will cost too much.

This heterogeneous treatment effect is important for the general public health and for equity:

[1]Herd immunity may also be difficult to reach if there are many people who would be willing to vaccinate but are deterred by other factors. For example, cost and access to health services affect vaccination rates. While many other preventive health measures are not cost-sensitive, (Kolstad and Kowalski 2012, table 4) find that seasonal influenza vaccine uptake is an exception. The out-of-pocket cost of any COVID-19 vaccine may therefore also affect uptake. Also, vaccines are contraindicated for people with immune system disorders and other health risks, which means herd immunity must arise from coverage within the remaining population.

Covid Economics 52, 15 October 2020: 39-67

Black, Native, and Latino Americans exhibit higher vaccine hesitancy than other groups, and have suffered higher infection and death rates from COVID-19 (Yancy 2020; Tai et al. 2020). A messaging campaign that increases vaccine uptake in this population could therefore have important public health benefits without any effect on non-Hispanic white uptake.

2 Research Design

This study uses a survey experiment to examine the effects of COVID-19 vaccine messaging on intended uptake. Messaging has been demonstrated to effectively increase vaccine compliance in other contexts; for example, simple differences in the ways pediatricians frame vaccination increases compliance with the pediatric vaccine schedule (Opel et al. 2013). Historically, public health officials have experimented with messages emphasizing the personal benefits conferred by vaccination on oneself and one's family, or the social duty of vaccine compliance (Colgrove 2006; Reich 2018), but we know little about which of these was more effective. Therefore, it is an open question whether the kind of messages that persuade parents to vaccinate their children against now-uncommon diseases will be the most effective in the context of a deadly viral pandemic.

2.1 A two-by-two message design

Four messages are structured to emphasize the risks of the virus or the safety of the vaccine, from the perspective of self-interest or of social responsibility. I join these two emphases and two perspectives into four messages using a two-by-two design. The same or similar sentences are combined across treatments to allow for clearer comparisons.[2] The *Risk-to-Self* message emphasizes the danger the virus poses to those who are not immune:

> It is important to vaccinate against COVID-19 because it is **dangerous**. The virus can harm the lungs, heart, brain, and other body systems. While elderly and ill people are at extra risk, COVID-19 is potentially deadly for anyone. Your vaccination **protects you from death or severe illness**.

It is logical that people would be more interested in vaccination when the risk of disease is higher. Previous research has shown that parents' compliance with pediatric vaccine schedules increases in response to local pertussis outbreaks (Oster 2018; Schaller et al. 2019). This response is likely driven by perceptions of risk and anticipated regret, which psychologists have found increase vaccine uptake (Brewer et al. 2017). In the specific context of COVID-19, Shoji et al. (2020) show that knowledge of heightened local infection risk increases social distancing behavior. Moreover, a survey experiment by Thunstrom et al. (2020) finds that risk communication increases intention to vaccinate against COVID-19.[3] A message emphasizing the respondent's personal risk

[2] The full text of the experiment is presented in the online appendix; here I provide the messages only alongside explanations of their potential effectiveness.

[3] On the other hand, Akesson et al. (2020) find that perceptions of greater infectiousness of the virus lead to *lower* social distancing, apparently due to fatalism about the certainty of eventual infection.

Covid Economics 52, 15 October 2020: 39-67

may therefore increase intent to vaccinate as well.

The Risk-to-Self message emphasizes the dangers posed by the virus and avoided by vaccination. The *Safety-for-Self* message keeps the first two sentences of Risk-to-Self but frames the benefit of vaccination in terms of the safety it allows, not the risk it avoids:

> It is important to vaccinate against COVID-19 because it is **dangerous**. The virus can harm the lungs, heart, brain, and other body systems. Vaccination **gives you immunity** without illness. Your vaccination can keep you **healthy and safe**.

This message emphasizes a positive frame on the decision to vaccinate as well as the safety of vaccination itself. Positive and negative framing of otherwise equivalent public-goods games can lead to very different pro-social outcomes (Andreoni 1995), including in public goods games that explicitly model the decision to vaccinate (Sorensen 2019). Moreover, fear of adverse vaccine outcomes drives much of vaccine noncompliance, well out of proportion to the actual risks involved. For example, negative media coverage about the safety of the human papillomavirus (HPV) vaccine decreased uptake not only for that vaccine but for all vaccines among Danish girls (Gørtz et al. 2020). For these reasons, a message emphasizing safety to oneself could yield a very different response to one emphasizing risks of viral infection, even if the factual information conveyed is little different.

Instead of focusing on risks and benefits to individuals, successful messages may want to emphasize the prosocial role of vaccination. The benefit of a vaccine redounds partially to the patient, who becomes immune, and partially to society, which faces lower rates of community spread. There is reason to think that altruism is important for vaccination against a pandemic virus—Campos-Mercade et al. (2020) show for example that a pro-social orientation predicts a person's willingness to wear a face covering, practice social distancing, and to donate to a COVID-19 charity. On the other hand, qualitative research finds while that anti-vaccine activists understand the "free rider" dilemma of herd immunity, they dispute that they should consider this an important reason to vaccinate their children (Reich 2018, pp. 236–238).

The *Risk-to-Others* message emphasizes the dangers that low population immunity poses to other people. The first two sentences emphasize the role of immunity in preventing viral spread. The last two sentences are almost identical to those in the Risk-to-Self treatment, but modified to focus on the risk to other people:

> It is important to vaccinate against COVID-19 because it is **contagious**. When a person becomes immune, they will not spread the virus to others. While elderly and ill people are at extra risk, COVID-19 is potentially **deadly for everybody**. Your vaccination **protects others from death or severe illness**.

Previous research has found that information about risks to others elicits behaviors to end the pandemic. Sjåstad and Van Bavel (2020) find that perceptions of risk to others correlate with greater intent to perform pro-social pandemic behaviors. Also, Abel et al. (2020) find that while providing information about the recipient's risk to healthy young people does not affect donations to a COVID-19 charity, information about risk to elderly people increases giving substantially.

CEPR PRESS

Lastly, the *Safety-for-Others* message combines the first two sentences of the Risk-to-Others message with a modified version of the last two sentences of the Safety-for-Self message, to deliver a message that emphasizes the pro-social externality of vaccination, in a frame that emphasizes keeping others safe:

> It is important to vaccinate against COVID-19 because it is **contagious**. When a person becomes immune, they will not spread the virus to others. Vaccination **gives you immunity** without illness. Your vaccination can **keep other people healthy and safe**.

After viewing a randomly assigned message and expressing whether they intended to vaccinate against COVID-19, respondents were shown a follow-up question conditional on their answer. Respondents who indicated that they planned to get the vaccine were asked how quickly they planned to receive it.[4] Respondents who did not indicate that they planned to receive the vaccine were asked to indicate their reason for hesitancy from a checklist of possible answers.

2.2 Survey implementation

These messages were presented to a nationally representative survey panel, administered by the Understanding America Study (UAS), a program of the Center for Economic and Social Research at the University of Southern California. Respondents were randomized to five experimental arms. Four treatment arms viewed one of the messages; a control group saw no message. Each group was then asked whether they wanted the future COVID-19 vaccine. The survey identifier was UAS-299. The research plan was approved by the USC Internal Revenue Board with identifier UP-14-00148-AM080. Survey responses were collected from July 13 to July 21, 2020.

The UAS survey allows researchers to link anonymous respondents across survey waves. I oversample respondents who were previously unsure whether they wanted vaccination against COVID-19 using a question from March 2020 survey UAS-230, which asked respondents to estimate their probability of vaccinating. The sampling design set a goal of 2200 responses, with subgroup goals of 400 respondents with self-reported March vaccine probability of 0%, 1300 respondents between 1 and 99%, and 500 respondents at 100%. UAS panel participants who did not answer UAS-230 were not invited to the survey. Analyses were rebalanced using survey probability weights to construct a nationally representative sample.[5]

The UAS survey includes demographic information gathered as part of the design of the panel, including gender, age, race, Hispanic/Latino ethnicity, educational attainment, marital status, and US citizenship. Additionally, I add a variable for political partisanship by linking respondents to survey UAS-221, which asked respondents their registered political party in January 2020. This method identifies partisanship for all respondents except for 314 who did not participate in UAS-221.

[4] This follow-up question was motivated by worry that people who were marginal "yes" respondents might have low urgency, which would suggest the true effect on actual uptake might be overstated. However, there is no discernible difference in uptake urgency across treatment groups conditional on a yes answer.

[5] Additional regressions presented in the online appendix demonstrate that the findings presented below are qualitatively the same in an equally weighted regression.

Covid Economics 52, 15 October 2020: 39-67

Covid Economics 52, 15 October 2020: 39-67

2.3 Analytical strategy

The primary analytical method is a linear probability model where the outcome is coded as 1 for respondents who answer "yes" when asked whether they want the COVID-19 vaccine and 0 for answers of "no" or "maybe / I don't know." Treatment effects are estimated using a weighted least squares regression of the form

$$\text{Yes}_i = \alpha + \sum_{t=1}^{4} \gamma_t \times \mathbf{1}\left\{\text{Treatment}_i = t\right\} + \mathbf{X_i}\boldsymbol{\beta} + \varepsilon_i \tag{1}$$

where Yes_i is equal to 1 if respondent i answered "Yes" and 0 if i answered "No" or "Don't Know / Not Sure." The coefficients of interest, denoted γ_t, are the differences in the yes-share for treatment t relative to the control group. $\mathbf{X_i}$ is a vector of control variables for demographics, education, partisanship, and March 2020 vaccine intent included in some specifications. The constant term α equals the yes-rate for the control group in specifications without additional controls. The residual is denoted ε_i.

This research design is preregistered in the American Economic Association randomized trial registry. Its unique identifying number is: "AEARCTR-0006133." A power analysis conducted as part of the preregistration suggests that the sampling design will reject the no-effect null for a true effect of 30% of a standard deviation at the 5% statistical significance level 80% of the time after adjusting for multiple hypothesis testing.[6]

3 Data

The survey received 2,334 complete responses and 4 partial responses from 2,905 invitations, for a response rate of 80.34%. Total usable responses were 2,336. This exceeded the target sample of 2,200 due to a better-than-expected response rate.

Table 1 presents raw counts, weighted sample proportions, and intention-to-vaccinate rates for the sample. Average intention to vaccinate is worrisomely low, consistent with public polling and previous research (Thunstrom et al. 2020).[7] Weighting for a nationally representative sample, just over half of respondents definitely want to vaccinate against COVID-19. Roughly 20% do not want to vaccinate and 30% are unsure.

This represents a decline in vaccine intent since March 2020. Of people who stated that they were 100% sure they wanted a vaccine in March, only 72% answered "yes" in this survey. At the other extreme, just 4% of those who expressed a 0% chance of vaccination in March answered "yes" in July.

[6]This power calculation assumed that most respondents' probability of answering "yes" would be close to their stated probability in the March UAS survey. This is a conservative adjustment: a power calculation assuming each respondent would be an i.i.d. draw from a binomial distribution with 2200 identical and independent draws rejects the null more than 90% of the time for a treatment of 25% of a standard deviation.

[7]A Gallup poll administered over July 20 to August 2 found 35% of Americans would refuse a free COVID-19 vaccine. (https://news.gallup.com/poll/317018/one-three-americans-not-covid-vaccine.aspx)

CEPR PRESS

Table 1: Summary Statistics

Respondent Group	Unweighted Count	Weighted Share	Share "Yes"
Do you want the COVID-19 vaccine?			
Yes	1023	50.6%	
No	598	19.5	
Maybe / I don't know	715	29.9	
March 2020: Probability of vaccination			
Probability = 0%	394	8.0	4.4%
1-99%	1395	47.4	38.5
100%	547	44.6	71.6
Message experiment control group	466	20.8	46.6
Treatment group: Risk to Self	473	20.9	49.1
Safety for Self	469	18.8	52.6
Risk to Others	461	19.9	50.3
Safety for Others	467	19.5	54.8
Race is white only	1782	76.7	53.5
Black only	216	12.7	24.6
American/Alaskan Native only	56	0.9	19.2
Asian only	120	5.6	74.9
Hawaiian/Pacific Islander only	21	0.1	57.7
More than one race	123	3.3	52.1
Race not reported	18	0.7	42.6
Hispanic / Latino ethnicity	370	16.7	47.5
White, Non-Hispanic	1519	62.6	54.5
Non-white or Hispanic	817	37.4	44.0
January 2020: Political Party			
Democratic	692	30.0	57.2
Republican	635	24.5	47.0
Independent	415	16.9	56.1
Minor Party	25	1.0	30.6
Not in UAS 221 survey	314	15.7	48.5
Education is less than high school	107	7.2	35.2
High school	989	47.0	41.8
Associate's degree	353	11.0	46.3
Bachelor's degree	523	19.1	61.5
Graduate degree	364	15.7	73.5
Age 18-29	234	12.7	46.7
30-39	414	25.1	41.6
40-49	451	15.8	46.7
50-59	506	16.6	52.5
60-69	437	17.6	54.7
70+	293	12.3	69.6
Female	1418	51.6	44.4
Male	918	48.4	57.2
Married	1263	54.3	53.9
Not Married	1073	45.7	46.7
US citizen	2278	97.6	50.5
Non-US citizen	58	2.4	52.6

Columns report the unweighted counts and survey-weighted share of respondents in each category, and the weighted share of each answering "yes" to the vaccination question.

Covid Economics 52, 15 October 2020: 39-67

Covid Economics 52, 15 October 2020: 39-67

Average intention to vaccinate varies widely by race and ethnicity. 75% of Asian respondents express intention to vaccinate, but only 25% of Black respondents and 19% of Native Americans and Alaskan Natives. Latino and Hispanic respondents had higher intention to vaccinate than these groups (48%), but still below the sample average. This is especially alarming considering that Black, Native, and Latino Americans have had much higher COVID-19 population mortality than other groups. White, Hawaiian/Pacific Islander, and bi-/multiracial respondents were close to the sample average, with a bit over half of each expressing intention to vaccinate.

Partisanship is correllated with intention to vaccinate. Democrats and Independents both express intention to vaccinate at higher rates (57% and 56%) than the sample average, while 47% of Republicans intend to get the vaccine. Minor-party members, such as registered Libertarians and Greens, have a low intention to vaccinate (31%), although this is based on a small sample of just 25 respondents.

Intention to vaccinate increases in education, from 35% of respondents with less than a high school education, to 74% of respondents with graduate degrees. Intention to vaccinate by age is lowest for people in their 30s (42%) and increasing in age thereafter, to a nearly 70% intention among people 70 and older. Men are more likely to plan to vaccinate than women (57% versus 44%), and married people are more likely than single people (54% versus 47%). Non-US citizens and US citizens have similar intention to vaccinate.

4 Experimental Results

Before proceeding to the primary analysis, I display the distribution of responses by treatment in figure 1, panel A. The "yes" share is lowest for the control group, although the differences between each treatment's yes rate and the control are visibly small. The largest mean difference is for the Safety-for-Others message, where the yes-rate is higher than the control by just over 8 percentage points. Similarly, the "no" response rate is higher in the control arm than in any of the treatments, although again the differences are modest.[8]

Regressions for the full sample are reported in table 2. None of the four messages have a statistically significant effect on the full sample. Point estimates vary somewhat: the risk-focused messages have estimated effect on the yes-rate ranging from slightly less than zero to just under five percentage points across specifications, while the Safety-for-Self message has an effect of four to six percentage points, and the Safety-for-Others treatment has point estimates of seven to eight percentage points. However, none of these effect sizes are statistically different from the control group, nor are they of the magnitude required for rejection of the null under the preregistered power calculations. Further, an effect of eight percentage points on a baseline vaccination rate of

[8]Respondents are offered a "maybe" option to avoid a forced decision between "yes" and "no." I think it is unlikely that moving subjects from a "no" answer to uncertainty predicts much actual increase in vaccine uptake rates. An alternative specification presented in the online appendix where the dependent variable is equal to 0 for "no" answers, ½ for "maybe / I don't know" and 1 for "yes." Estimated treatment effects are little changed compared the binary outcome regressions.

Figure 1: Response Distribution by Treatment and Race/Ethnicity

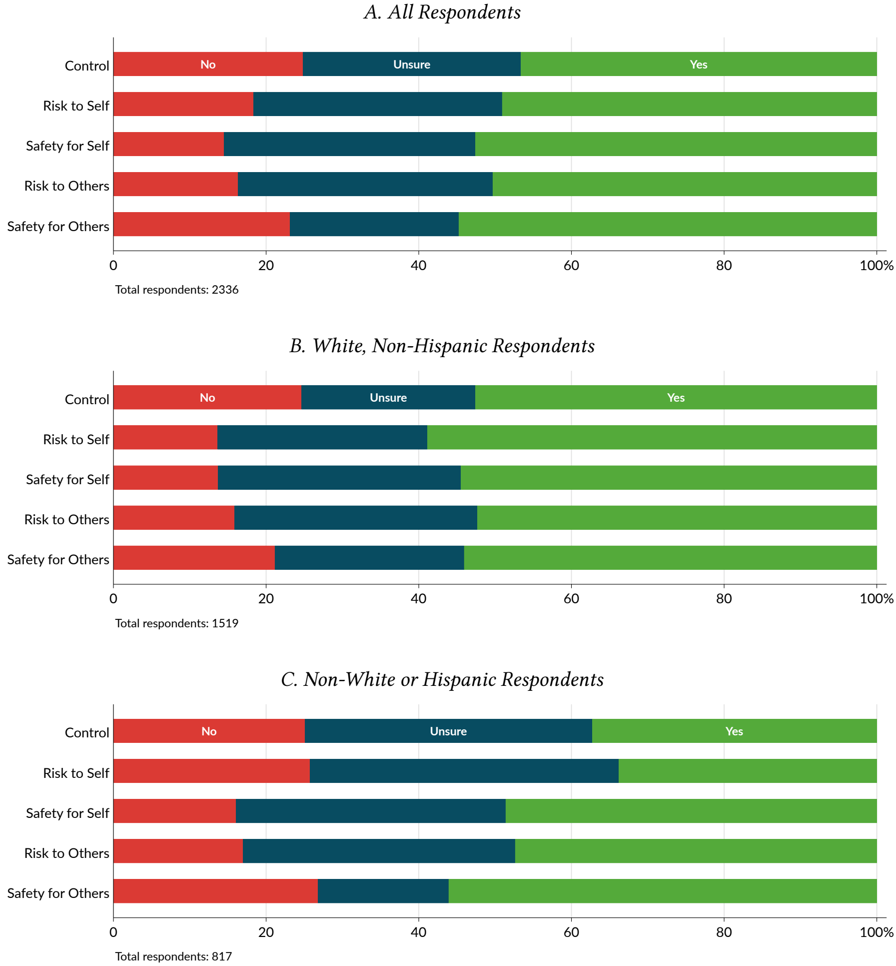

Notes: Panel A displays response by treatment for all respondents. Panel B displays proportions by treatment for white, non-Hispanic respondents, defined as persons whose race is classified as "white only" and who do not indicate Hispanic or Latino ethnicity. Panel C reports responses by treatment for all other respondents. Proportions are weighted by survey weights.

CEPR PRESS
Covid Economics 52, 15 October 2020: 39-67

47% would not be enough to achieve herd immunity.[9]

4.1 Treatment effects vary substantially by race and ethnicity

While the population-level effects are not statistically significant, they conceal important heterogeneity by race and ethnicity.

I divide the sample into two subgroups: one of respondents whose race is reported as white-only and of no Hispanic or Latino ethnicity, and one of respondents of Hispanic/Latino ethnicity or indicating a non-white racial identity.[10] I then reestimate the weighted least squares regression for these two subgroups. These estimates are reported in table 3.

The effects of messaging on white non-Hispanics is quite small. Four of the eight treatment point estimates are negative, and six of the eight estimates are between +2 and -2 percentage points. In column 2, where the addition of control variables improves the precision of the estimates, all 95% confidence intervals exclude effects larger than 10.1 percentage points.

Hispanic and non-white respondents have much larger estimated effects for three of the treatments. The Safety-for-Self and Risk-to-Others treatments have estimated effects over 10 percentage points, although these are not consistently different from zero at the 5% significance level. More strikingly, the Safety-for-Others message increases the share of Hispanic and nonwhite respondents who say "yes" by 19 to 22 percentage points, an estimate that *is* statistically different from zero at the 5% level in both specifications, both separately and in a joint test adjusted for multiple hypotheses.[11]

These differences are clearly visible if we replot the response distribution by race and ethnicity. In figure 2B, the "yes" share among non-Hispanic whites is almost indistinguishable across treatment arms, with the possible exception of the risk-to-self treatment, where it is slightly larger. Contrast this with the large differences in the "yes" rate by message in figure 2C, where the Safety-for-Others share of "yes" excluding non-Hispanic whites is much larger than for the control group's, primarily because of a much smaller share indicating that they are unsure about their response. The proportion of Safety-for-Others respondents who intend to get the vaccine (56.1%) is over 50% higher than in the control group (37.2%).

Regressions dividing the sample by political party, sex, and age do not reveal any notable heterogeneity on these dimensions. These additional regressions are reported in the online appendix.

[9] A commonly cited estimate of the herd immunity threshold for the United States is 70% of the population (Kwok et al. 2020).

[10] That is, respondents are placed in the second group if they are white and Hispanic/Latino, or if they identify as both white and another race.

[11] There are too few respondents assigned to each treatment arm to test effects on particular race and ethnicity groups with precision, but means comparison suggests that this effect is not driven by the responses of just one subgroup. The share of respondents who answer "yes" if they see the safety-for-others message, relative to the control, is 20.5%-points higher for Latino and Hispanic respondents, 18.4%-points higher for Black respondents, and 31%-points higher for Asian respondents; this compares to a difference of 1.5%-points among white, non-Hispanic respondents. The mean difference for Native Americans and Alaska Natives is -2.8%.

Covid Economics 52, 15 October 2020: 39-67

Table 2: Intention to Vaccinate by Treatment

	(1)	(2)	(3)	(4)
Risk to Self	0.025	0.009	0.047	0.032
	(0.048)	(0.042)	(0.043)	(0.040)
Safety for Self	0.060	0.050	0.053	0.046
	(0.049)	(0.043)	(0.043)	(0.040)
Risk to Others	0.037	0.017	0.014	-0.001
	(0.049)	(0.046)	(0.045)	(0.043)
Safety for Others	0.082	0.084	0.075	0.076
	(0.048)	(0.044)	(0.043)	(0.040)
Constant	0.466**	0.882**	-0.012	0.271*
	(0.034)	(0.121)	(0.045)	(0.124)
Demographic Controls		✓		✓
Vaccine Intent Controls			✓	✓
MHT-adjusted F-test p-value	0.497	0.305	0.413	0.288

Notes:
** – $p < 0.01$
* – $p < 0.05$
Each column reports a linear probability model regressing whether a respondent answered "Yes" regarding intention to vaccinate against COVID-19 (coded 1) or chose "No" or "Don't Know / Unsure" (coded 0). Demographic controls include sex, age, race, ethnicity, political partisanship, US citizenship, marriage, and education. Vaccine intent controls include the self-reported probability of vaccination as of March 2020 in a prior wave of the UAS survey, included as both a continuous variable and with indicators for probabilities equal to 0 or 1. A reported p-value at the bottom of the table is for a joint F-test that any of the four treatment variables are significantly different from zero. All regressions are weighted using survey probability weights.

Covid Economics 52, 15 October 2020: 39-67

Covid Economics 52, 15 October 2020: 39-67

Table 3: Intention to Vaccinate by Treatment and Race and Ethnicity

	White, Non-Hispanic		Non-white or Hispanic	
	(1)	(2)	(3)	(4)
Risk to Self	0.063	0.008	-0.035	0.062
	(0.057)	(0.048)	(0.078)	(0.068)
Safety for Self	0.019	0.001	0.113	0.159*
	(0.058)	(0.048)	(0.087)	(0.072)
Risk to Others	-0.002	-0.061	0.101	0.109
	(0.060)	(0.051)	(0.083)	(0.074)
Safety for Others	0.015	-0.005	0.188*	0.215**
	(0.058)	(0.048)	(0.083)	(0.070)
Constant	0.526**	0.120	0.373**	0.099
	(0.041)	(0.106)	(0.056)	(0.145)
Observations	1519	1519	817	817
Demographic Controls		✓		✓
Vaccine Intent Controls		✓		✓
MHT-adjusted F-test p-value	0.792	0.684	0.048	0.024

Notes:
** – $p < 0.01$
* – $p < 0.05$
Regressions are identical to those in table 2 except that demographic controls are limited to sex, age, political party, marriage, and education. The first two models are estimated only for white, non-Hispanic respondents, defined as persons whose race is classified as "white only" and who do not indicate Hispanic or Latino ethnicity. "Non-white or Hispanic respondents" are all respondents who are not white and non-Hispanic.

4.2 People are afraid the vaccine will not be safe

If respondents did not answer that they wanted to vaccinate against COVID-19, a follow-up question asked them to indicate the reason for their answer. Six prewritten reasons for vaccine refusal were offered, as well as space to offer a reason not listed. Respondents were allowed to choose more than one reason of refusal.

Table 4 reports the distribution of these responses by type of refusal and by race/ethnicity. The primary reason for vaccine hesitancy was safety concerns. Over half of the respondents who did not choose "yes" indicated worry that the COVID-19 vaccine would not be safe, including nearly 70% of respondents who answered "no." This concern was lower in white, non-Hispanic respondents (48%) than other respondents (55%).

Table 4: Reasons People Say Do Not Say "Yes"

	Vaccine Response			Race and Ethnicity	
	No	Maybe	Both	Non-Hispanic White	Hispanic or Non-white
I do not believe the COVID-19 vaccine will be safe	67.9%	39.3%	50.6%	47.5%	54.9%
I am not very concerned about the coronavirus	30.4	15.5	21.4	26.6	14.3
I expect the cost of the vaccine will be too high	10.4	25.3	19.4	15.9	24.2
I oppose all vaccines because of religious or personal beliefs	14.0	1.2	6.3	5.9	6.8
I have a health condition and cannot receive vaccines	5.6	3.7	4.4	5.2	3.4
I have had COVID-19 and I am already immune	3.4	1.1	2.0	1.0	3.4
Another reason not listed here	14.6	38.3	28.9	33.2	23.2

Columns report the share of respondents indicating reasons they did not select "yes" to the question asking whether they wanted to vaccinate against COVID-19. Respondents were permitted to indicate more than one reason and totals sum to more than 100 percent. Respondent shares are reported separately and pooled for the 598 "no" responses and 715 "Maybe / I don't know" responses. Proportions are weighted averages using survey weights.

Other important reasons for refusal or hesitancy included lack of concern about the coronavirus (21%) and worry that the vaccine would be too expensive (19%). These two reasons were noticeably different by race and ethnicity: lack of concern was almost twice as high among white, non-Hispanic respondents (27%) as others (14%), while cost concerns were much higher among non-white and Hispanic respondents (24%) than among non-Hispanic whites (16%). Respondents

Covid Economics 52, 15 October 2020: 39-67

CEPR PRESS

who were uninterested in the vaccine because of anti-vaccine sentiments generally, preexisting health conditions, or previous COVID infection each made up less than 7% of the non-yes responses.

A substantial number of the respondents, 29% of the hesitaters, indicated that they were hesitant for a reason not listed, and given space to write out their other reasons. Essentially none of these written comments suggested a reason overlooked by the prewritten responses. Instead, these responses either clarified an ambiguous response to particular prewritten answers (e.g. medical inability to take the vaccine only if it contains certain ingredients) or offering more detailed versions of the prewritten answers (e.g. three respondents wrote that they distrust the vaccine specifically because of conspiracy theories around the support of Bill Gates).

5 Discussion and conclusion

This study has important limitations. First, self-reported intention to vaccinate may not be predictive of actual uptake decisions (Leventhal et al. 1965; Falco and Zaccagni 2020). Yet, until an actual vaccine against COVID is available, uptake is hypothetical anyway. It is therefore prudent to design initial messaging based on self-reported intention-to-vaccinate surveys such as this one, and then reassess when actual uptake can be observed.

Second, the primary positive finding of this paper — that messages emphasizing the role of vaccines in safely protecting others is effective among non-white and Hispanic populations — is a post hoc finding.[12] A follow-up experiment designed to study this population specifically should be conducted to confirm and better understand this effect.

Third, this study is of a US population. The United States public health response to COVID, and its health system generally, are dissimilar from other countries' in important ways, and this may affect attitudes toward a COVID vaccine. A follow-up messaging experiment implemented in other countries might find a very different effect.

Despite these limitations, the effectiveness of this message for this group is plausible and, if valid, socially important. Black Americans in particular expect mistreatment and distrust claims of safety from the US health care system; such beliefs are predictive of dangerous underuse of health care resources (Alsan and Wanamaker 2017). Measures to persuade this group to vaccinate are also likely to have outsize importance given that infection and death rates from COVID-19 are much higher for Latinos, Native Americans, and Black Americans (Yancy 2020; Tai et al. 2020), and because vaccine hesitancy is highest for these groups (table 1). Such a strategy could potentially complement a "precision public health strategy" that implements targeted interventions for greater overall effectiveness (Rasmussen et al. 2020).

Messaging aside, this study has also shown that people of color disproportionately hesitate to vaccinate against COVID-19 because of concerns about vaccine safety and affordable access

[12]While the pre-registered analysis described in general terms an intent to estimate heterogeneous treatment effects by subgroups including race and ethnicity, gender, age, and partisanship, it left the exact implementation to be determined depending on the ultimate characteristics of the survey sample.

Covid Economics 52, 15 October 2020: 39-67

Covid Economics 52, 15 October 2020: 39-67

(table 4). Both of these reasons can be addressed with good public policy. Given the ongoing economic turmoil the pandemic has caused, programs to subsidize and widely distribute vaccine doses to willing but price-sensitive groups are likely to pass a cost-benefit analysis easily. And transparent safety testing of candidate vaccines, and effective messaging of that safety and the role of vaccines in protecting families and communities, is likely to be critical for encouraging the broadest possible uptake and protecting as many people as possible, particularly people of color.

Conflict of Interest Statement

The author has no conflicts of interest to disclose with respect to this research.

References

Abel, M., Byker, T., and Carpenter, J. (2020). Prosocial behavior in the time of Corona—the effect of role models. IZA Discussion Paper 13560.

Akesson, J., Ashworth-Hayes, S., Hahn, R., Metcalfe, R., and Rasooly, I. (2020). Fatalism, beliefs, and behaviors during the COVID-19 pandemic. NBER Working Paper 27245.

Alsan, M. and Wanamaker, M. (2017). Tuskegee and the Health of Black Men. *The Quarterly Journal of Economics*, 133(1):407–455.

Andreoni, J. (1995). Warm-Glow versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments. *The Quarterly Journal of Economics*, 110(1):1–21.

Brewer, N. T., Chapman, G. B., Rothman, A. J., Leask, J., and Kempe, A. (2017). Increasing vaccination: Putting psychological science into action. *Psychological Science in the Public Interest*, 18(3):149–207. PMID: 29611455.

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

Colgrove, J. (2006). *State of Immunity: The Politics of Vaccination in Twentieth-Century America*. University of California Press and Milbank Memorial Fund.

Falco, P. and Zaccagni, S. (2020). Promoting social distancing in a pandemic: Beyond the good intentions. *Covid Economics*, 28:159–195.

Gørtz, M., Brewer, N. T., Hansen, P. R., and Ejrnæs, M. (2020). The contagious nature of a vaccine scare: How the introduction of HPV vaccination lifted and eroded MMR vaccination in Denmark. *Vaccine*, 38(28):4432 – 4439.

Kolstad, J. T. and Kowalski, A. E. (2012). The impact of health care reform on hospital and preventive care: Evidence from Massachusetts. *Journal of Public Economics*, 96(11):909 – 929.

Kwok, K. O., Lai, F., Wei, W. I., Wong, S. Y. S., and Tang, J. W. (2020). Herd immunity — estimating the level required to halt the COVID-19 epidemics in affected countries. *Journal of Infection*, 80(6):e32–e33.

Leventhal, H., Singer, R., and Jones, S. (1965). Effects of fear and specificity of recommendation upon attitudes and behavior. *Journal of Personality and Social Psychology*, 2(1):20–29.

Opel, D., Heritage, J., Taylor, J., Mangione-Smith, R., Salas, H., DeVere, V., Zhou, C., and Robinson, J. (2013). The architecture of provider-parent vaccine discussions at health supervision visits. *Pediatrics*, 132(6):1037–1046.

Oster, E. (2018). Does disease cause vaccination? Disease outbreaks and vaccination response. *Journal of Health Economics*, 57:90 – 101.

Rasmussen, S. A., Khoury, M. J., and del Rio, C. (2020). Precision Public Health as a Key Tool in the COVID-19 Response. *JAMA*.

Reich, J. A. (2018). *Calling the shots: Why parents reject vaccines.* NYU Press.

Schaller, J., Schulkind, L., and Shapiro, T. (2019). Disease outbreaks, healthcare utilization, and on-time immunization in the first year of life. *Journal of Health Economics*, 67:102212.

Shoji, M., Cato, S., Iida, T., Ishida, K., Ito, A., and McElwain, K. (2020). COVID-19 and social distancing in the absence of legal enforcement: Survey evidence from Japan. MPRA Paper 100723.

Sjåstad, H. and Van Bavel, J. J. (2020). The best-case heuristic in risk prediction: Hopes and fears in a global health pandemic (COVID-19). PsyArXiv.

Sorensen, A. (2019). The effects of externality distribution and framing on individual vaccination decisions: Experimental evidence. *Economics Bulletin*, 39(4):2792–2812.

Tai, D. B. G., Shah, A., Doubeni, C. A., Sia, I. G., and Wieland, M. L. (2020). The Disproportionate Impact of COVID-19 on Racial and Ethnic Minorities in the United States. *Clinical Infectious Diseases*. ciaa815.

Thunstrom, L., Ashworth, M., Finnoff, D., and Newbold, S. (2020). Hesitancy towards a COVID-19 vaccine and prospects for herd immunity. *SSRN Electronic Journal*.

Troesken, W. (2015). *The Pox of Liberty: How the Constitution Left Americans Rich, Free, and Prone to Infection*. University of Chicago Press.

Yancy, C. W. (2020). COVID-19 and African Americans. *JAMA*, 323(19):1891–1892.

Covid Economics 52, 15 October 2020: 39-67

CEPR PRESS

Appendices

For Online Publication

Covid Economics 52, 15 October 2020: 39-67

A1 Additional Regressions

Table A1: Unweighted Regressions

	All Respondents		White, Non-Hispanic		Non-white or Hispanic	
	(1)	(2)	(3)	(4)	(5)	(6)
Risk to Self	0.032	0.043	0.010	0.009	0.071	0.095
	(0.032)	(0.028)	(0.040)	(0.034)	(0.054)	(0.049)
Safety for Self	0.068*	0.061*	0.024	0.008	0.149**	0.157**
	(0.032)	(0.028)	(0.040)	(0.033)	(0.055)	(0.050)
Risk to Others	0.037	0.022	0.004	-0.025	0.099	0.100*
	(0.033)	(0.028)	(0.041)	(0.035)	(0.053)	(0.048)
Safety for Others	0.068*	0.065*	0.033	0.015	0.130*	0.160**
	(0.032)	(0.028)	(0.041)	(0.034)	(0.054)	(0.049)
Constant	0.397**	0.261**	0.439**	0.137*	0.321**	-0.003
	(0.023)	(0.058)	(0.029)	(0.069)	(0.038)	(0.115)
Observations	2336	2336	1519	1519	817	817
Demographic Controls		✓		✓		✓
Vaccine Intent Controls		✓		✓		✓
MHT-adjusted F-test p-value	0.193	0.102	0.914	0.794	0.061	0.007

** – $p < 0.01$
* – $p < 0.05$
Regressions are identical to those in table 2 (columns 1 and 4) and table 3, except that estimation is unweighted, not weighted by survey probability weights.

Covid Economics 52, 15 October 2020: 39-67

Covid Economics 52, 15 October 2020: 39-67

Table A2: Vaccine Hesitancy Score Regressions

	All Respondents		White, Non-Hispanic		Non-white or Hispanic	
	(1)	(2)	(3)	(4)	(5)	(6)
Risk to Self	0.045	0.059*	0.086	0.050	-0.021	0.066
	(0.037)	(0.029)	(0.044)	(0.036)	(0.062)	(0.047)
Safety for Self	0.082*	0.073*	0.064	0.047	0.102	0.133*
	(0.037)	(0.029)	(0.044)	(0.035)	(0.064)	(0.053)
Risk to Others	0.061	0.028	0.043	-0.004	0.091	0.084
	(0.038)	(0.031)	(0.047)	(0.037)	(0.063)	(0.053)
Safety for Others	0.049	0.049	0.025	0.004	0.086	0.121*
	(0.039)	(0.031)	(0.047)	(0.038)	(0.069)	(0.055)
Constant	0.609**	0.427**	0.640**	0.265**	0.561**	0.308**
	(0.027)	(0.064)	(0.034)	(0.083)	(0.044)	(0.106)
Observations	2336	2336	1519	1519	817	817
Demographic Controls		✓		✓		✓
Vaccine Intent Controls		✓		✓		✓
MHT-adjusted F-test p-value	0.263	0.105	0.325	0.311	0.164	0.098

** – $p < 0.01$
* – $p < 0.05$
Regressions are identical to those in table 2 (columns 1 and 4) and table 3, except that the dependent variable is equal to one-half for respondents who chose "Maybe / I Don't Know." Responses of "Yes" are still coded 1 and "No" coded 0.

CEPR PRESS

Covid Economics 52, 15 October 2020: 39-67

Table A3: Intention to Vaccinate by Age Group

	Age 18 to 51		Age > 51	
	(1)	(2)	(3)	(4)
Risk to Self	0.056	0.040	-0.016	0.026
	(0.064)	(0.053)	(0.069)	(0.058)
Safety for Self	0.124	0.088	-0.027	-0.005
	(0.066)	(0.054)	(0.070)	(0.056)
Risk to Others	0.087	0.011	-0.024	-0.021
	(0.066)	(0.056)	(0.072)	(0.064)
Safety for Others	0.126	0.074	0.023	0.077
	(0.065)	(0.055)	(0.069)	(0.056)
Constant	0.370**	0.258	0.593**	0.103
	(0.045)	(0.135)	(0.049)	(0.228)
Observations	1206	1206	1130	1130
Demographic Controls		✓		✓
Vaccine Intent Controls		✓		✓
MHT-adjusted F-test p-value	0.271	0.412	0.951	0.467

** – $p < 0.01$
* – $p < 0.05$
The sample is divided into equally sized subsamples at median age 51. Regressions are identical to those in table 2, columns 1 and 4, except that demographic controls are limited to gender, race, ethnicity, citizenship, marriage, and education.

COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

CEPR PRESS

Covid Economics 52, 15 October 2020: 39-67

Table A4: Intention to Vaccinate by Gender

	Female		Not Female	
	(1)	(2)	(3)	(4)
Risk to Self	0.021	0.066	0.023	0.003
	(0.063)	(0.053)	(0.071)	(0.059)
Safety for Self	0.083	0.075	0.039	0.009
	(0.064)	(0.053)	(0.074)	(0.058)
Risk to Others	0.079	0.077	-0.018	-0.044
	(0.068)	(0.058)	(0.071)	(0.059)
Safety for Others	0.038	0.045	0.152*	0.105
	(0.064)	(0.053)	(0.070)	(0.059)
Constant	0.402**	0.122	0.538**	0.354*
	(0.044)	(0.167)	(0.051)	(0.167)
Observations	1418	1418	918	918
Demographic Controls		✓		✓
Vaccine Intent Controls		✓		✓
MHT-adjusted F-test p-value	0.658	0.59	0.107	0.138

** – $p < 0.01$
* – $p < 0.05$
Regressions are identical to those in table 2 except that demographic controls are limited to race, ethnicity, citizenship, age, marriage, and education.

COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

Covid Economics 52, 15 October 2020: 39-67

Table A5: Regressions by Party Registration

	Republicans		Democrats		Others	
	(1)	(2)	(3)	(4)	(5)	(6)
Risk to Self	-0.018	-0.015	0.009	-0.016	0.054	0.070
	(0.096)	(0.089)	(0.091)	(0.076)	(0.069)	(0.053)
Safety for Self	0.023	-0.024	0.035	0.039	0.094	0.077
	(0.096)	(0.084)	(0.093)	(0.068)	(0.072)	(0.059)
Risk to Others	0.059	-0.008	0.018	-0.034	0.018	0.004
	(0.098)	(0.077)	(0.090)	(0.079)	(0.074)	(0.065)
Safety for Others	0.086	0.062	0.128	0.112	0.047	0.072
	(0.097)	(0.077)	(0.090)	(0.071)	(0.070)	(0.060)
Constant	0.439**	0.403**	0.535**	0.440**	0.441**	0.313**
	(0.070)	(0.137)	(0.065)	(0.145)	(0.047)	(0.107)
Observations	635	635	692	692	1009	1009
Demographic Controls		✓		✓		✓
Vaccine Intent Controls		✓		✓		✓
MHT-adjusted F-test p-value	0.806	0.797	0.591	0.25	0.746	0.521

** – $p < 0.01$
* – $p < 0.05$
Regressions are identical to those in table 2 (columns 1 and 4) except that demographic controls exclude political party.

Covid Economics 52, 15 October 2020: 39-67

A2 Text of Survey Experiment

Question 1, Control Group:

About the COVID-19 Vaccine

This short survey will ask about your plans to vaccinate against the coronavirus that causes the COVID-19 illness. Right now, there is no vaccine against the coronavirus. However, one may be available in the coming months.

Do you want to get the COVID-19 vaccine?

◯ Yes

◯ No

◯ Maybe / I Don't Know

Choose One. If "Yes," go to question 2. If "No" or "Maybe / I Don't Know," go to question 3.

Covid Economics 52, 15 October 2020: 39-67

Question 1, Risk-to-Self Treatment Group:

About the COVID-19 Vaccine

This short survey will ask about your plans to vaccinate against the coronavirus that causes the COVID-19 illness. Right now, there is no vaccine against the coronavirus. However, one may be available in the coming months.

The following is a potential message to promote the vaccine when it becomes available:

> It is important to vaccinate against COVID-19 because it is **dangerous**. The virus can harm the lungs, heart, brain, and other body systems. While elderly and ill people are at extra risk, COVID-19 is potentially **deadly for anyone**. Your vaccination **protects you from death or severe illness**.

Please read the message above carefully before answering the survey.

Do you want to get the COVID-19 vaccine?

◯ Yes

◯ No

◯ Maybe / I Don't Know

Choose One. If "Yes," go to question 2. If "No" or "Maybe / I Don't Know," go to question 3.

CEPR PRESS

Question 1, Safe-for-Self Treatment Group:

About the COVID-19 Vaccine

This short survey will ask about your plans to vaccinate against the coronavirus that causes the COVID-19 illness. Right now, there is no vaccine against the coronavirus. However, one may be available in the coming months.

The following is a potential message to promote the vaccine when it becomes available:

> It is important to vaccinate against COVID-19 because it is **dangerous**. The virus can harm the lungs, heart, brain, and other body systems. Vaccination **gives you immunity** without illness. Your vaccination can **keep you healthy and safe**.

Please read the message above carefully before answering the survey.

Do you want to get the COVID-19 vaccine?

◯ Yes

◯ No

◯ Maybe / I Don't Know

Choose One. If "Yes," go to question 2. If "No" or "Maybe / I Don't Know," go to question 3.

Covid Economics 52, 15 October 2020: 39-67

COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

Covid Economics 52, 15 October 2020: 39-67

Question 1, Risk-to-Others Treatment Group:

About the COVID-19 Vaccine

This short survey will ask about your plans to vaccinate against the coronavirus that causes the COVID-19 illness. Right now, there is no vaccine against the coronavirus. However, one may be available in the coming months.

The following is a potential message to promote the vaccine when it becomes available:

> It is important to vaccinate against COVID-19 because it is **contagious**. When a person becomes immune, they will not spread the virus to others. While elderly and ill people are at extra risk, COVID-19 is potentially **deadly for everybody**. Your vaccination **protects others from death or severe illness**.

Please read the message above carefully before answering the survey.

Do you want to get the COVID-19 vaccine?

◯ Yes

◯ No

◯ Maybe / I Don't Know

Choose One. If "Yes," go to question 2. If "No" or "Maybe / I Don't Know," go to question 3.

CEPR PRESS
Covid Economics 52, 15 October 2020: 39-67

Question 1, Safe-for-Others Treatment Group:

About the COVID-19 Vaccine

This short survey will ask about your plans to vaccinate against the coronavirus that causes the COVID-19 illness. Right now, there is no vaccine against the coronavirus. However, one may be available in the coming months.

The following is a potential message to promote the vaccine when it becomes available:

> It is important to vaccinate against COVID-19 because it is **contagious**. When a person becomes immune, they will not spread the virus to others. Vaccination **gives you immunity** without illness. Your vaccination can **keep other people healthy and safe**.

Please read the message above carefully before answering the survey.

Do you want to get the COVID-19 vaccine?

◯ Yes

◯ No

◯ Maybe / I Don't Know

Choose One. If "Yes," go to question 2. If "No" or "Maybe / I Don't Know," go to question 3.

COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

Question 2, For "Yes" Answers to Question 1:

When the vaccine becomes available, how quickly will you want it?

◯ As soon as possible

◯ I want to wait before getting vaccinated

◯ Not sure / I don't know

Choose One. After answer, end survey

Covid Economics 52, 15 October 2020: 39-67

Question 3, For "No" and "Maybe" Answers to Question 1:

You indicated that do not want the COVID vaccine, or that you are not sure. Why not? Check all that apply

☐ I have had COVID-19 and I am already immune.

☐ I have a health condition and cannot receive vaccines.

☐ I oppose all vaccines because of religious or personal beliefs.

☐ I do not believe the COVID-19 vaccine will be safe.

☐ I expect the cost of the vaccine will be too high.

☐ I am not very concerned about the coronavirus.

☐ Another reason not listed here.

Choose at least one. After answer, end survey

Covid Economics 52, 15 October 2020: 39-67

Covid Economics Issue 52, 15 October 2020

The geography of pandemic containment[1]

Elisa Giannone,[2] Nuno Paixão[3] and Xinle Pang[4]

Date submitted: 10 October 2020; Date accepted: 10 October 2020

How does interconnectedness affect the pandemic? What are the optimal within and between states containment policies? We embed a spatial SIR model into a multi-sector quantitative trade model. We calibrate it to US states and find that interconnectedness increases the death toll by 73,200 lives. A local within-state containment policy minimizes welfare losses relative to a national policy or to one that reduces mobility between states. The optimal policy combines local within- and between-state restrictions and saves 132,200 lives. It includes a peak reduction in mobility of 33% saving approximately 40,000 lives. Different timing of policies across states is key to minimize losses. States like Arizona might have imposed too early internal lockdowns while too late travel restrictions.

1 This paper previously circulated under the title "Pandemic in an Interregional Model". First draft: April 2020. The views expressed in this paper are those of the authors and do not reflect those of the Bank of Canada.

2 Pennsylvania State University.
3 Bank of Canada.
4 Pennsylvania State University.

CEPR PRESS
Covid Economics 52, 15 October 2020: 68-95

1 Introduction

Interconnectedness through trade and mobility across states is a pillar of the Constitution of the United States. Exceptionally, the COVID-19 pandemic has challenged such long-standing paradigm. Some policy makers have advocated for the limitation of mobility of goods and individuals across states as a way to mitigate the pandemic. Concerns that interconnectedness exacerbates the diffusion of the disease and dampens the effects of containment policies are rising. Therefore, in this paper we consider interconnectedness for two reasons. First, to understand how it impacts the propagation of the disease and the economic activity in comparison to an one-region economy. Second, to study containment policies that restrict the movement of people and goods across states, resembling traveling restrictions or quarantines. In specific, it permits to analyze the extra benefit of between-state policies in saving lives and improving welfare.

This paper provides a quantitative multi-region framework with spatial infection diffusion to study the evolution of pandemics and related economic consequences. The model is calibrated to the US states using state level data on COVID-19 cases, inter-state trade flows and mobility of people across states through mobile phone tracking. Through the lens of the model, we analyze a battery of optimal containment policies imposed at different geographical levels. In particular, we analyze and compare nation-wide, state level and between-state policies.

The economic block of the model features two sectors, a regular consumption good and a social good sector. Each heterogeneous location produces a differentiated regular consumption good that is traded across cities, generating an economic link across locations.[1] The SIR block builds on Eichenbaum et al. (2020) that assumes that individuals internalize how their actions impact their own probability of getting infected leading to an endogenous change in consumption and labor supply even without mitigation policies. We depart from them in two dimensions. First, infection transmission is sector specific. The probability of getting infected through working and consuming in the social sector is higher than in the regular good sector. Second, we add a spatial component by assuming that agents in one state can be exposed to infected people in other states. The exposure across states is directly related to the size of trade flows and people's mobility. This framework allows us to highlight the role of interconnectedness on the spread of the disease and its impact on economic activity. Specifically, while mobility of goods and people favors economic activity, it simultaneously

[1]For simplicity, and given the short-term nature of the questions we are after, we assume that agents do not migrate or change sector. Although these assumptions can be easily relaxed, its inclusion would increase substantially computational complexity. Moreover, other frameworks as Giannone et al. (2020) are more suitable to analyze these decisions in a mid/long term horizon.

Covid Economics 52, 15 October 2020: 68-95

contributes to a faster spread of the disease, creating a tension between economic and health outcomes.

We present a set of positive and normative results. On the positive side, we find that the dynamics of the pandemic measured in terms of health and economic outcomes are more severe in a model with interconnectedness relative to one with isolated states. Total deaths as percentage of population are 0.0223p.p. higher, which corresponds approximately to 73,200 extra deaths. The peak drop in consumption is 6.2% in the model with connected states while 3.6% when we consider isolated states. These differences are substantially larger in states with lower initial infection and population and larger trade openness. In terms of welfare, we find that the welfare loss generated by the pandemic is 2.52p.p. larger in the economy with connected states. Another important features of our model is the behavioral response of agents that internalize how their actions impact their probability of getting infected. We find that a model that doesn't consider this behavioural response overestimates the total death toll by 0.01p.p., while the consumption peak drop is 4.44p.p. lower, which shows the importance of this feature in designing the optimal policy.

On the normative side, we study within-state optimal containment policies that resemble lockdowns. We differentiate between homogeneous (henceforce, national) and heterogeneous (henceforce, local) lockdowns across states. We also bring to the table a between-state mitigation policy that echoes traveling restrictions or quarantines. There are three main takeaways. First, local lockdown policies mitigate the pandemic more effectively than national ones. We highlight that the key factor determining the success of optimal local lockdown is *time* flexibility. The national lockdown would be imposed too early for small and low infection states as Arizona and too late for states with high population and infections as New York. Under both national and local optimal policies, lockdowns are almost exclusively imposed on the social sector. Second, a policy that restricts trade and mobility across states mitigates welfare losses but it doesn't reduce significantly the total death toll. This suggests that given the internal spread of the pandemic, limiting between-state mobility alone is not able to mitigate the pandemic. Third, combining local lockdowns and travel restrictions is the most effective policy. This policy would save 132,200 lives, which is approximately 32,000 more lives saved under a optimal local within-state lockdown.

This paper closely speaks to the fast growing literature of papers on COVID-19 that in the last few months contributed to understand the economic and health trade off of COVID-19 and optimal policy responses (e.g., Alvarez et al. 2020, Atkeson 2020, Atkeson et al. 2020, Eichenbaum et al. 2020, Faria-e-Castro 2020, Jones et al. 2020, Glover et al. 2020, Guerrieri et al. 2020, Piguillem et al. 2020). A few papers in this literature, like ours, have analyzed the spatial dimension of the COVID-19 crisis studying several economic and policy implications of the spread of the disease (among others, Antràs et al. 2020, Argente et al. 2020, Cuñat and Zymek 2020 and Fajgelbaum et al. 2020).

We contribute to the literature above in three dimensions. First, we develop a quantitative model of interregional trade and geographic mobility where agents internalize the impact of their actions in their own probability of getting infected. Second, through the lens of the calibrated model to US states, we study and compare optimal state-specific versus national containment policies. Third, we bring to the table the study of a between-state containment policy that could be interpreted in light of required quarantines and travelling restrictions that have been put in place in the recent months by several states.

Covid Economics 52, 15 October 2020: 68-95

2 Model

We build a two-sector quantitative trade model to study the role of interconnectness in the transmission of a pandemic. Agents internalize how their actions impact their own probability of getting infected and optimally choose consumption and labor supply. On the epidemiological side, we add an infection diffusion process across space and assume that each production sector has different infection's transmission rates.

2.1 Economic Environment

Space The economy is defined by L locations indexed by l. Every location produces a tradable differentiated regular consumption c and a non-tradable social good x. Locations differ in size, sector specific productivities and labor force distribution.

Preferences Prior to the pandemic, all agents across regions are identical and maximize a similar lifetime utility function:

$$U_l = \sum_{t=0}^{\infty} \beta^t u(c_{l,t}, x_{l,t}, n_{l,t}),$$

where the flow utility function is assumed to be:

$$u(c_l, x_l, n_l) = \log\left(\left(\phi^\rho c_l^{1-\rho} + (1-\phi)^\rho x_l^{1-\rho}\right)^{\frac{1}{1-\rho}} \right) - \gamma \frac{n_{l,t}^{1+\theta}}{1+\theta}.$$

$\beta \in (0,1)$ denotes the discount factor and $c_{l,t}$, $x_{l,t}$, $n_{l,t}$ denote the regular good's consumption, social good's consumption and hours worked, respectively. Regular good's consumption c_l is defined as a bundle of traded goods from different regions combined through the CES aggregator:

$$c_l = \left(\sum_{j=1}^{L} \alpha_{j,l} \tilde{c}_{j,l}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} \tag{1}$$

where $\epsilon > 0$ is the elasticity of substitution across products from different origins. $\tilde{c}_{j,l}$ denotes the consumption in region l of regular good produced in region j and $\alpha_{l,j}$ denotes the region specific measure of relative taste for goods of different regions. This introduces economic linkages across regions. A supply's disruption in one region imposes a utility cost elsewhere due to the lack of perfect sustainability across goods. Moreover, a negative income shock propagates across space due to lower demand.

Production Each location produces c and x according to the following CRS technologies:

$$Y_l^k = Z_l^k N_l^k \quad \text{for} \quad k = \{c, x\}.$$

where N_l^k is the labor demand in each sector. Labor cannot move across sectors and locations. Z_l^k is sector-location specific productivity.

Prices are region and sector specific, $\tilde{p}_{l,t}^c$ and $p_{l,t}^x$, respectively, for sector c and x. Wages and prices are fully flexible, but restrictions on labor mobility across sectors and regions induce a wage differential across sectors within region. In specific, perfect competition implies $w_{l,t}^c = Z_l^c \tilde{p}_{l,t}^c$ and $w_{l,t}^x = Z_l^x p_{l,t}^x$.

2.2 SIR with spatial diffusion

In this section, we augment the canonical SIR model with a spatial diffusion component similar to Gatto et al. (2020) and allow for economic decisions to have an impact on the probability of becoming infected. Given the heterogeneity across regions and social contact intensity across sectors, the probability of becoming infected is region-sector specific. It depends on the region's characteristics and increases with the intensity of the economic activity, the number of infected in the region and also the number of infected agents in other regions, especially those with stronger economic links.

We assume that agents are in one of the following health states: Susceptible, Infected, Recovered and Deceased. In a given region l, the total number of agents of sector k in these four groups are given by $S_{l,t}^k$, $I_{l,t}^k$, $R_{l,t}^k$ and $D_{l,t}^k$, respectively. Susceptible agents, those that haven't contracted the disease, may become infected by interacting with infected people. Infected people recover at rate π_r or die at rate π_d that are assumed to be common across sectors and regions. The evolution of the number of individuals in each state is given by the

Covid Economics 52, 15 October 2020: 68-95

following set of equations:[2]

$$
\begin{aligned}
S^k_{l,t+1} =& S^k_{l,t} - H^k_{l,t} \\
I^k_{l,t+1} =& I^k_{l,t} + H_{l,t} - (\pi_r + \pi_{l,d}) I^k_{l,t} \\
R^k_{l,t+1} =& R^k_{l,t} + \pi_r I^k_{l,t} \\
D^k_{l,t+1} =& D^k_{l,t} + \pi_{d,l} I^k_{l,t} \\
Pop^k_{l,t+1} =& Pop^k_{l,t} - D^k_{l,t}
\end{aligned}
$$

where the number of newly infected, $H^k_{l,t} = h^k_{l,t} S^k_{l,t}/Pop_{l,t}$, is given by the number of susceptible in each sector times the probability of becoming effect, $h^k_{l,t}$, which is defined as follows:[3]

$$
\begin{aligned}
h^k_{l,t} =& \pi_1 c^{s,k}_{l,t}(I_{l,t} C^i_{l,t}) + \pi_2 x^{s,k}_{l,t}(I_{l,t} X^i_{l,t}) + \pi_3 n^{k,s}_{l,t}(I^k_{l,t} N^{k,i}_{l,t} + 1_{(k=x)} I_{l,t} X^i_{l,t}) \\
&+ \pi_{4,l}\left(\gamma_{l,l} I_{l,t} + \sum_{j \neq l} (\gamma_{l,j} + \gamma_{j,l}) \frac{\tilde{c}_{l,j,t} + \tilde{c}_{j,l,t}}{\tilde{c}_{l,j} + \tilde{c}_{j,l}} I_{j,t}\right)
\end{aligned}
\tag{2}
$$

Susceptible people can contract the disease by meeting infected people while purchasing regular goods, consuming social goods, working or meeting infected people outside working and consumption activities. Following Eichenbaum et al. (2020), we assume that the probability of contacting people while purchasing goods is directly related with the shopping intensity and number of both infected and susceptible people. π_1 and π_2 relate with the probability of contracting the disease per encounter during shopping of regular and social goods, respectively.

Similarly, the likelihood of becoming infected while at work in the regular sector is proportional to the number of agents and hours worked by infected and susceptible. Agents in the the social, instead, besides interacting with co-workers, they are also exposed to potential infected clients. We then assume that the number of infections depends both on hours worked and amount of social goods consumed by infected agents, as a proxy for total number of potential interactions with infected clients. We assume that the probability of becoming infected in case of meeting one infected person at work, π_3, is the same in both sectors. But as workers in the social sector meet on average more people, the effective probability of contracting the virus is higher in the s sector.

The last component of equation 2 defines the infection spatial diffusion. As people move,

[2]The total population in a given sector-region declines with the number of the deceased.

[3]For simplicity, we denote the total number of agents in a given state in a region by $Y_{l,t} = Y^c_{l,t} + Y^x_{l,t}$ for $Y \in \{S, I, R, D, Pop\}$ and the total population in each sector is given by $Pop^k_{l,t} = S^k_{l,t} + I^k_{l,t} + R^k_{l,t}$. Aggregate consumption in each state $Y \in \{S, I, R\}$ is defined as $Y_{l,t} C^Y_{l,t} = Y^c_{l,t} C^{Y,c}_{l,t} + Y^x_{l,t} C^{Y,x}_{l,t}$ and $Y_{l,t} X^Y_{l,t} = Y^c_{l,t} X^{Y,c}_{l,t} + Y^x_{l,t} X^{Y,x}_{l,t}$.

Covid Economics 52, 15 October 2020: 68-95

susceptible people may be exposed to infected ones from different regions. We assume that the likelihood of meeting an infected person from another region is directly related to the fraction of population that moves and the level of economic linkage between the two regions. $\gamma_{l,j}$ in equation (2) is the average share of the population of region j present in region l prior to the beginning of the pandemic. Then, we interpreted $\gamma_{l,j} + \gamma_{l,j}$ as the movement of people across regions consistent with the steady-state gross trade flows $\tilde{c}_{l,j} + \tilde{c}_{j,l}$. Therefore, the expected number of infected people from region j that a susceptible person in region l may meet is given by $(\gamma_{l,j} + \gamma_{j,l})\frac{\tilde{c}_{l,j,t}+\tilde{c}_{j,l,t}}{\tilde{c}_{l,j}+\tilde{c}_{j,l}} I_{j,t}$ that decreases as trade-flows fall. π_4 reflects the probability of becoming infected conditional on randomly meeting someone infected.

2.3 Optimization

Mobility frictions across locations and sectors and the absence of any insurance mechanism against the risk of infection make the budget constraint location-sector-health specific. We assume that the budget constraint of an agent in region l, sector k and health status $a \in \{s, i, r\}$ is:

$$(1 + \tau^c_{l,t})p_{l,t}c_{l,t} + (1 + \tau^x_{l,t})p^x_{l,t}x^{k,a}_{l,t} = w^k_{l,t}\nu^a n^{k,a}_{l,t} + T^{k,a}_{l,t} \tag{3}$$

where $(1+\tau^c_{l,t})p_{l,t}c_{l,t}$ denotes the total cost of purchasing aggregate regular good $c_{l,t}$ in location l and is defined as

$$(1 + \tau^c_{l,t})p^c_{l,t}c_{l,t} = \sum_{j=1}^{L}(1 + \tau^c_{l,j,t})\tilde{p}_{j,t}\tilde{c}_{j,l,t}$$

ν^a determines the effective hours worked for different health states. We set $\nu = 1$ for the susceptible and recovered people and $\nu < 1$ for infected people. $\tau^x_{l,t}$ is the consumption tax on social good and $\tau^c_{l,j,t}$ is the tax rate in state l of goods from region j. $T^{k,a}_{l,t}$ are location-sector specific transfers. We assume that government runs a balance budget every period and rebates the revenues generated in location-sector to the workers of the same location-sector. Taxes on foreign goods are rebated for both sectors in the state.[4]

Agents face a dynamic problem during the pandemic as their consumption and labor decisions impact the future probability of becoming infected.[5] In case of becoming infected the agent faces two consequences. First, she has lower labor productivity which translates in

[4] Rebating foreign taxes solely to sector c underperforms in terms of reducing welfare losses.

[5] Although total regular consumption c, social consumption x and total hours worked n are chosen taking into consideration the dynamic component of the problem, the allocation of the consumption of c across goods produced in different locations is purely a static problem. Given the consumption aggregator defined in (1), any agent in region l at time t in sector k and health status s demands from region j: $\tilde{c}_{j,l,t} = \left(\frac{(1+\tau^c_{j,l,t})\tilde{p}_{j,t}}{\alpha_{j,l,t}(1+\tau^c_{l,t})p^c_{l,t}}\right)^{-\epsilon} c_{l,t}$. The price level for c-sector goods in city l and given by $(1 + \tau^c_{l,t})p^c_{l,t} = \left[\sum_{j=1}^{L} \alpha_{j,l,t}{}^{\epsilon}\left((1 + \tau^c_{j,l,t})\tilde{p}_{j,t}\right)^{1-\epsilon}\right]^{\frac{1}{1-\epsilon}}$.

Covid Economics 52, 15 October 2020: 68-95

less effective hours of work and, therefore, income. Second, in case of becoming infected she faces a positive probability of death and, therefore, forgone utility.

Susceptible People A susceptible person s in location l in sector $k \in \{c, x\}$ chooses consumption $c_l^{k,s}$ and $x_l^{k,s}$ and hours worked $n_l^{k,s}$ that solve the following optimization problem:

$$U_{l,t}^{k,s} = \max_{\{c_{l,t}^{k,s}, x_{l,t}^{k,s}, n_{l,t}^{k,s}\}} u(c_{l,t}^{k,s}, x_{l,t}^{k,s}, n_{l,t}^{k,s}) + \beta \left[(1-\tau_{l,t}^{k})U_{l,t+1}^{k,s} + \tau_{l,t}^{k} U_{l,t+1}^{k,i}, \right] \quad \text{s.t.} \quad (3) \tag{4}$$

where, $\tau_{l,t}^{k}$, the probability of becoming infected is defined in equation (2). We assume that susceptible people take as given aggregate variables, but understand how their consumption and working decisions impact their own probability of becoming effect. However, they don't internalize how their decisions impact the aggregate variables, giving origin to an infection externality.

Infected People We implicitly assume that the cost of death is the foregone utility of life and that infected people don't take into consideration that they may infect others. Therefore, infected people solve the following problem:

$$U_{l,t}^{k,i} = \max_{\{c_{l,t}^{k,i}, x_{l,t}^{k,i}, n_{l,t}^{k,i}\}} u(c_{l,t}^{k,i}, x_{l,t}^{k,i}, n_{l,t}^{k,i}) + \beta \left[(1-\pi_r - \pi_{d,l})U_{l,t+1}^{k,i} + \pi_r U_{l,t+1}^{k,r}, \right] \quad \text{s.t.} \quad (3) \tag{5}$$

Recovered People Similarly to infected people, the decisions of the recovered people are also static and satisfy the following problem:[6]

$$U_{l,t}^{k,r} = \max_{\{c_{l,t}^{k,r}, x_{l,t}^{k,r}, n_{l,t}^{k,r}\}} u(c_{l,t}^{k,r}, x_{l,t}^{k,r}, n_{l,t}^{k,r}) + \beta U_{l,t+1}^{k,r} \quad \text{s.t.} \quad (3)$$

2.4 Equilibrium Definition

Given the initial labor allocations across sectors and space, $\left\{N_l^k\right\}_{l=\{1,..,L\}}^{k=\{c,x\}}$ and a sequence of taxes and transfers, $\left\{\tau_{l,t}^c, \tau_{l,t}^x, T_{l,t}^c, T_{l,t}^x\right\}_{t=\{1,..,\infty\}}^{l=\{1,..,L\}}$, the equilibrium consists of a set of prices $\{\tilde{p}_{l,t}^c, p_{l,t}^x, w_{l,t}^c, w_{l,t}^x\}_{t=1}^{\infty}$ and allocations $\{c_{l,t}^{k,r}, x_{l,t}^{k,r}, n_{l,t}^{k,r}\}_{t=1}^{\infty}$ for each sector $k \in \{c, x\}$ and region $l \in \{1, ..., L\}$ that solve the agents' maximization problems and satisfy the goods and labor markets clearing conditions defined as:

[6]The solutions to agents' problem are contained in the Online Appendix.

Covid Economics 52, 15 October 2020: 68-95

$$\sum_{k\in\{c,x\}} S^k_{l,t}X^{k,s}_{l,t} + I^k_{l,t}X^{k,i}_{l,t} + R^k_{l,t}X^{k,r}_{l,t} = Y^x_{l,t}$$
$$\sum_{j\in\{1,\ldots,L\}}\sum_{k\in\{c,x\}} S^k_{l,t}\tilde{C}^{l,s}_{l,j,t} + I^k_{l,t}\tilde{C}^{l,i}_{l,j,t} + R^k_{l,t}\tilde{C}^{l,r}_{l,j,t} = Y^c_{l,t}$$
$$S^k_{l,t}\nu^s N^{k,s}_{l,t} + I^k_{l,t}\nu^s N^{k,i}_{l,t} + R^k_{l,t}\nu^s N^{k,r}_{l,t} = N^k_{l,t}, \quad \text{for} \quad k \in \{c,x\}$$

3 Taking the Model to the Data

3.1 Parameters Values

We calibrate the model at weekly frequency and to the pre-pandemic U.S states' characteristics. The decision to make a state-specific model is driven by the fact that most containment policies, such as lockdowns and quarantine, are implemented at state-level.

In the Online Appendix, we describe in detail the full calibration. Here we restrict our attention to the parameters related to the spatial and SIR components. In specific, we set the elasticity of substitution across states, ϵ, to 5 as estimated by Ramondo et al. (2016). The relative taste for goods of different states, α's, are chosen to match the share of imported goods from each state, using shipments data between-states from the Commodity Flow Survey. We measure the share of people moving across states pre-pandemic, γ, using cell phone data from Couture et al. (2020). This data reports among the smartphones that pinged in a given state in a certain day, the share of those devices that pinged in each of the other 50 state at least once during the previous 14 days.

To calibrate the SIR parameters, we use a similar approach as in Eichenbaum et al. (2020). We match the probability of death of 0.5% and assume that 18 is average number of days to recover or die. π_1, π_2, π_3 and $\pi_{4,l}$ in equation (2) are jointly estimated to match different transmission rates across activities. Using data from the Time Use Survey, we find that 18% and 30% of the time spent on general community is used for the purchase of "goods and services" and "eating and drinking outside", respectively. According to Ferguson et al. (2006) 33% of virus transmission are likely occur in the general community, thus, we set the average number of infections originated by consumption of c to 6% (0.33×0.18) and those originate by the consumption of x to 10% (0.33×0.3). 17% of infections occur in the work place with the largest share occurring in the social sector. Fernandez-Villaverde and Jones (2020) shows a wide heterogeneity in the basic reproduction number, R_0, across states at the beginning of

Covid Economics 52, 15 October 2020: 68-95

the pandemic. We calibrate the four parameters above to match such state-level estimates.[7]

Finally, we take into consideration the heterogeneity in the number and timing of the first cases in each state. In specific, we initialize the pandemic in the model with 0.1% of the population infected in New York, Massachusetts and District of Columbia,[8] while in the other states we initiate the model with 0.01% of population infected. These three states had the highest number of reported cases in the early stage of the pandemic.

3.2 Understanding The Model's Mechanisms

In this section we highlight the main mechanisms at play in our model and the role of interconnectedness.[9] In the first part we show a large degree of heterogeneity across states in health and economic outcomes generated by the pandemic. The top panel of figure 1 presents a map of cumulative infections and deaths. We find that the most affected states are hit 4 times more than the least ones, with most affected states concentrated in the Northeast. States with a larger number of cases and deaths per capita have, on average, higher levels of population, R_0 and openness.[10]

The bottom panel of figure 1 presents the results for economic outcomes such as hours worked and consumption. We find that the most affected states had a decline in labor and consumption around 4 to 5 times larger than the least impacted ones.[11] States with a larger drop of labor supply and consumption have, on average, higher levels of population and R_0 but lower openness. Finally, we find only a mild positive relationship between health and economic outcomes. While Northeast states face a large number of cases and large drop in economic conditions, states like California, Texas and Florida see large decline in labor and consumption despite relative low levels of infections. This is mostly due to the endogenous response that makes individuals stop consuming good from the sector when the pandemic hits even when cases are low.[12] This analysis emphasizes the large degree of spatial heterogeneity in the pandemic outcomes and points in the direction of state-specific interventions.

We now analyze the dynamics of health and economic outcomes and show that interconnectedness plays an important role in the evolution of the pandemic. Figure 2 shows the

[7]Any ratio between π_1, π_2, π_3 are common across states and $\pi_{4,l}$ varies across states to match the state-level R_0 estimates from Fernandez-Villaverde and Jones (2020).

[8]Eichenbaum et al. (2020) also initializes the model with 0.1% of the population infected.

[9]In this section we assume no policy intervention, $\tau^c_{l,t} = \tau^x_{l,t} = 0$ for any l and t.

[10]State openness refers both to trade and people's mobility. We defined it as $(\gamma_{l,j} + \gamma_{j,l})\frac{\tilde{c}_{l,j,t}+\tilde{c}_{j,l,t}}{\text{income}_l}$.

[11]We exclude District of Columbia from this calculation since it is a strong outlier. DC is a degree of openness that is 5 times larger than the second one.

[12]The Online Appendix reports the correlations between-states' characteristics and deaths as well as cumulative consumption drop.

Covid Economics 52, 15 October 2020: 68-95

Figure 1: Geographic Heterogeneity in Pandemic Impact

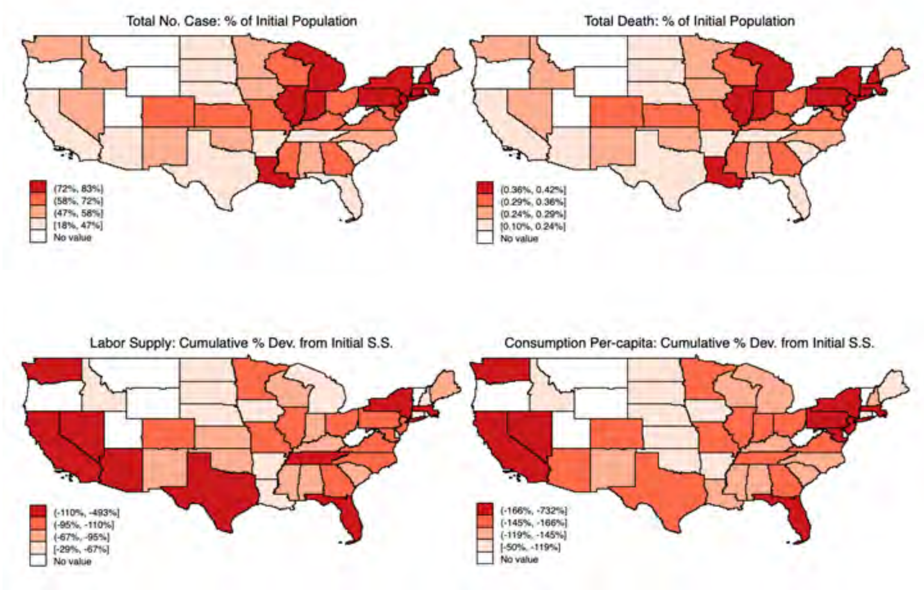

number of cumulative cases and deaths as percentage of initial population, as well as labor supply and consumption per-capita in percentage deviation from the initial steady-state. We report these outcome for the “Baseline" economy (solid red line) and for an economy without trade and geographic mobility, denoted by “Isolated Economy" (dashed blue line). The top panel reports the results for the evolution of infections at aggregate level in column one, and for New York, Ohio and Arizona in the second, third, fourth columns, respectively. These three states represent extreme cases of high, medium and low initial infections level and population. In row one, we find a 2p.p. difference in the pick of infection rates between the baseline economy and the isolated one. Besides generating more deaths, interconnectedness anticipates the peak of infections. By observing the plots for the three states separately, we find that the largest differences are generated by Ohio and Arizona, while there are nearly zero differences between models for New York. Similarly, in the second row, we show the evolution of deaths over time. The baseline economy produces approximately 73,200 extra deaths more the isolated one. By analyzing the three graphs on the right, we find that the largest overall death toll occurs in New York with similar values under both economies. Instead, there are larger differences for Ohio and Arizona, where interconnectedness generates

Covid Economics 52, 15 October 2020: 68-95

0.019 and 0.043p.p. more deaths per-capita, respectively. Overall, interconnectedness impacts relatively more states with lower R_0 and smaller population as these states, like Arizona, import relatively more infections per-capita.

Figure 2: Health and Economic Outcomes of COVID-19 Crisis

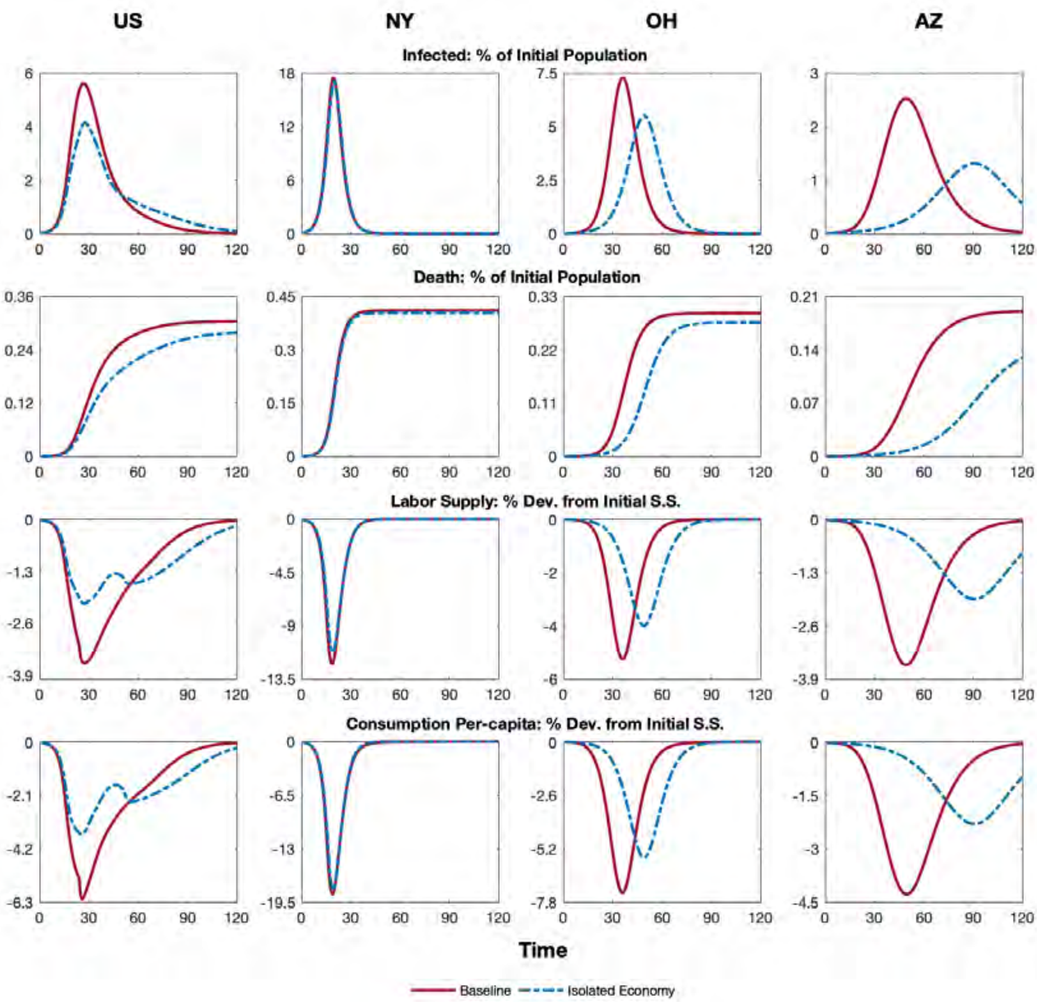

Rows three and four of figure 2 report the evolution of labor and consumption per-capita in percentage deviation from the pre-pandemic steady-state, respectively. Individuals voluntarily contract consumption and labor supply as the virus spread in order to mitigate the probability of becoming infected. We find that labor supply and consumption drop the most around the

Covid Economics 52, 15 October 2020: 68-95

Covid Economics 52, 15 October 2020: 68-95

time of the infection peak. These drops in labor supply and consumption are, respectively, 1.55p.p. and 2.55p.p. smaller in the economy where states are not connected. At state level, we observe that New York displays the largest drop in labor supply and consumption, followed by Ohio and Arizona. In New York we find almost no difference between an isolated and interconnected economy while the largest differences are displayed for Arizona where in the connected case, the drop in labor supply and consumption occurs earlier and at a greater extent than in an isolated economy.

3.3 The Geography of Optimal Containment Policies

In this section, we analyze and compare within-state and between-state containment policies. Since agents are atomistic, they don't internalize the impact of their behavior on the disease transmission. Therefore, the competitive equilibrium is not Pareto Optimal and there is room for government intervention. The social planner maximizes the social welfare in the entire country by imposing a set of tax instruments that constrain economic activity and the disease dynamics. The social planner can chooses a sequence of consumption tax rates that can be sector and state specific. In specific, for each state l, it can tax consumption of social goods, $\tau^s_{l,t}$, own-state regular good, $\tau^c_{l,t}$, and for regular goods imported from each of the other states, $\left\{\tilde{\tau}^c_{l,j,t}\right\}_{j\neq l}$ for T periods.[13] The aggregate social welfare, U_0, is defined as a weighted average of the lifetime utility of the different agents in each health status:[14]

$$U_0 = \sum_{l=1}^{l=L} \left[S^c_{l,0} U^{s,c}_{l,0} + S^x_{l,0} U^{s,c}_{l,0} + I^c_{l,0} U^{i,c}_{l,0} + I^x_{l,0} U^{i,x}_{l,0} \right]$$

$U^{s,k}_{l,0}$ and $U^{i,k}_{l,0}$ are the lifetime utility at time 0 (beginning of the pandemic) of susceptible and infected, respectively, in each state l and sector k. Those are the solution to optimization problems (4) and (5) given the sequence of tax rates imposed by the government.[15]

If we consider $T = 250$, it would imply a choice of 451,500 parameters, which is computationally very challenging. Therefore, we approximate the optimal time paths by a generalized logistic function of time:[16]

$$\tau(t) = \kappa_1 \frac{\kappa_2 \kappa_3 e^{\kappa_3 (t-\kappa_4)}}{\left[1 + e^{\kappa_3 (t-\kappa_4)}\right]^{1+\kappa_2}}$$

κ_1 determines the highest level of the mitigation and κ_2 its persistence, κ_3 controls mitigation in the earlier periods and κ_4 determines the period with the highest mitigation policy.

[13] After period T all rates are set to 0.

[14] Note that at the initial period, there are no deaths or recovered people.

[15] To solve this Ramsey problem, we guess a sequence of tax rates and solve for the competitive equilibrium. We then evaluate the social welfare function and iterate on this sequence until we find the optimum.

[16] We consider alternative functional forms, but they under performed compared to the one selected.

Covid Economics 52, 15 October 2020: 68-95

Figure 3: Optimal Within-State Containment Policy

Panel A: Spatial Comparison

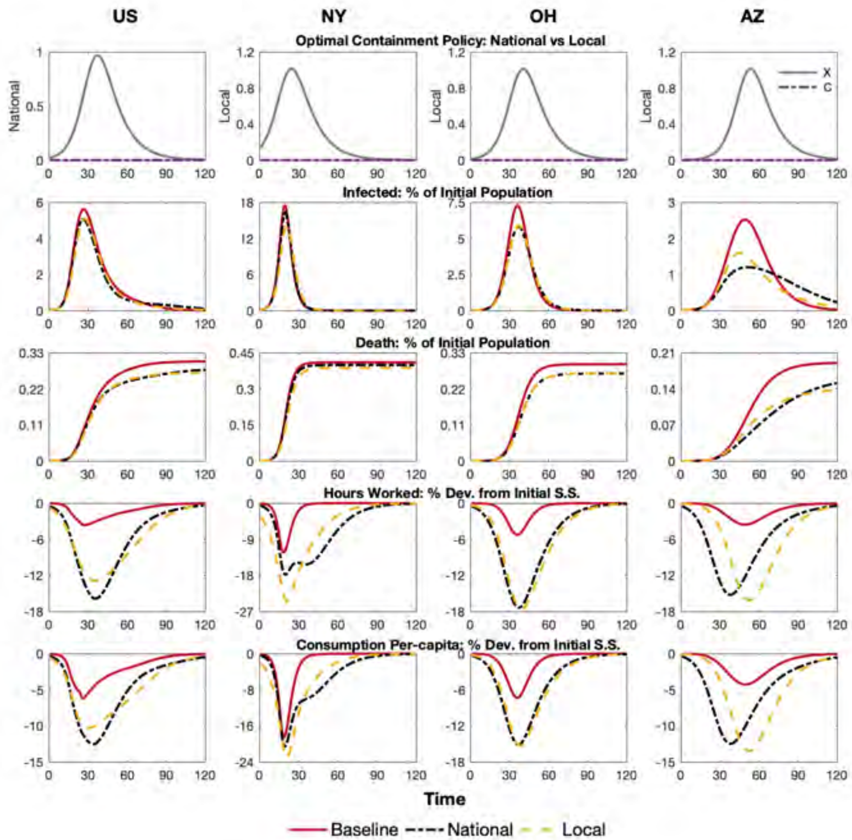

Panel B: Timeline Comparison

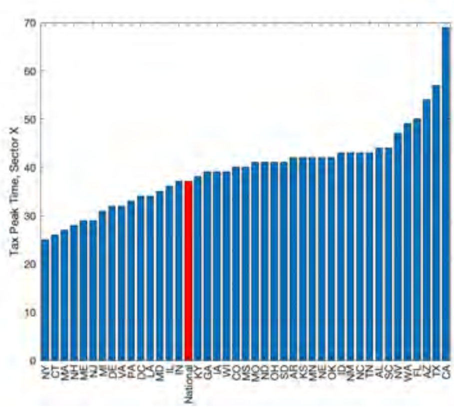

Below we study and compare different optimal containment policies with different characteristics. We first consider policies that focus on within-state consumption behavior, denominated within-state policies. Second, we study a policy that targets trade flows across regions, called between-state policy. We consider both the cases in which policies are equally implemented everywhere, national policies, and policies that vary across states, local policies. At last, we look at the optimal policy that combines both within and between-state policies.

Optimal Within-State Containment Policy Figure 3 shows the tax path and the evolution of health and economic outcomes under national and local optimal within-state containment policies. The local policy consists of state specific consumption taxes on social and regular consumption good. Regular goods are taxed equally regardless of their origin. The national policy imposes the same tax rate in all the states, but it can vary across sectors. The top left panel shows the optimal national tax rates for sectors X and C. The other three plots of the first row show the local tax rates for New York, Ohio and Arizona, respectively. First, we find that taxing c is not optimal under both policies. Second, perhaps surprisingly, the maximum tax rate of the social sector is the same in all the states, regardless of their total death toll.

The second and third rows report the evolution of infections and cumulative deaths. We find lower infection and death rates under both optimal policies than in the baseline, but a local policy is able to save roughly 38% more lives than the national one. Overall, we find that the the local policy would reduce in 100,500 the number of overall deaths. When we take a closer look at the different states, we observe that there are small difference between national and local policies for New York but quite large for Arizona. The largest improvement in death rates by imposing local or national policies would come from states with smaller population and lower initial reproduction number, R_0, as Ohio and Arizona. In New York, despite the fact that the optimal policy reduces slightly the infection peak, the high initial R_0 prevents policies to significantly reduce death rate.

Rows four and five report the evolution of labor and consumption, respectively. Both policies generate a larger drop in hours worked and consumption than in the baseline economy. At the aggregate level, national policy amplifies the peak drop in hours worked and consumption by 2.97p.p. and 2.26p.p. relative to the local policies. The differential impact is more pronounced in states like New York and Arizona, which is mainly explained by the timing of the different optimal local policies.

As previously mentioned the maximum local tax rate is the same in all the states, however, the maximum tax level is reached at different time across states. Panel B of figure 3 compares the timelines of the peak of the optimal policies across states. This highlights that time is a key margin through which state-specific optimal policies operate. In specific, local tax rates

Covid Economics 52, 15 October 2020: 68-95

Covid Economics 52, 15 October 2020: 68-95

on x follow closely the evolution of cases in each state. While cases are low, tax rate is low and increases as the number of cases and deaths go up. The maximum value of the optimal policy occurs when the state reaches their peak. While the national policy has a pick at approximately 30 weeks, New York under a local policy would have taxed the most around week 20 while Texas and Arizona around week 60, implying a gap of approximately 40 weeks.

Therefore, a homogeneous policy across states would imposed a lockdown too late in some states and too early in others. This result stresses that a premature lockdown can be economically very costly with little benefits in reducing the death toll.

Optimal Between-State Containment Policy We now study the optimal containment policy that restricts the movement of goods and people across states. This policy consists of taxing goods from other states, which translates into lower trade flows, lower mobility of individuals and lower spatial infection diffusion. The blue line in figure 4 reports the evolution of health and economic outcomes under this between-state tax rate. The red line reports the baseline economy and the yellow line the outcomes under the optimal local tax policy previously analyzed.[17]

The middle graph of the first row plots the optimal between-state containment policy. The social planner finds optimal to tax foreign goods but at a much lower rate than service goods under the within-state policy. As the optimal local tax on service goods follows closely the infection cases at state level, this local between-state policy does the same.

As reported in row three, the optimal between-state tax rate induces a higher death toll than the local within-state policy. A between-state tax alone reduces the overall death toll as percentage of initial population from 0.3 to 0.287p.p., but above the 0.263p.p. deaths generated by the local within-state policy. This policy brings relatively larger gains to Arizona, the most affected state by interconnectedness among these three. This occurs because between-state tax directly impacts the degree of openness of the state, as this tax reduces foreign demand, trade-flows and consequently movement of people and infection diffusion across space. In specific, it induces a reduction in cross-state mobility of 57% at the peak and 9% on average during the pandemic period. These numbers compare with a reduction of mobility of 9% at the peak and 1% on average in the absence of any policy.

This policy generates less economic losses per life saved as hours worked and consumption decline substantially less than under the within-state policy. However, this policy alone does not have capacity to save as many lives as other policies. This policy targets movement of goods and people across regions, but disease spreads within-state even if borders are completely closed. Although reduction in trade flows attenuates infection diffusion internally,

[17] The green line reports the optimal policy combining local consumption tax and between-state containment policies. This policy is discussed later.

a policy that does not consider the social good sector faces limitations in the number of deaths that can avoid.

Figure 4: Optimal Between-State and Overall Policy

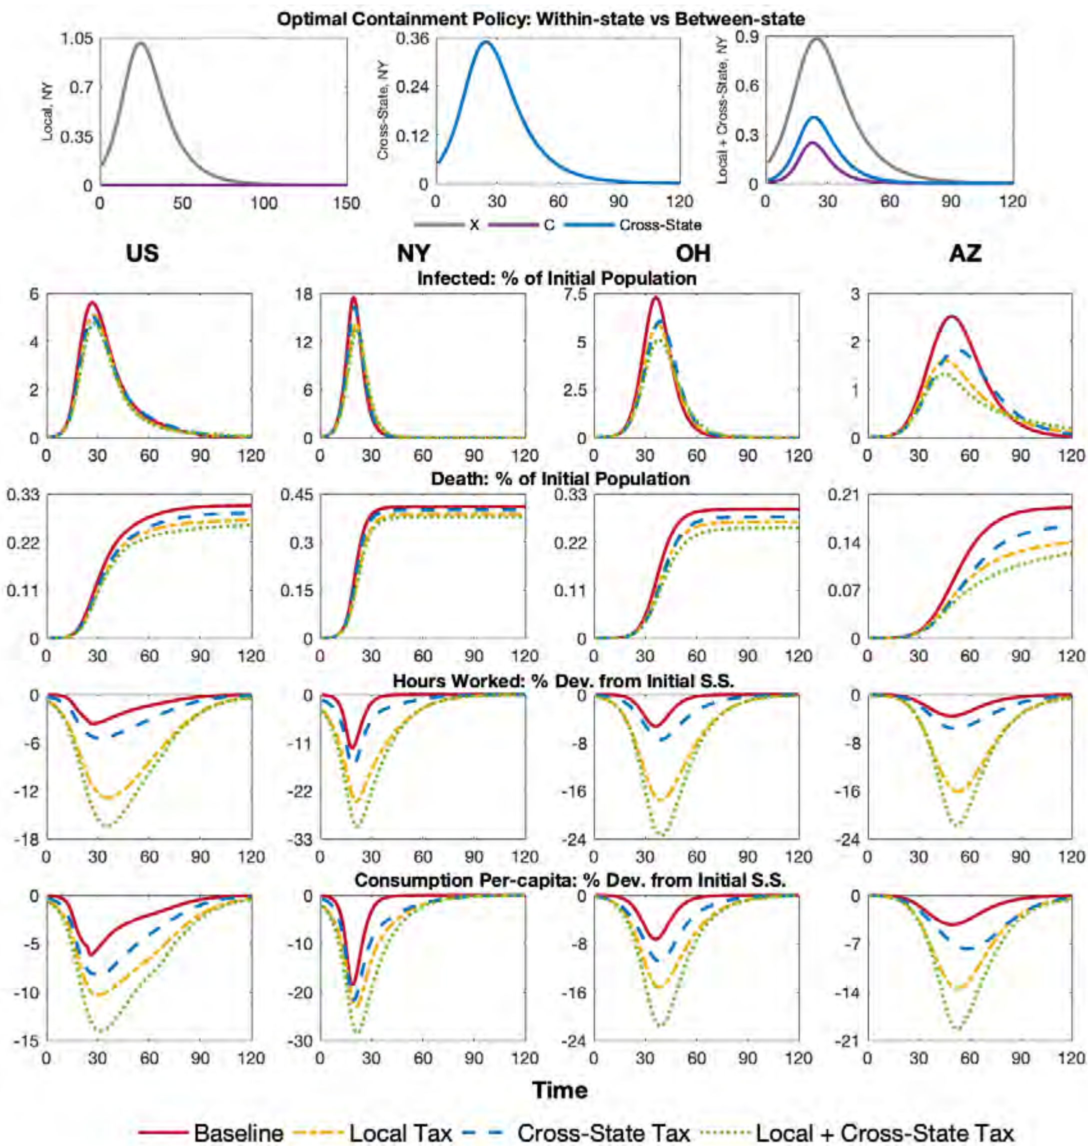

Covid Economics 52, 15 October 2020: 68-95

Covid Economics 52, 15 October 2020: 68-95

Optimal Within & Between State Containment Policy We now analyze the case where the social planner can jointly choose the optimal combination of local within- and local between-states consumption tax rates. The optimal tax paths for New York are reported in the third graph of the first row of figure 4. We find that the peak of the taxation on c internally happens earlier than in foreign goods and sector x, which are taxed at much higher extent. When analyzing the health dynamics of the policy we find that the optimal overall policy would reduce infection peak by 0.04p.p. compared to the baseline with gains happening across states. Similarly, the death toll would decrease by 132,200 lives compared to the baseline and 31,650 lives compared to the local within-state policy. As the previous analysis suggests, most of the gain occurs by implementing within-state policies. These saved lives is follow by a stronger economic drop. The optimal policy would lead to a peak drop in hours worked and consumption of 16.4 and 14p.p., respectively. Mobility and trade flows would drop approximately 33% at the peak and 5% on average.

These results suggest that despite some substitutability between within- and between-state policies, they mainly tackle different issues. While between-state policies can attenuate the pandemic by limiting the number of cases imported, it alone is not able to mitigate substantially the pandemic. Once cases are already within the state, only within-state policies can be effective.

Welfare Comparisons Table 1 reports the welfare losses attributable to the pandemic comparing different models and mitigation policies. To understand how our model differs in terms of welfare from others, we first compare our baseline economy to an isolated economy and one without behavioral response.[18] Column named "No Policy" reports the welfare impact of the pandemic for these three models. The results show that the baseline economy produces a welfare loss of 2.52p.p. larger than the isolated economy and 0.34p.p. larger than a model without behavioral response. Thus, we conclude that not including these margins when studying optimal policies biases substantially the aftermath of pandemics. We, then, compare the welfare effects of the battery of optimal policies discussed above. Regarding within-state policies, we find that the optimal national within-state containment policy would ameliorate welfare losses by 2.29p.p. while the optimal local level one would improve it in 3.09p.p.. These results highlight that a policy that could resemble a state-specific lockdown works better than a national lockdown. The key dimension through which this happens is time flexibility. We, then, report the welfare effect of a between-state tax on consumption that is either homogeneously (national) or heterogeneously (local) applied across states. The welfare improvement are very modest, approximately 0.2p.p., for both policies, showing that the

[18] In a model without behavioral agents do not respond by adjusting labor and consumption when they observe the infection rate going up.

Covid Economics 52, 15 October 2020: 68-95

best between-state optimal policy alone would not have the same welfare effects as a local within-state consumption tax. Moreover, a national within-state lockdown is better than a simple local between-state policy. Finally, when the planner is allowed choose the optimal combination of within and between-state policy instruments, welfare gains increase relative to the optimal within-state policy. We find that this policy applied at national and local level would mitigate the welfare losses by 2.34p.p. and 3.27p.p.. We conclude that the optimal policy is a combination of within local within-state and between-state policies.

Table 1: Welfare Impact of the Pandemic

Model	Welfare Loss (%)						
	No Policy	Optimal Policy					
		Within-state		Between-state		Overall	
		National	Local	National	Local	National	Local
Baseline	-31.65%	-29.36%	-28.56%	-31.48%	-31.39%	-29.31%	**-28.38%**
Isolated	-29.13%						
No Behavioral	-31.31%						

4 Conclusions and Future Work

We highlight how interconnectedness amplifies the severity of pandemics. We find that if the US was constituted of isolated states, there would be approximately 73,200 less deaths and the peak of consumption drop would be attenuated by 2.5p.p.. Therefore, we stress that the optimal containment policy must take into account interconnectedness and consider policies that temporarily limit the movement of people and goods across US states. We find that the optimal policy combines within and between-states restrictions and it saves approximately 132,200 lives. A promising application of this framework is the study of the optimal travelling restriction policies among countries. Finally, understanding whether the pandemic might have consequences on globalization by reducing trade and movement of people for long time spells is a long-term goal of this research agenda.

References

Alvarez, Fernando, David Argente, and Francesco Lippi, "A Simple Planning Problem for COVID-19 Lockdown," *COVID Economics 14*, 2020.

Antràs, Pol, Stephen J Redding, and Esteban Rossi-Hansberg, "Globalization and Pandemics," *COVID Economics 49*, 2020.

Argente, David O, Chang-Tai Hsieh, and Munseob Lee, "The Cost of Privacy: Welfare Effect of the Disclosure of COVID-19 Cases," 2020.

Atkeson, Andrew, "What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios," Technical Report, National Bureau of Economic Research 2020.

_ **et al.**, "How Deadly is COVID-19? Understanding the Difficulties with Estimation of its Fatality Rate," 2020.

Couture, Victor, Jonathan Dingel, Allison Green, and Jessie Handbury, "Device Exposure Index based on PlaceIQ Data," *Working Paper*, 2020.

Cuñat, Alejandro and Robert Zymek, "The (Structural) Gravity of Epidemics," *Covid Economics 17*, 2020.

Eichenbaum, Martin, Sergio Rebelo, and Mathias Trabandt, "The macroeconomics of epidemics," *Working Paper, National Bureau of Economic Research*, 2020.

Fajgelbaum, Pablo, Amit Khandelwal, Wookun Kim, Cristiano Mantovani, and Edouard Schaal, "Optimal lockdown in a commuting network," 2020.

Faria-e-Castro, Miguel, "Fiscal policy during a Pandemic," *COVID Economics 2*, 2020.

Ferguson, Neil M, Derek AT Cummings, Christophe Fraser, James C Cajka, Philip C Cooley, and Donald S Burke, "Strategies for mitigating an influenza pandemic," *Nature*, 2006, *442* (7101), 448–452.

Fernandez-Villaverde, Jesus and Charles I. Jones, "Estimating and Simulating a SIRD Model of COVID-19 for Many Countries, States, and Cities," *Working Paper, National Bureau of Economic Research*, 2020.

Gatto, Marino, Enrico Bertuzzo, Lorenzo Mari, Stefano Miccoli, Luca Carraro, Renato Casagrandi, and Andrea Rinaldo, "Spread dynamics of the Covid-19 Epidemic in Italy: Effects of Emergency Containment Measures," 2020.

Giannone, Elisa, Qi Li, Nuno Paixao, and Xinle Pang, "Unpacking Moving," *Working Paper*, 2020.

Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Vıctor Rıos-Rull, "Health versus Wealth: On the Distributional Effects of Controlling a Pandemic," *COVID Economics 6*, 2020.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning, "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?," 2020.

Jones, Callum, Thomas Philippon, and Venky Venkateswaran, "Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home," 2020.

Kaplan, Greg, Benjamin Moll, and Giovanni L Violante, "The Great Lockdown and the Big Stimulus: Tracing the Pandemic Possibility Frontier for the US," 2020.

Piguillem, Facundo, Liyan Shi et al., "The optimal covid-19 quarantine and testing policies," 2020.

Ramondo, Natalia, Andrés Rodríguez-Clare, and Milagro Saborío-Rodríguez, "Trade, domestic frictions, and scale effects," *American Economic Review*, 2016, *106* (10), 3159–84.

Covid Economics 52, 15 October 2020: 68-95

Covid Economics 52, 15 October 2020: 68-95

A Online Appendix

A.1 Figures and Tables

Figure 5: Correlations between Cumulative Deaths and State Characteristics

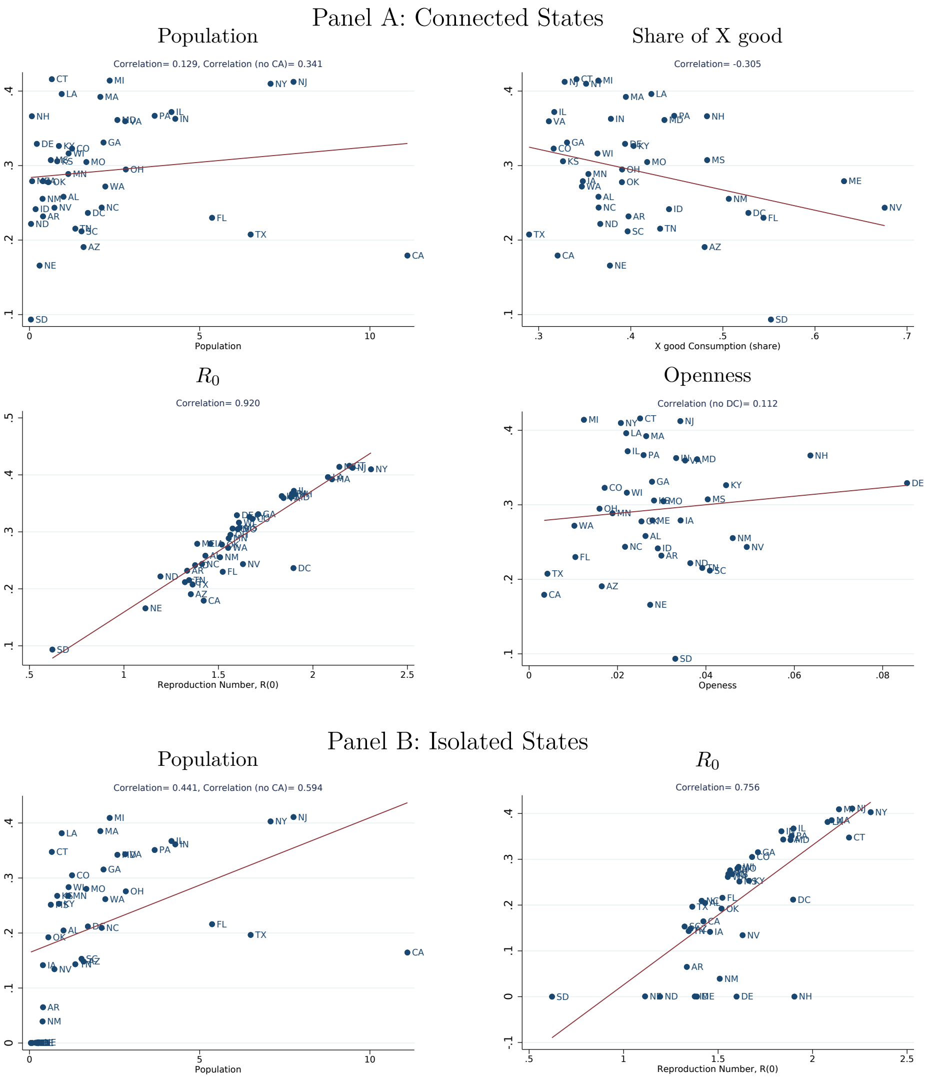

Covid Economics 52, 15 October 2020: 68-95

Figure 6: Correlations between Consumption Drop and State Characteristics

Panel A: Connected States

Population

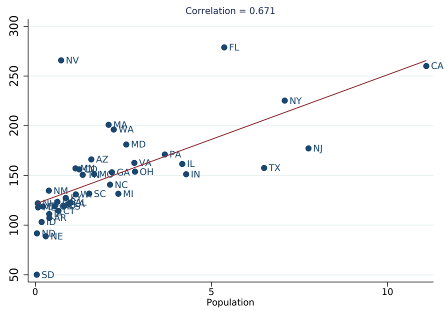

Share of X good

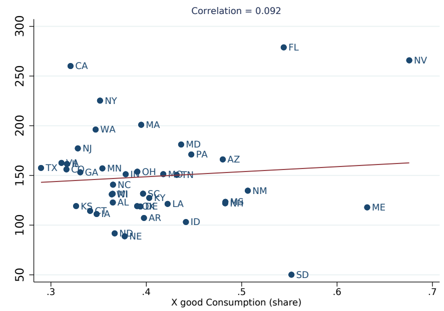

R_0

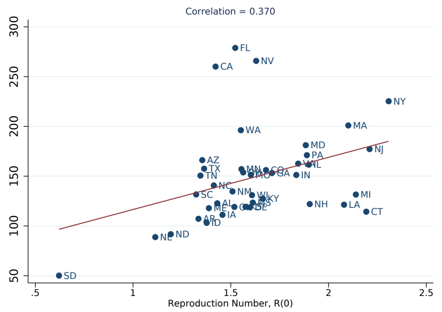

Openness

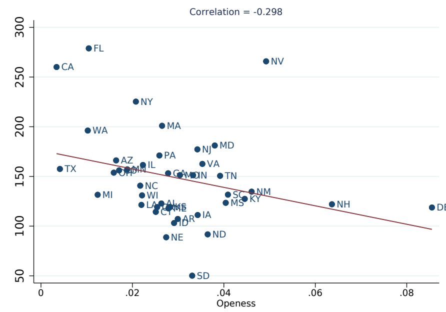

Panel B: Isolated States

Population

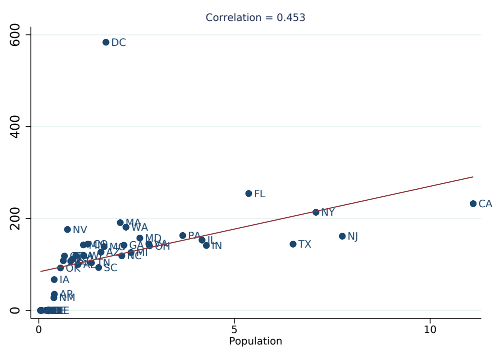

R_0

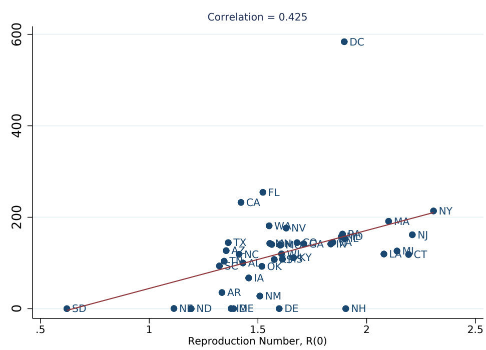

Covid Economics 52, 15 October 2020: 68-95

A.2 Optimization Problems

This section describes the optimization problem faced by the agents of this economy.

We start by discussing the consumption of regular goods from different regions. As widely known, the allocation of consumption across different varieties for a given level of expenditure is a static problem. An individual in location l, allocates the aggregate consumption of regular good, c_l, according to the following problem:

$$u(c_l) = \max_{\{c_{j,l}\}_{j=\{1,\ldots,L\}}} \left(\sum_{j=1}^{L} \alpha_{j,l} \tilde{c}_{j,l}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}$$

$$s.t. \quad \sum_{j=1}^{L} (1+\tau_{l,j}^c)\tilde{p}_j \tilde{c}_{j,l} = p_l^c c_l$$

There first order conditions are:

$$c_l^{\frac{1}{1-\epsilon}} \alpha_{j,l} \tilde{c}_{j,l}^{-\frac{1}{\epsilon}} = \lambda(1+\tau_{j,l}^c)\tilde{p}_j$$

After some algebra and defining the aggregate regular good price index after taxes in location l as,

$$(1+\tau_l^c)p_l^c = \left[\sum_{j=1}^{L} \alpha_{j,l}{}^{\epsilon} \Big((1+\tau_{j,l}^c)\tilde{p}_{j,t} \Big)^{1-\epsilon} \right]^{\frac{1}{1-\epsilon}},$$

we obtain that an agent in location l consumes from location j:

$$\tilde{c}_{j,l} = \left(\frac{(1+\tau_{j,l}^c)\tilde{p}_j}{\alpha_{j,l}(1+\tau_l^c)p_l^c} \right)^{-\epsilon} c_l$$

We are now left to solve for the aggregate consumption of regular and social good and hours worked for individuals of different health status, location and sectors across time.

Susceptible People A susceptible person s in location l in sector $k \in \{c, x\}$ chooses consumption $c_l^{k,s}$ and $x_l^{k,s}$ and number of hours worked $n_l^{k,s}$ that solves the following optimization problem:

$$U_{l,t}^{k,s} = \max_{\{c_{l,t}^{k,s}, x_{l,t}^{k,s}, n_{l,t}^{k,s}\}} u(c_{l,t}^{k,s}, x_{l,t}^{k,s}, n_{l,t}^{k,s}) + \beta \left[(1-\tau_{l,t}^k) U_{l,t+1}^{k,s} + \tau_{l,t}^k U_{l,t+1}^{k,i}, \right] \tag{6}$$

$$s.t. \quad (1+\tau_l^c)p_{l,t}^c c_{l,t}^{k,s} + (1+\tau_l^x)p_{l,t}^x x_{l,t}^{k,s} = w_{l,t}^k \nu^i n_{l,t}^{k,is} + T_{l,t}^{k,s}$$

where $\tau_{l,t}^k$, the probability of becoming infected is defined in equation (2). We assume that susceptible people take as given aggregate variables, but they understand how they

Covid Economics 52, 15 October 2020: 68-95

consumption and working decisions impacts their own probability of becoming effect. However, they don't internalize how their decisions impact the aggregate variables, giving origin to infection externality.
The first order conditions are:

$$u_1(c^{k,s}_{l,t}, x^{k,s}_{l,t}, n^{k,s}_{l,t}) = \lambda^{k,s}_{l,t}(1+\tau^c_l)p^c_{l,t} + \beta(U^{k,s}_{l,t+1} - U^{k,i}_{l,t+1})\pi_1 I_{l,t} C^i_{l,t}$$

$$u_2(c^{k,s}_{l,t}, x^{k,s}_{l,t}, n^{k,s}_{l,t}) = \lambda^{k,s}_{l,t}(1+\tau^c_l)p^x_{l,t} + \beta(U^{k,s}_{l,t+1} - U^{k,i}_{l,t+1})\pi_2 I_{l,t} X^i_{l,t}$$

$$\gamma\left(n^s_{l,t}\right)^\theta = \lambda^{k,s}_{l,t}\nu^s w^k_{l,t} - \beta(U^{k,s}_{l,t+1} - U^{k,i}_{l,t+1})\pi_3(I^k_{l,t}N^{k,i}_{l,t} + 1_{(k=x)}I_{l,t}X^i_{l,t})$$

where $\lambda^{k,s}_{l,t}$ is the Lagrangian multiplier associated with the budget constraint. As expected the shadow price of each good is not only the market price but also the impact of one extra unit of consumption/leisure on the probability of becoming infected. The change in probability weights $\beta(U^{k,s}_{l,t+1} - U^{k,i}_{l,t+1})$, the forgone future utility in case becomes infected. This forward looking component it's the crucial element that makes the problem of the susceptible dynamic even in the absence of any asset.

Infected People Infected people solves the following problem:

$$U^{k,i}_{l,t} = \max_{\{c^{k,i}_{l,t}, x^{k,i}_{l,t}, n^{k,i}_{l,t}\}} u(c^{k,i}_{l,t}, x^{k,i}_{l,t}, n^{k,i}_{l,t}) + \beta\left[(1-\pi_r-\pi_{d,l})U^{k,i}_{l,t+1} + \pi_r U^{k,r}_{l,t+1},\right] \tag{7}$$

$$s.t. \quad (1+\tau^c_l)p^c_{l,t}c^{k,i}_{l,t} + (1+\tau^x_l)p^x_{l,t}x^{k,i}_{l,t} = w^k_{l,t}\nu^r n^{k,i}_{l,t} + T^{k,i}_{l,t}$$

As in Eichenbaum et al. (2020) we implicitly assume that the cost of death is the foregone utility of life and that infected people don't take into consideration that they may infect other people. Therefore, the infected people's problem become static with the following first order conditions:

$$u_1(c^{k,i}_{l,t}, x^{k,i}_{l,t}, n^{k,i}_{l,t}) = \lambda^{k,i}_{l,t}(1+\tau^c_l)p^c_{l,t}$$

$$u_2(c^{k,i}_{l,t}, x^{k,i}_{l,t}, n^{k,i}_{l,t}) = \lambda^{k,i}_{l,t}(1+\tau^x_l)p^x_{l,t}$$

$$\gamma\left(n^i_{l,t}\right)^\theta = \lambda^{k,i}_{l,t}\nu^r w^k_{l,t}$$

Recovered People Similarly to infected people, the decisions of the recovered people are also static and satisfy the following problem:

$$U^{k,r}_{l,t} = \max_{\{c^{k,r}_{l,t}, x^{k,r}_{l,t}, n^{k,r}_{l,t}\}} u(c^{k,r}_{l,t}, x^{k,r}_{l,t}, n^{k,r}_{l,t}) + \beta U^{k,r}_{l,t+1} \tag{8}$$

$$s.t. \quad (1+\tau^c_l)p^c_{l,t}c^{k,r}_{l,t} + (1+\tau^x_l)p^x_{l,t}x^{k,r}_{l,t} = w^k_{l,t}\nu^i n^{k,r}_{l,t} + T^{k,r}_{l,t}$$

CEPR PRESS

where the first order conditions resemble the ones from the infected people.

A.3 Parameters Values

Space We calibrate the model to US states, therefore, there will be 51 locations. The decision to make a state-specific model is driven by the fact that several policies are implemented by state-level government rather than other units of geographies. We normalized the population in Alabama, the smallest state to 1.

Preferences Regarding the labor supply, we set γ to 0.001275 and the Frisch elasticity θ to 1 as in Eichenbaum et al. (2020). We set the discount factor β to be 0.9994, which reflects a yearly discount rate of 0.97 adjusted to account for the weekly model. This discount factor implies a value of a life of 9.3 million 2019 dollars in the pre-epidemic steady state, which is consistent with the economic value of life used by U.S. government agencies in their decisions process.

We consider that social consumption goods are the sum of healthcare expenditures, entertainment, food outside the house, education, apparel, personal services and personal care products and services following a similar definition as Kaplan et al. (2020), and the rest fall into the category of tradable consumption goods. From the Consumer Expenditure Survey we pin down ϕ to match the share of expenditure on tradable vs social consumption goods in 2018.

Regarding the economic linkages across states, we set the elasticity of substitution across states, ϵ to 5 as estimated by Ramondo et al. (2016). Following the trade literature, we parametrize $\alpha_{j,l}$ as a log-linear function of bilateral distance between states $\alpha_{j,l} = \alpha_0 dist^{\alpha_1}$ for $j \neq l$ and set $\alpha_{l,l} = 1$. This implies gravity equation on bilateral trade flows:

$$\log E_{j,l} = (\epsilon - 1)\alpha_1 \log(dist_{j,l}) + \delta_j + \delta_l + \eta_{j,l},$$

where $E_{j,l}$ is the expenditure of state l on state j's tradable goods and δ_j and δ_l are the origin and destination fixed effects. Using between-states shipments data from the 2017 Commodity Flow Survey, we estimate $(\epsilon - 1)\alpha_1$ to be -1.31. α_0 is then chosen to match the expenditure share of tradable goods in each state coming from the other states.

Production To estimate the productivity by sector in each state, z_l^c and z_l^x, we match the wage in the model with the pre-disease data on wage.

SIR In order to calibrate the parameters related to the SIR model we use a similar approach as in Eichenbaum et al. (2020). Therefore, we match the probability of recovery, π_r, and the probability of death, π_d, to the values reported in this paper. They claim that a 0.5%

Covid Economics 52, 15 October 2020: 68-95

death rate in the US after adjusting for the population structure. Taking into account that our model is weekly, we set π_d to be 0.00194, which is the equivalent of $7 \times 0.005/18$, where 18 is the average number of days that takes to recover or die. Hence, the probability of recovery if infected is set to $7 \times 0.995/18$.

Using the data from the Time Use Survey and the definition of "time-use in general community activities" of Eichenbaum et al. (2020), we find that 18% and 30% of the time spent on general community is used for the purchase of "goods and services" and "eating and drinking outside the home", respectively. Since according to Ferguson et al. (2006) 33% of virus transmission are likely occur in the general community, we set the average number of infections originated by consumption of regulator good C to 6% (0.33×0.18) and those originate by the consumption of social good X to 10% (0.33×0.3).

We also follow Eichenbaum et al. (2020) and assume that 17% of infections occur in the work place. The functional form assumed in 2 generates higher transmission rates while working in the social sector than in the regular good sector.

Finally, most of the transmissions occur at home or by randomly meeting people in activities not related to consumption or working. We departure from the literature that the likelihood of getting infected does not depend only on the number of infected people in the region but also depends on the likelihood of contacting with an infected person from other state. Travelling for leisure, regular commuting or the performance of professional duties as meeting with clients, conferences or simply the transportation of goods generates a large flow of people across regions. To disciple how likely we are to meet a person from the home-state versus a different state, we use data from Couture et al. (2020). This data set uses cell phone data to measure the movement across regions. In specific, they report among the smartphones that pinged in a given state in a certain day, the share of those devices that pinged in each of the other 50 state at least once during the previous 14 days.

Moreover, Fernandez-Villaverde and Jones (2020) shows a wide heterogeneity in the basic reproduction number, R_0, across states at the beginning of the pandemic. We calibrate the four parameters above to match such state-level estimates.

To sum up, π_1, π_2, π_3 and $\pi_{4,l}$, are chosen to match the following four equations:

$$\sum_{l}^{N} \frac{Pop_l}{Pop} \frac{\pi_1 X_l^2}{T_l} = 0.1$$

$$\sum_{l}^{N} \frac{Pop_l}{Pop} \frac{\pi_2 C_l^2}{T_l} = 0.06$$

$$\pi_3 \sum_{l}^{N} \frac{Pop_l}{Pop} \frac{\left(\frac{Pop_l^c}{Pop_l}\right)^2 (N_l^c)^2 + \left(\frac{Pop_l^x}{Pop_l}\right)^2 (N_l^x)^2 + \left(\frac{Pop_l^x}{Pop_l}\right)^2 N_l^x X_l}{T_l} = 0.17$$

Covid Economics 52, 15 October 2020: 68-95

CEPR PRESS

$$R_{0,l} = \frac{\frac{T_{0,l}}{I_{0,l}}}{\pi_d + \pi_r}$$

where

$$T_l = \pi_{1,l} X_l^2 + \pi_{2,l} C_l^2 + \pi_{3,l} \left(\left(\frac{Pop_l^c}{Pop_l} \right)^2 (N_l^c)^2 + \left(\frac{Pop_l^x}{Pop_l} \right)^2 (N_l^x)^2 + \left(\frac{Pop_l^x}{Pop_l} \right)^2 N_l^x X_l \right)$$

$$+\pi_{4,l} \left(\gamma_{l,l} + \sum_{j \neq l} (\gamma_{l,j} + \gamma_{j,l}) \frac{I_j}{I_l} \right)$$

Table 2: Parameter Values

Parameter	**Interpretation**	**Internal**	**Value**
Space			
N	Number of Locations	N	51
Preferences			
θ	Labor Disutility	N	0.001275
ϕ	consumption good c share	Y	0.735
β	Discount factor	Y	$0.97^{1/52}$
ρ	Elast. substitution between c and s	N	0.5
$\alpha_{i,j}$	Share of c from other states	Y	
ϵ	Elast. substitution between c from diff. states	N	5
Technology			
z^s	Productivity in s	Y	
z^c	Productivity in c	Y	
SIR			
π_r	Probability of recovery	N	$7 \times 0.995/18$
π_d	Probability of dying	N	$7 \times 0.005/18$
π_1	Infection by X	Y	1.33×10^{-4}
π_2	Infection by C	Y	2.86×10^{-5}
π_3	Infection by Working	Y	1.64×10^{-4}
$\pi_{4,l}$	Infection by General contact	Y	

Note: This table reports the parameters' values used in the calibration stating whether they are internal or externally calibrated. The model is calibrated at a weekly frequency.

Covid Economics 52, 15 October 2020: 68-95

Pandemics and income inequality: A historical review

Adham Sayed[1] and Bin Peng[2]

Date submitted: 8 October 2020; Date accepted: 9 October 2020

This paper examines the effects of pandemics on income inequality, specifically those pandemics that claimed more than 100,000 lives. Given that pandemics are events that rarely occur, we have use data spanning over the last 100 years (1915-2017) and relating to four pandemics. The study includes four countries that had income inequality data covering that period. Using panel data methods – Fixed Effects and Augmented Mean Group estimators – we found a significant effect of these pandemics on declining income inequality. The study argues that based on the characteristics of the COVID-19 pandemic, namely that fatalities are highly concentrated in older age groups, we can neither expect a labor scarcity nor a sharp decline in productivity; however, we could expect a reduction in consumption, the possibility of savings, high unemployment rates, and high public debt ratios. The ultimate effects of COVID-19 on inequality remain unclear so far, as some of its inherent characteristics push for an increase in inequality. In contrast, others push toward a narrowing of the income gap.

Covid Economics 52, 15 October 2020: 96-117

1 Huazhong University of Science and Technology.
2 Huazhong University of Science and Technology.

Introduction

The number of COVID-19 related infections and deaths continue to rise, accounting, by mid-September 2020, for more than 28 million confirmed cases and around 910,000 deaths. The coronavirus is now the cause of the world's most severe recession in decades (World Bank, 2020). Faced with this reality, there is significant uncertainty about the eventual coronavirus outbreak's effects on the world economy. Yet, researchers agree that deadly pandemics have detrimental effects that adversely affect economies and societies alike in both the short and medium terms.

In contrast, it is believed that major crises, such as wars and pandemics, have direct impacts on income distribution in any society. Piketty and Saez (2014) argue that the First and Second World Wards were behind the narrowing of income gaps in the last 150 years. Milanovic (2016) agrees with this conclusion and adds that pandemics also reduce income inequality and represent a force that pushes towards increased income equality. Similarly, Alfani and Murphy (2017) conclude that lower-income inequality is the expected outcome of crises with high mortality rates. However, opinions differ on the relationship between pandemics and income distribution. Studies examining this relationship find that the subsequent consequences of pandemics on inequality are primarily related to the characteristics that distinguish one virus (or any other cause of a pandemic) from another.

Due to the difficulty of predicting the eventual results of the COVID-19 pandemic, we follow Barro et al. (2020), and Jordà et al. (2020), in examining the worst-case scenario. We extend our study to include pandemics whose death toll exceeded 100,000 fatalities in the last century. Thus, as adopted in Jordà et al. (2020)'s work, we will focus on pandemics (caused by a virus or bacteria) with 100,000 deaths or more. We will start with the case of the Spanish flu pandemic that spread worldwide after World War I in four successive waves lasting from 1918 until 1920 and claimed 40 to 100 million lives.[1] We will then consider the 2009 H1N1 pandemic after looking at the Asian Flu and the Hong Kong Flu (see Table 1).

For reasons related to the long-term availability of data, this study focuses on a few European countries in addition to the United States of America. The first panel includes four developed countries - the United States, the United Kingdom, France, and Germany - between 1915 and 2017. The second contains 49 US states for the period 1917-2015. We will also conduct a

[1] There is no unified number of the deaths resulted from this pandemic. In their study, Barro et al. (2020) estimated it at 40 million, while Jordà et al. (2020) adopt the 100 million deaths.

historical review[2] of the effects of pandemics on inequality, namely the Black Death (also known as the Black Plague) and the Great influenza (Spanish flu), based on several classical theories and modern studies, as well as historical data that until recently was absent from researchers' works. We mainly rely on Alfani's research results, who is considered one of the few researchers who studied this relationship in the long term (see Alfani, 2015; Alfani and Percoco, 2016; Alfani and Murphy, 2017; Alfani and Ammannati, 2017).

Table 1: Pandemic events with at least 100,000 deaths, 1918-2018.

Event	Start	End	Deaths
Spanish Flu	1918	1920	100,000,000
Asian Flu	1957	1958	2,000,000
Hong Kong Flu	1968	1969	1,000,000
H1N1 Pandemic	2009	2009	203,000

Source: Jordà et al. (2020)

Based on Pesaran and Smith (1995) (Mean Group estimator, MG), and Eberhardt and Bond (2009) (Augmented Mean Group, AMG) and in order to capture unobserved effects (such as economic shocks and wars) that may affect income inequality; we adopt heterogeneous interactive effects panel data models controlling for unobservable common factors. Our results confirm that, pandemic events with more than 100,000 deaths over the past century reduced income inequality in post-pandemic years. These results will help bridge the large gap in the literature examining the impact of pandemics on income inequality.

The rest of the paper is organized as follows. Section 2 introduces the literature review; Section 3 presents our data; Section 4 describes the empirical model; Section 5 presents the empirical results; Section 6 discusses the expected impact of COVID-19 on income inequality. Section 7 concludes.

Literature review

The direct impact of the pandemic on income distribution stems from the fact that the spread of a deadly virus in a society leads to the death of a large number of workers, as poverty-stricken and low-income groups are most vulnerable to disease (Furceri et al., 2020; Galletta and Giommoni, 2020). For example, the relative impact of the SARS-CoV-2 virus in the United

[2] This review will focus on two historical events considered as the worst epidemiological scenarios that humankind has experienced in the past 700 years.

Covid Economics 52, 15 October 2020: 96-117

States was far greater on poor communities than on affluent ones (Schmitt-Grohé et al., 2020). So far, studies have concluded that the economic and health impact of COVID-19 on the most vulnerable individuals –the poor, the homeless, etc.– has been much greater than on other sections of the population (Bell et al., 2020[3]; Adams-Prassl et al., 2020; Alon et al., 2020). Moreover, the 1918 flu pandemic witnessed very high mortality rates among workers (Brainerd and Siegler, 2003). Consequently, a high mortality rate among workers leads to a scarcity of labor supply, which in turn results in higher wages (Alfani and Murphy, 2017). Similarly, a pandemic can lead to the emergence of a production crisis and a reduction in consumption as people tend to save more, thus directly affecting the return on capital that the rich benefit from (the top 10% income share and above). Such a pandemic effect leads, therefore, to the narrowing of the income gap.

In his study, Cipolla (1964, p 524) argued that the Black Death might have been the leading cause in the redistribution of income through the increase of wages. Based on long term data as well as evidence generated by research on real wages (Pamuk, 2007; Pamuk and Shatzmiller, 2014). Alfani and Murphy (2017) corroborated Cipolla's conclusion, arguing that labor scarcity led to an increase in real wages, thus narrowing the income gap. In their opinion, this was one of the direct effects of the Black Death on society. Moreover, these results refuted misconceptions that prevailed when Black Death brought about a rise in inequality (see Herlihy, 1978).

The decline in inequality after the pandemic argument is summarized in Figure 1. It assumes that productivity remains constant, while the shock to the labor supply caused by the deadly pandemic (the Plague, for example) causes a decline in curve S to S', which leads to an increase in the wages of those workers who survive the pandemic, from W to W'. This argument proposed by Capasso and Malanima (2007) was used by Li and Li (2017) in their study of the subsequent impact of the Plague that struck families immigrating to the Manchuria region in northern China in 1910-1911. The study examined a family that arrived in a village afflicted with the Plague and another that came to a neighboring village, free of the Plague. The study found that the wages of the first family were higher than those of the second family. To explain this difference, researchers used the classical economic theory of supply and demand (see Figure 1). For example, consider that two villages had the same wage level before the Plague. In the town affected by the Plague, the ensuing deaths and the delayed arrival of seasonal

[3] https://voxeu.org/article/prepare-large-wage-cuts-if-you-are-younger-and-work-small-firm

Covid Economics 52, 15 October 2020: 96-117

migrants lead to a decrease in the labor supply, thus shifting the supply curve from S to S'. It results that wages in this town increase from W to W'. Whereas for the town that is not affected by the Plague, wage levels remain unaffected.

Figure 1: The effect of the Plague on labor supply

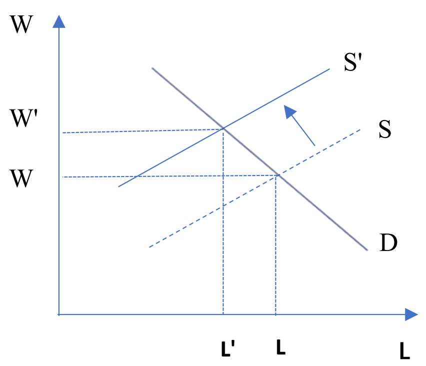

However, this assumption cannot be generalized to all pandemics. Productivity itself is not always stable since pandemics may lead to a decline in productivity. Alfani and Percoco (2016) argue that there is no reason to assume that every pandemic (defining the Plague) has a positive impact on wages because the ultimate effect depends on the magnitude of the labor supply shock relative to the productivity shock caused by the pandemic. This can be shown in Figure 2, where the shock in productivity shifts the curve for labor demand from D to D ', leading to a decrease, stability, or a rise in wages. Therefore, if the curve D falls below the curve S, we should expect the pandemic to result in higher wages. The opposite will result in lower wages.

Figure 2: The combined effect of the Plague on labor supply and productivity

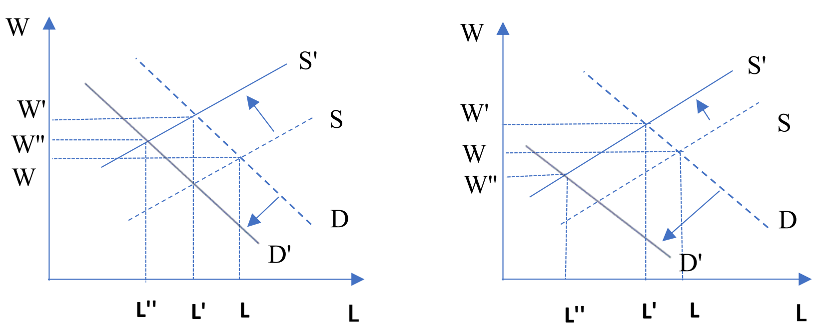

Accordingly, the subsequent effects of the pandemic on income inequality may differ according to the characteristics of the pandemic itself. Consequently, it cannot be inferred that every

Covid Economics 52, 15 October 2020: 96-117

pandemic leads to a decrease in inequality. In this context, Alfani (2015) and Alfani and Ammannati (2017) assert that although in the very long term, inequality has always shown a tendency to rise, the Black Death was the only recorded event that led to a decrease in income inequality in northwest Italy in the period of 1300-1800, as compared to other epidemiological events that did not reduce inequality. In fact, since the seventeenth century, events with a high death toll have had entirely different effects from those of the Black Death. Consequently, Alfani (2015) concluded that the ensuing effects of different plague pandemics on inequality levels were not the same.

The effects of the Black Death in Italy are similar to those of Spain. Álvarez-Nogal and De La Escosura (2013) provided historical data dating back to the year 1200 (see Figure 3) and showing the levels of inequality in Spain in the very long term. Based on this data, Milanovic (2016) concluded that income equality declined due to the Black Death, and wages in northwestern Europe increased significantly because of this pandemic.

Figure 3: Land Rent-Wage Rate Ratio in Spain (1282-1850)

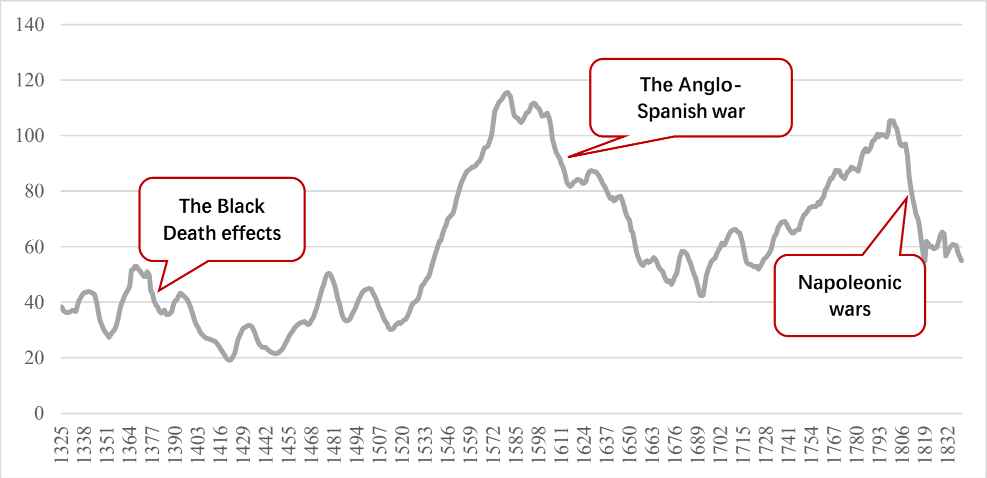

Note: Data is from Álvarez-Nogal and De La Escosura (2013). Comments on the path of inequality are from Milanovic (2016).

This result is confirmed by Clark (2007; 2010) in England, where the Black Death led to a decrease in the labor supply by 25% to 40%, while real wages rose by almost 100%. On the other hand, this pandemic decreased the rates of return on land from 5% to around 8%. As for Jordà et al. (2020), they examined the macroeconomic cost of 15 pandemics, particularly the return rates on assets, based on data dating back to the fourteenth century. Their study found that the after-effects of a pandemic can last for up to four decades, during which the natural rate of interest drops drastically (1.5%), and in the same period, real wages rise by up to 5%. The study also concludes that the response of real wages mirrors that of the natural interest rate (which also measures the return on capital). Their findings coincide with the neoclassical model's predictions and are also consistent with historical narratives: One of the Black Death effects was labor scarcity in the European economies, which is believed to have led to the rise in real wages. In parallel, this has been accompanied by lower returns on capital.

The income effects of the Black Death are not much different from the grave pandemic that spread worldwide between 1918 and 1920 - the Spanish flu. The results of Jordà et al. (2020) that include the 1918-1920 pandemic, confirm that each pandemic with a death toll of more than 100 thousand people can have a positive impact on reducing disparities in income. However, the Spanish Flu virus, which wiped out 2% of the world's population at the time, still has its unique features and a far greater impact on inequality than other pandemics. The age curve of mortality took a "W" shape for the Spanish flu, meaning that the number of deaths among individuals aged between 15 and 44 years was high, in contrast to the "U-shaped" mortality curve of the influenza virus (Brainerd and Siegler, 2003).

In addition to these factors, Barro et al. (2020) found that the high mortality rates of the 1918 flu pandemic led to a decrease in real per capita consumption by 8%, while real returns on stocks and short-term government bonds decreased drastically. Therefore, based on the modern economic theory on labor scarcity and under-consumption mentioned previously (see also, Ramsey, 1928; Rachel and Smith, 2017; Jordà et al., 2020), it is likely that the 1918 flu pandemic contributed to a reduction in income inequality.

However, the literature on the impact of pandemics on inequality is still scarce, and this may be due to the relative scarcity of pandemics in the past 100 years. As for empirical research examining the relationship between income inequality and pandemics, Alfani (2015) and Alfani and Ammannati (2017) studied the impact of pandemics in some Italian cities during the pre-industrial era and concluded that the Black Death had a fundamental role in the decline of inequality in Italy. Additionally, we refer to Furceri et al. (2020), who studied the impact of 5 epidemics that appeared in the past twenty years (2000-2020), namely SARS 2003, H1N1 2009, MERS 2012, Ebola 2014, and Zika 2016. Although these events had a relatively weak impact on the economy and society, only a few countries were affected, and the mortality rate

Covid Economics 52, 15 October 2020: 96-117

was moderately low compared to previous pandemics[4]—the 2009 pandemic was the only exception. Researchers found that these events led to a rise in inequality in the countries covered by the research. However, we believe that the characteristics of the viruses that have increased in the last two decades make it difficult to draw general conclusions about pandemics' impact on inequality. However, Galletta and Giommoni (2020) who studied the impact of the 1918 flu pandemic on inequality in Italian towns, concluded that in both short and medium terms, income inequality would increase in the towns that were most affected by the pandemic, due to the decrease in the income share held by the poor.

Data and Variables

To study the income inequality responses to the pandemic events during the last 100 years, we use available historical data on income inequality, measured by the top 10% income share collected from the world inequality database (WID) available at www.wid.world. Our sample is divided into two panels; the first one comprises four developed countries, mainly France, USA, UK, and Germany, covering the period from 1915 to 2017. Those countries are the only countries that have a nearly complete-time series since 1915. For all other countries, the time series is significantly shorter and mostly start in 1980. This limitation is the same for the Gini index.

For the pandemic events, following Jordà et al. (2020), we focused on four events (see Table 1). We constructed a dummy variable for a pandemic, which takes the value 1 when the number of deaths caused by the pandemic is more than 100,000, and 0 otherwise.

For the robustness check, we used two control variables. The first one is the level of economic growth measured by GDP per capita (GDPPC) in 2011 US$, 2011 benchmark, derived from two sources. For the period between 1915 and 1959, the data were collected from the Maddison Project Database, 2018 version (available at www.ggdc.net/maddison), while the rest was collected from the World Bank Database. The second control variable was the population (POP) derived from the Maddison Project Database, 2018 version, and the World Bank Database.

The data in figure 1 was sourced from Álvarez-Nogal and De La Escosura (2013). Table A1 in the Appendix illustrates descriptive statistics for the variables.

[4] The global deaths caused by these pandemics and the number of countries affected by the virus are, respectively: SARS (774 Deaths; 27 countries), H1N1 (200,000 Deaths; 148 countries), MERS (858 Deaths; 22 countries), Ebola (11,325 Deaths; 5 countries), and Zika (18 Deaths; 18 countries).

Covid Economics 52, 15 October 2020: 96-117

Empirical Model

To estimate the impact of pandemics on income inequality, we followed the methods proposed by Pesaran (2006) and Eberhardt and Bond (2009). Models 1 and 2 give our econometrics representation of the income inequality responses to the pandemic events. The first model is a Fixed effects model, with intercept and slope coefficients that vary across countries:

$$y_{it} = \alpha_i + \delta_t + \beta_i PAN_{it} + \aleph_i x_{it} + u_{it} \,, \quad i = 1,2,\ldots, N \,; t = 1,2,\ldots, T \qquad (1)$$

where y_{it} are a vector of dependent variables includes the top 10% share pre-tax national income (INQ_{it}) and its first difference (ΔINQ_{it}), for country i at time t; PAN_{it} is a dummy variable indicating a pandemic event that affects country i in year t; x_{it} is a vector that includes two lags of the pandemic dummy and the dependent variable. In the baseline, as for Barro et al. (2020) and Furceri et al. (2020), we do not control for other factors affecting inequality; Then, the regressions do not include any independent variables. That is, our focus is on disaster pandemics, which are treated as exogenous shocks. α_i is a country-specific effect; δ_t is a country-specific time trend; and u_{it} is the error term.

Recent literature argues that the history of income distribution has been political in-depth, and has had a strong correlation with the political and financial policies, globalization, war and pandemics, and other unknown common shocks; see Acemoglu and Robinson (2002); Piketty and Saez (2003, 2014); Milanovic (2016); Alfani (2015); and Alvaredo et al. (2018). It further leads to the presence of cross-sectional dependence. Ignorance of common global shocks may lead to inconsistent and biased estimates (Ditzen, 2018). Therefore, we adopted heterogeneous interactive effects panel data models controlling for unobservable common factors.

To account for all the unobserved cofounders, we used an interactive effects panel data model introduced by Pesaran (2006). Then we considered a model with an unobserved common factor and a heterogeneous factor loading, represented as:

$$INQ_{it} = \alpha_i + \delta_t + \beta_i PAN_{it} + \aleph_i x_{it} + \gamma_i f_t + e_{it} \,, \quad i = 1,2,\ldots,N; t = 1,2,\ldots,T \qquad (2)$$

where, f_t is a $(m \times 1)$ vector of unobserved common effects, where m is the number of unobserved factors; γ_i is a heterogeneous factor loading, and e_{it} is a cross-section country-specific IID error term. f_t, γ_i and e_{it} are all unobserved. In this case, common factors are possibly non-stationary and are maybe correlated with the regressors. e_{it} is allowed to be serially correlated over t and weakly correlated across i (Chudik et al., 2011).

Covid Economics 52, 15 October 2020: 96-117

Eberhardt and Bond (2009) proposed an Augmented Mean Group (AMG) estimate for macro panel data that allows for slope heterogeneity, non-stationarity, and accounts for cross-sectional dependence through the inclusion of a "common dynamic effect" in the country regression, which is extracted from the year dummy coefficients (D_t) of a pooled regression in first differences (FD-OLS) and represents the levels-equivalent mean evolvement of unobserved common factors across all countries. The first stage is a standard FD-OLS regression with ($T - 1$) year dummies in the first differences (Δ), from which the year dummy coefficients ($\hat{\mu}_t^{\circ}$) were collected. The AMG estimations are derived as averages of the individual country estimations.

Empirical Results

Before we modeled the cross-sectional dependence, we checked its existence. We considered LM_{BP}, LM_{adj}, and LM_{CD} tests proposed by Breusch and Pagan (1980), Pesaran et al. (2008), and Pesaran (2004), respectively. Table 2 shows the results of the cross-sectional dependence tests. We rejected the null hypothesis of no cross-sectional dependence for both panels.

Table 2: Cross-sectional Dependence Tests Results

Tests	***Panel 1 (4 countries)***		***Panel 2 (49 States)***	
	Statistics	***P-Values***	***Statistics***	***P-Values***
LM_{BP} test	359.4	0.000	6.20E+04	0.000
LM_{adj} test	580.3	0.000	6981	0.000
LM_{CD} test	18.86	0.000	241.1	0.000

Note: if T (time dimension) is larger than N (cross-sectional dimension), we may use for these purposes the Breusch and Pagan's (1980) Lagrange multiplier (LM) test (LM_{BP}). But if N is larger than T, we may use for these purposes the Pesaran's (2004) CD test (LM_{CD}).

Due to the existence of cross-sectional dependence, the Im-Pesaran-Shim panel unit-root test (CIPS) proposed by Pesaran (2007) was used to check the stationarity of the series. Table 3 reports the results and shows that both panels fail to reject the null hypothesis; therefore, all variables are non-stationary.

Covid Economics 52, 15 October 2020: 96-117

Table 3: Im-Pesaran-Shin (CIPS) Panel Unit Root Test Results

Variables	***Panel 1***		***Panel 2***	
	Value	***P-Values***	***Value***	***P-Values***
INQ	-0.3688	0.3561	0.2324	0.4081
ΔINQ	-3.8451	0.0001	-27.4413	0.000

Note: the null hypothesis is the presence of unit root in panel data with cross-sectional dependence in the form of common factor dependence.

Since the variables are non-stationary, we needed to test for co-integration. Table 4 reports the results of the Kao (1999) panel co-integration tests. These tests can accommodate serially correlated error terms, country-specific intercepts, and trends. For all tests, the results indicated the rejection of the null hypothesis of no co-integration.

Table 4: Kao Panel Co-Integration Tests Results

	Panel 1		**Panel 2**	
	Statistic	**p-value**	**Statistic**	**p-value**
Modified Dickey-Fuller t	-26.3196	0.000	-7.1831	0.000
Dickey-Fuller t	-16.9713	0.000	-6.1137	0.000
Augmented Dickey-Fuller t	-10.8971	0.000	-2.9823	0.0014
Unadjusted modified Dickey Fuller t	-62.9571	0.000	-25.1922	0.000
Unadjusted Dickey-Fuller t	-20.4795	0.000	-12.1549	0.000

After we verified the presence of co-integration in our model, Tables 5 and 7 show results from the estimation of the panels 1 and 2 (respectively) with fixed effects (FE) and augmented mean group (AMG) estimators.

Table 5 reports the results for panel 1 with INQ as the dependent variable is in the first and second columns, and those with ΔINQ as the dependent variable are in the next two columns. We found that estimated coefficients of the dummy variables representing the pandemics events (PAN) are all significantly negative, at least at the 5% level, indicating that pandemics with more than 100,000 deaths lead to a significant decrease in income inequality. Then, the pandemics in the last 100 years are estimated to have reduced income inequality by about 8% and 5% when we controlled for the unobserved common shocks (using AMG estimator for both INQ and ΔINQ respectively), and about 0.5% using FE estimator for both INQ and ΔINQ.

For lag variables, using FE, the estimated coefficient on the first and second pandemics lag variables are positive but insignificant. In contrast, using AMG, both first and second lag

Covid Economics 52, 15 October 2020: 96-117

variables are significantly negative, indicating that the redistributive effects of pandemics on income inequality tend to last for some time.

This finding is consistent with results regarding the post-pandemic effect on inequality distribution, real wages, and return on capital, such as those reported in Clark (2007; 2010), Jordà et al. (2020), Alfani and Murphy (2017), and Li and Li (2017). For a panel of countries, these studies found that pandemic with unique characteristics (e.g., Hight death rate) lead to a decline in income inequality. Besides, our findings are contrary to the results obtained from other studies, which tend to show that pandemics increase the income gap (e.g., Furceri et al. (2020), and Galletta and Giommoni (2020)**)**.

Table 5: Panel Data Estimates for Determinants of Income Inequality. Panel 1

Independent variables	***INQ***		***ΔINQ***	
	FE	***AMG***	***FE***	***AMG***
PAN	-0.0048** (-2.08)	-0.08138* (-20.73)	-0.00538** (-2.27)	-0.05203* (-14.06)
$\boldsymbol{PAN_{i,t-1}}$	0.00304 (1.16)	-0.03012* (-14.19)	0.002987 (1.12)	-0.026976* (-12.96)
$\boldsymbol{PAN_{i,t-2}}$	0.0016 (0.7)	-0.0137* (-8.88)	0.00115 (0.48)	-0.04024* (-30.77)
$\boldsymbol{INQ_{i,t-1}}$	1.1539* (23.7)	0.536096* (8.9)	______	______
$\boldsymbol{INQ_{i,t-2}}$	-0.1961* (-4.03)	-0.28109** (-2.12)	______	______
$\boldsymbol{\Delta INQ_{i,t-1}}$	______	______	0.186997* (3.74)	-0.16401* (-3.49)
$\boldsymbol{\Delta INQ_{i,t-2}}$	______	______	-0.05557 (-1.12)	-0.21049* (-3.42)
Constant	0.0152* (3.41)	0.3295* (8.66)	5.02E-05 (0.07)	0.052921* (77.78)
n. obs	410	410	409	409

*Notes: *Significant at 0.01 level; **at 0.05 level; ***at 0.10 level. t-statistics in parentheses.*

We carried out a robustness check of these findings by including two control variables in the regression – real GDP per capita and population – and we repeated the regression using the same estimation methods. The results are reported in Table 6, and the estimated parameters for pandemic events (PAN) are close to those in Table 5, showing that pandemics lead to a reduction in income inequality. The only difference is that using AMG; the results show that the impact of pandemics is negatively highest when we control for economic growth and population (pandemics reduce inequality by around 12%).

Table 6: Panel Data Estimates. With control variables

Independent variables	***INQ***		***ΔINQ***	
	FE	***AMG***	***FE***	***AMG***
PAN	-0.0048** (-2.06)	-0.134* (-33.91)	-0.00533** (-2.25)	-0.1231* (-36.69)
$PAN_{i,t-1}$	0.00303 (1.16)	-0.074* (-34.88)	0.00296 (1.11)	-0.1276* (-61.04)
$PAN_{i,t-2}$	0.0016 (0.7)	-0.072* (-29.27)	0.00114 (0.48)	-0.0839* (-69.01)
$INQ_{i,t-1}$	1.1496* (23.54)	0.4606* (7.05)	———	———
$INQ_{i,t-2}$	0.1917* (-3.93)	-0.3329* (-2.76)	———	———
$\Delta INQ_{i,t-1}$	———	———	0.18261* (3.65)	-0.1542* (-2.82)
$\Delta INQ_{i,t-2}$	———	———	-0.0595 (-1.2)	-0.1818* (-2.62)
GDPPC	3.67E-08 (0.53)	-1.88E-06 (-0.7)	5.85E-08 (0.83)	7.84E-07* (3.73)
POP	1.11E-08 (0.45)	3.48E-06 (1.43)	4.54E-09 (0.18)	1.35E-06* (3.24)
Constant	0.0133* (2.84)	0.208* (2.44)	-0.0016 (-0.89)	-0.01843 (-1.25)
n. obs	410	410	409	409

*Notes: *Significant at 0.01 level; **at 0.05 level; ***at 0.10 level. t-statistics in parentheses.*

Covid Economics 52, 15 October 2020: 96-117

Table 7 repeats the analysis for 49 states of the United States of America (Panel 2). The sample size is bigger than that of panel 1. The main results are similar to those of panel 1 (see Table 5), although the estimated effects on income inequality are more extensive in magnitude.

Table 7: Panel Data Estimates for Determinants of Income Inequality. Panel 2

Independent variables	*INQ*		*ΔINQ*	
	FE	*AMG*	*FE*	*AMG*
PAN	-0.00659*	-0.10838*	-0.00834*	-0.0845*
	(-3.63)	(-63.78)	(-4.45)	(-55.06)
$PAN_{i,t-1}$	0.00625*	0.072629*	0.006316*	0.0853*
	(3.11)	(78.02)	(3.04)	(86.46)
$PAN_{i,t-2}$	0.00343**	-0.15851*	0.004329*	0.00558*
	(1.91)	(-114.2)	(2.34)	(6.31)
$INQ_{i,t-1}$	0.772279*	0.06254	_______	_______
	(53.28)	(1.5)		
$INQ_{i,t-2}$	0.07459*	-0.17061*	_______	_______
	(5.15)	(-7.7)		
$\Delta INQ_{i,t-1}$	_______	_______	-0.16787*	-0.4896*
			(-11.58)	(-18.85)
$\Delta INQ_{i,t-2}$	_______	_______	-0.11128*	-0.2287*
			(-7.67)	(-12.04)
Constant	0.0572*	0.4363*	-0.0002	-0.0874*
	(18.31)	(29.38)	(-0.37)	(-329.5)
n. obs	4849	4849	4848	4848

*Notes: *Significant at 0.01 level; **at 0.05 level; ***at 0.10 level. t-statistics in parentheses.*

What is the expected impact of COVID-19 on income inequality?

The previous parts offered the presentation of a worst-case scenario caused by a pandemic. By comparison, the COVID-19 pandemic presents certain characteristics that distinguish it from previous pandemics. In what follows, we will discuss some of the most important factors that may control the eventual impact of the current outbreak on income inequality.

Mortality rates

The SARS-CoV-2 virus is considered one of the most dangerous viruses to infect the respiratory system since the beginning of the last century, specifically since the 1918 flu pandemic (Ferguson et al., 2020). McKibben and Fernando (2020) used the same model[5] as McKibbin and Sidorenko (2006) to explore seven different scenarios of how COVID-19 might evolve in the coming year. They found that for the lowest and the highest level of the pandemic scenarios, the death toll estimates range between 15 and 68 million deaths, respectively. However, after more than nine months since the start of the COVID-19 pandemic and given the medical advancement and health measures taken to control the spread, the death toll of COVID-19 is not expected anymore to reach the levels of the 1918 flu pandemic. In fact, according to Barro et al. (2020), the 1918 flu pandemic led to the death of 2% of the world's population at that time - in today's numbers, that is equivalent to 150 million people. Regardless of the projected death toll, the highest mortality rates are so far concentrated in the seniors age group - that is 65 years of age and above. Thus, a direct impact on the labor supply is not expected.

Pandemics and the Economy: from the "Great Depression" to the "Great Lockdown"

The rapid and severe shock of the COVID-19 pandemic and the lockdown measures taken to contain it plunged the global economy in the midst of the worst recession since World War II. According to World Bank projections, the US economy will contract by 6.1%, while the Euro area is expected to contract by 9.1% in 2020. The average per capita income is also expected to decrease by 3.6%, causing millions of people to fall into extreme poverty. Other scenarios predict further deteriorating conditions where global GDP may fall by up to 8% this year (World Bank, 2020a). The International Monetary Fund has called this unprecedented crisis "The Great Lockdown" and classified it as the worst recession since the years of the "Great Depression" and much worse than the global financial crisis of 2008-2009; a first for a pandemic. The cumulative loss in global GDP due to this crisis is expected to be about 12 trillion US dollars over 2020 and 2021 (IMF, 2020a).

In this context, McKibben and Fernando (2020) showed the losses that can be incurred by the GDP in the global economy. However, their estimates showed that the COVID-19 pandemic

[5] Uses a modelling technique to explore four different pandemic influenza scenarios: "mild" scenario in which the pandemic is similar to the 1968-69 Hong Kong Flu; "moderate" scenario which is similar to the Asian flu of 1957; "severe" scenario based on the Spanish flu of 1918-19 (lower estimate of the case fatality rate), and an "ultra" scenario similar to Spanish flu but with upper-middle estimates of the case fatality rate.

Covid Economics 52, 15 October 2020: 96-117

effects are greater than what their study predicted. According to the adopted scenarios, the study expected losses of around 2.4 trillion US dollars based on a low-end pandemic pattern. As for the worst-case scenarios similar to the Spanish flu, the global GDP is expected to drop by about 9 trillion US dollars during the first year of the pandemic. This number remains lower than the 12 trillion US dollars predicted by the IMF. Therefore, this pandemic is expected to have one of the worst economic repercussions ever. Despite the pessimistic outlooks at the beginning of the virus' spread, the results were far more disastrous.

The present pandemic has led to a significant decline in consumption, as isolation measures pushed people to reduce their consumption to the lowest level and opt for more precautionary savings (World Bank, 2020b), which will have a negative impact on returns on capital (Barro et al., 2020; Jordà et al., 2020). Furthermore, the total closure of countries contributed to a complete or partial halt in productivity, leading to a further decline in the returns on capital, all of which will reduce inequality. However, the impact on productivity may also lead to an increase in unemployment figures. According to ILO (2020a) estimates, about 40% of the global workforce are employed in sectors facing a high risk of worker displacement[6], thus forecasting severe adverse effects on income.

Public debt levels

The COVID-19 pandemic has led to a number of unprecedented financial interventions by governments aiming to limit the adverse consequences of the pandemic; this has contributed to a significant increase in public debt-to-GDP ratios (IMF, 2020b). According to several studies, public debt leads to an increase in income inequality, as countries work to borrow from those who own wealth to pay off the debt and its interests. Countries intend to impose new taxes, mostly targeting people with low income, thus leading to an inverted redistribution of income, from the poorest to the richest (Piketty, 2014; You and Dutt, 1996**)**. However, the ultimate consequences of debt depend on the measures taken by governments to repay this debt. In previous crises such as World War II, many countries imposed exceptional taxes on the rich to cover the cost of the war, which was a contributing factor in reducing inequality (Piketty and Zucman, 2014; Milanovic, 2016). In other events, such as the 2008 financial crisis, governments did not take the same approach, which led to an increase in inequality.

[6] An estimated 155 million full-time jobs were lost during the first quarter of 2020, and it estimated to reach 400 million for the second quarter (ILO, 2020b).

Based on the characteristics of the COVID-19 pandemic, namely that fatalities are highly concentrated in older age groups, hence, we can neither expect a labor scarcity nor a sharp decline in productivity. But we could expect a reduction in consumption, the possibility of savings, high unemployment rates, and high public debt ratios. The ultimate effects of COVID-19 on inequality remain unclear so far, as some of its inherent characteristics push for an increase in inequality while others push toward a narrowing of the income gap.

Conclusion

The COVID-19 pandemic has caused a massive economic crisis described as the worst in decades, and whose effects are even worse than those of the 2008 financial crisis. This pandemic has revealed the high level of economic, health, and social inequality that afflict almost all societies. Although the final effects of this pandemic cannot be foreseen yet, there is no doubt that it will have significant impacts on income inequality.

Accordingly, this paper aimed to learn from the past experiences by examining the impact of past pandemics on income inequality, relying on the definition provided by Jordà et al. (2020) of a pandemic as a health crisis (virus or bacteria) that leads to more than a hundred thousand deaths. The study period covered four pandemics that occurred between 1915 and 2017.

This study provided new evidence that the pandemics that occurred over the last 100 years contributed to a decline in income inequality in the years following the pandemics. These results also contributed to fill the void left by the lack of research in this field. The study argued that the final effects of the pandemic on income inequality are mainly related to the characteristics of the pandemic, especially in terms of its impact on labor supply, productivity and consumption, and public debt.

Based on these results, the final effects of the COVID-19 pandemic on income inequality remain unclear as this pandemic has characteristics that differentiate it from previous scenarios. Future research may help us extract the ultimate and final effects of the SARS-COV-2 virus on income inequality. Lastly, it is important to emphasize that governments can play a more significant role in containing the catastrophic effects of this pandemic by supporting the most affected vulnerable groups. That would contribute to mitigating the negative effects of the pandemic on income distribution.

Covid Economics 52, 15 October 2020: 96-117

References

Acemoglu, D., Robinson, J.A., 2002. The Political Economy of the Kuznets Curve. Rev. Dev. Econ. 6, 183–203. https://doi.org/10.1111/1467-9361.00149

Adams-Prassl, A., Boneva, T., Golin, M., Rauh, C., 2020. Inequality in the impact of the coronavirus shock: Evidence from real time surveys. J. Public Econ. 189, 104245. https://doi.org/10.1016/j.jpubeco.2020.104245

Alfani, G., 2015. Economic Inequality in Northwestern Italy: A Long-Term View (Fourteenth to Eighteenth Centuries). J. Econ. Hist. 75, 1058–1096. https://doi.org/10.1017/S0022050715001539

Alfani, G., Ammannati, F., 2017. Long-term trends in economic inequality: The case of the Florentine state, c. 1300-1800. Econ. Hist. Rev. 4, 1072–1102. https://doi.org/10.1111/ehr.12471

Alfani, G., Murphy, T.E., 2017. Plague and Lethal Epidemics in the Pre-Industrial World. J. Econ. Hist. 77, 314–343. https://doi.org/10.1017/S0022050717000092

Alfani, G., Percoco, M., 2016. Plague and long-term development: the lasting effects of the 1629–30 epidemic on the Italian cities, EHES Working Paper No. 106. https://doi.org/10.1111/ehr.12652

Alon, T., Doepke, M., Olmstead-Rumsey, J., Tertilt, M., 2020. The Impact of Covid-19 on Gender Equality, Covid Economics, 4, 62-85.

Alvaredo, F., Chancel, L., Piketty, T., Saez, E., Zucman, G., 2018. Global Inequality and Policy: The Elephant Curve of Global Inequality and Growth. AEA Pap. Proc. 108, 103–08. https://doi.org/10.1257/pandp.20181073

Álvarez-Nogal, C., De La Escosura, L.P., 2013. The rise and fall of Spain (1270-1850). Econ. Hist. Rev. 66, 1–37. https://doi.org/10.1111/j.1468-0289.2012.00656.x

Barro, R.J., Ursúa, J.F., Weng, J., 2020. The Coronavirus and the Great Influenza Pandemic: Lessons from the “Spanish Flu” for the Coronavirus’s Potential Effects on Mortality and Economic Activity, NBER Working Paper Series No. 26866.

Brainerd, E., Siegler, M., 2003. The Economic Effects of the 1918 Influenza Pandemic, CEPR Discussion Paper No. 3791.

Covid Economics 52, 15 October 2020: 96-117

Breusch, T.S., Pagan, A.R., 1980. The Lagrange Multiplier Test and Its Applications to Model Specification in Econometrics. Rev. Econ. Stud. 47, 239–253.

Capasso, S., Malanima, P., 2007. Economy and Population in Italy:1300-1900. Pop. E Stor. 2/2007, 15–40.

Chudik, A., Pesaran, M.H., Tosetti, E., 2011. Weak and strong cross-section dependence and estimation of large panels. Econom. J. 14, 45–90. https://doi.org/10.1111/j.1368-423X.2010.00330.x

Cipolla, C.M., 1964. Economic Depression of the Renaissance? Econ. Hist. Rev. 16, 519–524. https://doi.org/https://doi.org/10.1111/j.1468-0289.1964.tb01745.x

Clark, B.G., 2007. The long march of history: Farm wages, population, and economic growth, England 1209–1869. Econ. Hist. Rev. 60, 97–135.

Clark, G., 2010. The macroeconomic aggregates for England, 1209-2008, Research in Economic History. Elsevier. https://doi.org/10.1108/s0363-3268(2010)0000027004

Ditzen, J., 2018. Estimating Dynamic Common-Correlated Effects in Stata. Stata J. 18, 585–617. https://doi.org/10.1177/1536867x1801800306

Eberhardt, M., Bond, S., 2009. Cross-Section Dependence in Non-stationary Panel Models : A Novel Estimator, Munich Personal RePEc Archive (MPRA) Paper No. 17692.

Ferguson, N.M., Laydon, D., Nedjati-Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunubá, Z., Cuomo-Dannenburg, G., Dighe, A., Dorigatti, I., Fu, H., Gaythorpe, K., Green, W., Hamlet, A., Hinsley, W., Okell, L.C., Van Elsland, S., Thompson, H., Verity, R., Volz, E., Wang, H., Wang, Y., Gt Walker, P., Walters, C., Winskill, P., Whittaker, C., Donnelly, C.A., Riley, S., Ghani, A.C., 2020. Report 9 - Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand, Imperial College COVID-19 Response Team. https://doi.org/10.25561/77482

Furceri, D., Loungani, P., Ostry, J.D., Pizzuto, P., 2020. Will Covid-19 affect inequality? Evidence from past pandemics, in: Covid Economics. pp. 138–57.

Galletta, S., Giommoni, T., 2020. The Effect of the 1918 Influenza Pandemic on Income Inequality: Evidence from Italy. Covid Economics, 33, 73-109.

Covid Economics 52, 15 October 2020: 96-117

Herlihy, D., 1978. The Distribution of Wealth in a Renaissance Community: Florence 1427, in: In Towns in Societies: Essays in Economic History and Historical Sociology, Edited by Philip Abrams and Edward Anthony Wrigley. pp. 131–157.

ILO, 2020a. ILO Monitor 2nd edition: COVID-19 and the world of work, International Labour Organization. https://doi.org/10.18356/ba5cc386-en

ILO, 2020b. COVID-19 and the world of work.Fifth edition, International Labour Organization.

IMF, 2020a. A Crisis Like No Other, An Uncertain Recovery. World Econ. Outlook Updat. June.

IMF, 2020b. World Economic Outlook: The Great Lockdown, International Monetary Fund.

Jordà, Ò., Singh, S.R., Taylor, A.M., 2020. Longer-run Economic Consequences of Pandemics, Covid Economics, 1, 1-15.

Kao, C., 1999. Spurious regression and residual-based tests for co-integration in panel data. J. Econom. 90, 1–44. https://doi.org/10.1016/S0304-4076(98)00023-2

Li, D., Li, N., 2017. Moving to the right place at the right time: Economic effects on migrants of the Manchuria Plague of 1910–11. Explor. Econ. Hist. 63, 91–106. https://doi.org/10.1016/j.eeh.2016.11.001

McKibben, W., Fernando, R., 2020. The Global Macroeconomic Impacts of COVID-19: Seven Scenarios, Covid Economics, 10, 116-156.

McKibbin, W.J., Sidorenko, A.A., 2006. Global Macroeconomic Consequences of Pandemic Influenza, CAMA Working Papers 2006-26, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.

Milanovic, B., 2016. Income inequality is cyclical. Nature. https://doi.org/10.1038/537479a

Pamuk, Ş., 2007. The Black Death and the origins of the 'Great Divergence' across Europe, 1300–1600. Eur. Rev. Econ. Hist. 11, 289–317. https://doi.org/10.1017/S1361491607002031

Pamuk, Ş., Shatzmiller, M., 2014. Plagues, wages, and economic change in the Islamic Middle East, 700-1500. J. Econ. Hist. 74, 196–229. https://doi.org/10.1017/S0022050714000072

Covid Economics 52, 15 October 2020: 96-117

Pesaran, M.H., 2007. A Simple Panel Unit Root Test in the Presence of Cross-Section Dependence. J. Appl. Econom. 22, 265–312. https://doi.org/10.1002/jae.951

Pesaran, M.H., 2006. Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure. Econometrica 74, 967–1012.

Pesaran, M.H., 2004. General Diagnostic Tests for Cross Section Dependence in Panels, CESifo Working Paper Series No. 1229, IZA Discussion Paper No. 1240. https://doi.org/10.1103/PhysRevB.85.241102

Pesaran, M.H., Smith, R., 1995. Estimating long-run relationships from dynamic. J. Econom. 68, 79–113.

Pesaran, M.H., Ullah, A., Yamagata, T., 2008. A bias-adjusted LM test of error cross-section independence. Econom. J. 11, 105–127. https://doi.org/10.1111/j.1368-423X.2007.00227.x

Piketty, T., 2014. Capital in the Twenty-First Century, Harvard University Press. Cambridge.

Piketty, T., Saez, E., 2014. Inequality in the Long Run. Science (80-.). 344, 838–843.

Piketty, T., Saez, E., 2003. Income Inequality in the United States, 1913-1998. Q. J. Econ. CXVIII, 1–39. https://doi.org/10.1162/00335530360535135

Piketty, T., Zucman, G., 2014. Capital is Back: Wealth-Income Ratios in Rich Countries 1700-2010. Q. J. Econ. 129, 1255--1310. https://doi.org/10.1093/qje/qju018.Advance

Rachel, L., Smith, T.D., 2017. Are Low Real Interest Rates Here to Stay? Int. J. Cent. Bank. 13, 1–42.

Ramsey, F.P., 1928. A Mathematical Theory of Saving. Econ. J. 38, 543–559. https://doi.org/10.4324/9781315199818-8

Schmitt-Grohé, S., Teoh, K., Uribe, M., 2020. Covid-19: Testing Inequality in New York City. Covid Economics, 8, 27-43.

World Bank, 2020a. Global Economic Prospects. https://doi.org/10.1596/978-1-4648-1553-9

World Bank, 2020b. Projected poverty impacts of COVID-19 (coronavirus).

You, J., Dutt, A.K., 1996. Government debt, income distribution and growth. Cambridge J. Econ. 20, 335–351.

Appendix:

Table A1: Descriptive statistics

Variable	*Obs*	*Mean*	*Std. Dev.*	*Min*	*Max*
*INQ**	412	0.3614	0.0529	0.2778	0.5046
*INQ***	4851	0.3758	0.63603	0.1819	0.9696
GDPPC	412	20978.5	12846.3	3941	53382.7
POP	412	94426.6	70652.4	38542	324985.5

*Note: * and ** represent Top 10% national income share for panel 1 (4 countries) and panel 2 (49 states) respectively. GDPPC and POP are for panel 1.*

Covid Economics 52, 15 October 2020: 96-117

CEPR PRESS

Covid Economics Issue 52, 15 October 2020

Worker productivity during lockdown and working from home: Evidence from self-reports[1]

Ben Etheridge,[2] Li Tang[3] and Yikai Wang[4]

Date submitted: 10 October 2020; Date accepted: 13 October 2020

We examine self-reported productivity of home workers during lockdown using survey data from the UK. On average, workers report being as productive as at the beginning of the year, before the pandemic. However, this average masks substantial differences across sectors, by working-from-home intensities, and by worker characteristics. Workers in industries and occupations characterized as being suitable for home work according to objective measures report higher productivity on average. Workers who have increased their intensity of working from home substantially report productivity increases, while those who previously always worked from home report productivity declines. Notable groups suffering the worst average declines in productivity include women and those in low-paying jobs. Declines in productivity are strongly associated with declines in mental well-being. Using stated reasons for productivity declines, we provide evidence of a causal effect from productivity to well-being.

1 This paper draws on data from Understanding Society, distributed by the UK Data Service. Understanding Society is an initiative funded by the Economic and Social Research Council and various Government Departments, with scientific leadership by the Institute for Social and Economic Research, University of Essex, and survey delivery by NatCen Social Research and Kantar Public. Preliminary working paper version; comments welcome. All errors remain the responsibility of the authors.

2 Lecturer in Economics, University of Essex.
3 PhD student in Economics, University of Birmingham.
4 Lecturer in Economics, University of Essex.

Covid Economics 52, 15 October 2020: 118-151

CEPR PRESS

Covid Economics 52, 15 October 2020: 118-151

1 Introduction

The Covid-19 pandemic has caused widespread disruption to working practices. The most noticeable change has been the vast increase in working from home. The share of the labour force working from home increased from around 5% to over 40% in the U.S. during the lockdown (Bloom, 2020). While some changes to working practices are probably temporary, many could very likely be persistent. Even after the pandemic ends, home working in particular is expected to be much more prevalent than previously.[1]

A key policy issue, therefore, both in the near and the far term, is how these changes in labour practices impact worker productivity. Despite previous research on the effects of working from home (Bloom et al., 2015), given the size of the changes seen during the pandemic the evidence base is inevitably thin. Most research since the onset of the pandemic has focused on characteristics of jobs through objective measures such as those provided by O*NET.[2] There is little direct evidence on productivity in the new working environment and how it varies not only across job types, but also worker characteristics.

In this paper we use the Covid-19 module from the UK Household Longitudinal Survey (UKHLS), which provides representative data on home workers' self-reported productivity towards the end of the lockdown period in the UK, in June 2020.[3] In this survey, anyone working from home (WFH) at least some of the time was asked about changes in their productivity since before the pandemic period, at the beginning of the year. These data allow us to examine how productivity changes vary across job and worker types and are influenced by, for example, the presence of children. The advantage of using individual-level reported productivity over data obtained from, say, characteristics of jobs, is that we obtain a more direct measure of the key object of interest. The advantage of using individual-level over aggregate data reported in national statistics is that we can examine the rich causes of productivity changes at the micro level, as well as examining effects on other outcomes of interest. Overall we find that workers report being approximately as productive as before the pandemic, on average. However, productivity varies substantially across socioeconomic groups, industries and occupations.

In more detail, we find that workers in industries and occupations that are less suitable for working from home report lower productivities than before the pandemic. Consistently with this, and with the literature, females and low earners also report lower productivity at home on average. The opposite

[1] Again, see Bloom (2020). For a wider discussion see also the dedicated discussion of the literature below.

[2] See, for example, Dingel and Neiman (2020) and Mongey et al. (2020).

[3] In the UK, the official 'lockdown' began on March 23 when a widespread stay-at-home order was introduced. The lockdown eased as the incidence of Covid declined, over May and June. On June 1, restrictions were lifted which allowed people to meet with up to six others from separate households in outdoor places. An accepted date for the end of lockdown is July 4, when many businesses, especially in retail and food were allowed to re-open.

CEPR PRESS

types of workers, e.g., those in the "right" occupations and with high incomes, report higher productivities than previously. More specifically, we incorporate external measures of feasibility of home work from Dingel and Neiman (2020); Adams-Prassl et al. (2020b), and need for physical proximity to others from Mongey et al. (2020). The sector-level correlations between our reported productivity changes and these job-based measures are always of the expected sign. In fact they are higher when comparing occupations than industries: For example, the correlation with feasibility of home work across occupations is 0.56, and across industries it is 0.23. This difference suggests that while occupational job characteristics provide quite accurate information about the impact of working from home on productivity, the industry characteristics are more noisy; it is at the job-task level that most impacts of the pandemic have been felt. Our direct measure of productivity changes allow us to understand how well those measures — feasibility of home working and physical proximity — capture the realized productivity changes in different contexts.

In addition, workers' productivity changes correlate with other aggregate outcomes: occupational job losses recorded in early lockdown, and aggregate labor productivity changes at the industry level.

We then examine individual characteristics in further detail. Females, low earners and the self-employed report worse productivity outcomes than their counterparts. Their low productivity is not only related to their job characteristics, but is also directly affected by their socioeconomic conditions. For example, while females are more likely to work in occupations less suitable for home work (Adams-Prassl et al., 2020b), productivity of females is also more negatively affected by the presence of children. This finding shows the strength of the measure used here over those based purely on characteristics of the job.

Third, we find that home workers' productivity during the lockdown is related to the *intensity* of working from home and its change since the prior period. Those who previously worked from home at least sometimes and then increased the intensity of home-working experienced a productivity increase. Those who did not increase their home-working frequency or never worked from home before the pandemic report a large productivity decline. This pattern is partly explained by the occupational characteristics of the jobs in each category. However it also suggests two countervailing forces: a positive productivity effect of increased home working alongside a direct negative effect of the pandemic itself. The productivity decline reported by those who have always worked from home is evidence of this latter phenomenon.

A noteworthy feature of the pandemic period has been a decline in mental well-being, observed particularly in the UK (Banks and Xu, 2020; Etheridge and Spantig, 2020). We therefore assess the association of workers' mental well-being with productivity changes. We find strong correlations between the two: those who state they get much less done at home report declines in well-being comparable

Covid Economics 52, 15 October 2020: 118-151

Covid Economics 52, 15 October 2020: 118-151

to the effect of an unemployment shock. We also find evidence of a causal effect from productivity to well-being: using ineffective equipment as an instrument for productivity declines, we find a 1 standard deviation lower productivity causes a 0.24 standard deviation lower mental well-being, as measured by general health questionnaire scores. This result is consistent with Etheridge and Spantig (2020) who find that females and low income groups have experienced large deteriorations in mental well-being compared to their counterparts. Our paper therefore offers a novel explanation for the recent declines in mental well-being among certain groups. It also suggests that policies that target workers in the vulnerable socioeconomic groups or certain jobs with large productivity drops may not only boost productivity but also mental well-being on aggregate.

Before proceeding to the general literature review, we briefly discuss the most relevant papers. Our findings on homeworking productivity during the lockdown in the UK are closely related to Felstead and Reuschke (2020). They utilize the same survey data to document overall patterns of homeworking in the UK, including homeworking intensity, productivity and mental well-being. Our paper focuses on understanding the heterogeneity of productivity and how it is related to workers' personal characteristics (gender, income, childcare, etc.) and more importantly, job characteristics (industries and occupations). Therefore, our study is also closely linked to the literature pioneered by Dingel and Neiman (2020), which use job characteristics to construct feasibility and predicted productivity of homeworking, and we contribute by measuring the realized productivity by sectors and documenting how it relates to the constructed measures.

The remainder of the paper is organized as follows. In the rest of this section, we review the relevant literature. Section 2 describes data used in this paper. Section 3 presents and discusses the results. Section 4 concludes.

Related Literature

Our paper is related to four strands of literature: (1) working from home as an alternative practice; (2) sector-specific productivity changes and optimal policies during the current pandemic; (3) inequality across gender and socioeconomic groups, especially during difficult times such as the current pandemic and other recessions; and (4) mental well-being during the current pandemic.

First, working from home and its impact on productivity have been getting increasing attention in recent years, and especially since the Covid-19 outbreak. Bloom et al. (2015) study workers' productivity and attitude towards working from home using a random experiment on call-center workers in a Chinese travel agency. They find that home-working led to a 13% performance increase and that, after the experiment, over half of the workers chose to switch to home-working. Battiston et al. (2017) con-

Covid Economics 52, 15 October 2020: 118-151

duct an experiment at the Operational Communications Branch of a police station. They conclude an opposite finding: face-to-face working increases productivity, and attribute the contrast to the complexity of the tasks in their experiment. Dutcher (2012) recruit a group of students as experiment subjects and find that working from home improves productivity of creative tasks but compromises productivity of dull tasks. From the survey of Upwork Future Workforce Report, Ozimek (2020) find that more hiring managers, including executives, VPs, and managers, report productivity increases from home working than report productivity decreases. While Bloom et al. (2015), Battiston et al. (2017) and Ozimek (2020) focus on a few occupations, the Covid-19 outbreak and the lockdowns in many countries have dramatically increased the prevalence of working from home in almost all occupations. Since the home-working style is likely to persist even after the Covid-19 pandemic ends, this paper, by looking at productivity changes by sectors (occupations and industries), provides relevant policy insights, such as the possible productivity outcomes after implementation of home working across different sectors, and a direct way to model supply shocks in multi-sectoral macroeconomic models.

Felstead and Reuschke (2020) use the survey data in the UK and document that while 5% of workers worked from the home before the pandemic, the share increased to 45% in April 2020, remaining high thereafter. They find that workers' productivity at home on average does not change much during the pandemic. The same patterns — increasing home-working and not much change in workers' average productivity at home — are also found in Europe and North America (see Rubin et al., 2020 for Netherlands, Eurofound, 2020 for the Europe as a whole, and Bartik et al., 2020; Brynjolfsson et al., 2020 for the US). In Japan, where working culture and organization structures are quite different, workers report that their productivities at home during the pandemic are on average 30–40% lower than before (Masayuki, 2020). Our findings confirm this average productivity pattern in the UK and other western countries, and, on top of it, illustrate the underlying heterogeneity in productivity change across sectors and by some personal characteristics of interest.

The second strand of the literature is the sector-specific productivity of working from home, and optimal sectoral policies. The existing papers pioneered by Dingel and Neiman (2020) use characteristics of jobs to provide predictions on home-working productivities across occupations and industries. Dingel and Neiman do this by constructing a measure of feasibility to work from home across industries and occupations using data from O*NET. Adams-Prassl et al. (2020b) follow this by eliciting a conceptually similar measure derived using individual self-reports. Again similarly, Mongey et al. (2020) also use O*NET to construct a measure of need for physical proximity to co-workers to carry out one's work effectively. The direct evidence of productivity changes provided in the current paper can be used to understand how well the measures constructed from job characteristics capture the real productivity changes across sectors, and can potentially be used in macro models of the pandemic with

CEPR PRESS

sector-specific shocks and optimal policies.

In this way, estimates of productivity changes by sector are important for macroeconomic models that try to capture the sectoral and aggregate labor and output changes during the Covid-19 pandemic, e.g., Baqaee and Farhi (2020). Bonadio et al. (2020) study the impact of the Covid-19 pandemic on GDP growth and the role of the global supply chains. They discipline the labor supply shock across sectors using the fraction of work that can be done from home across generations measured by Dingel and Neiman. While the correlation of this measure with our measure of realized labor productivity is reasonably high, there is space for improvement by obtaining better measures of realized labor productivity changes.

Third, the differential impacts of working from home across sectors and socioeconomic groups implies that inequality is strongly affected by enforced home working in the pandemic. Income inequality has also been increasing since the 1980s both in the U.S. (Heathcote et al., 2010), and in the UK (Blundell and Etheridge, 2010). Inequality has often been found to increase during recessions (see Perri and Steinberg, 2012 for a discussion of the great recession after 2008). In the current pandemic, it is also the economically disadvantaged groups, such as low-income groups and females, that are suffering larger declines in economic outcomes.

In this vein Alon et al. (2020) study the potentially different impacts of Covid-19 pandemic on the employment of men and women given the gender differences in occupation and childcare. They predicted that women's employment would suffer disproportionately. Adams-Prassl et al. (2020b) document that female workers report a lower ability to work from home, and Adams-Prassl et al. (2020a) document that women are more likely to lose their jobs in the UK and in the US (though not in Germany, around early April 2020). They also find worse outcomes for lower earners. Our paper contributes to this strand of the literature by studying inequality of worker productivity across gender and socioeconomic groups. Our findings confirm the prediction of this literature: Females and low income groups have suffered larger productivity declines while working from home during the lockdown, indicating an increase in inequality.

The fourth strand of related literature is that on mental well-being during Covid-19. Early in the pandemic, international organizations and researchers warned about the resulting psychological effects (Holmes et al., 2020; World Health Organization, 2020). The pandemic imposes large risks and potential damages to mental well-being through a variety of channels. Anxiety is caused by the disease's spread: Fetzer et al. (2020) conduct a survey covering 58 countries and show, by exploiting time variation in country-level lockdown announcements, that people's perception of the spread of the disease causes lower mental well-being. Lower mental well-being is also caused by adverse economic shocks (see Chang et al., 2013; Dagher et al., 2015 for the 2008 recession, and Janke et al., 2020 for UK during

Covid Economics 52, 15 October 2020: 118-151

2002–2016). Finally, loneliness and social isolation can be induced by quarantine (Brooks et al., 2020) and lockdown (Brodeur et al., 2020; Knipe et al., 2020; Tubadji et al., 2020).

Banks and Xu (2020) and Etheridge and Spantig (2020) document decline in mental well-being during the Covid-19 pandemic in the UK using the same dataset as the current paper. We add to these papers by documenting an association between mental well-being and worker productivity. More widely the literature on the relation between economic conditions and mental well-being is vast; see for example, Janke et al. (2020) who study how macroeconomic conditions affect health condition, especially mental health conditions, using British data over the period 2002–2016.

2 Data

We use the Covid-19 module from the UK Household Longitudinal Survey (UKHLS), administered monthly from April 2020. The analysis makes specific use of the Covid module's third wave, conducted in June, which includes questions on self-reported productivity. These interviews were conducted in the seven days from Thursday June 25, with around 75% of interviews completed within the first three days. We merge these data with the April and May waves of the Covid module as well as with wave 9 of the 'parent' UKHLS (also known as 'Understanding Society'), a large-scale national survey administered yearly from 2009. Wave 9 of the parent survey was itself administered between 2017 and 2019.

The UKHLS Covid module is conducted as a web survey. The underlying sampling frame consists of all those who participated in the UKHLS main survey's last two waves. To conduct the fieldwork, the sample was initially contacted using a combination of email, telephone, postal and SMS requests. Of those eligible, and who responded to the main survey wave 9, the response rate was a little under 50%. To adjust our analysis for non-response, we use the survey weights provided. In addition, to allow for the fact that many respondents are related either through primary residence or through the extended family, we cluster all regressions at the primary sampling unit level. For a further discussion of the Covid module and underlying UKHLS design see (Institute for Social and Economic Research, 2020).

The main variable of interest is self-reported productivity while working from home and compared to a stated baseline. To elicit this the survey includes a bespoke question. Precisely, respondents are asked as follows:

"Please think about how much work you get done per hour these days. How does that compare to how much you would have got done per hour back in January/February 2020?"

If the respondent did not work from home before the pandemic, then the question ends with:

"...when, according to what you have previously told us, you were not working from home?"

Covid Economics 52, 15 October 2020: 118-151

Interviewees are then asked to respond on a Likert-type scale of 1 to 5 ranging from "I get much more done" to "I get much less done".

We transform the variable as follows: we invert so that responses are increasing in productivity; we re-centre so that the response "I get about the same done" is valued 0, and we divide the distribution by its standard deviation. In this way the mean response across the population can be interpreted in terms of standard deviations away from a neutral effect. When discussing results we sometimes term the resulting variable as a 'semi-standardized' productivity change.

It is worth discussing the question and resulting data in more detail. Notice first that the question explicitly attempts to ask about productivity per hour, and so corresponds to a concept of labour productivity. We examine the relationship between the variable here and aggregate productivity data from the National Accounts in more detail in section 3. Notice further that the question actually makes no reference to working from home itself, except in the qualifier referencing prior working location. In principle, therefore, this question could be asked of workers in any location. It was in fact only asked of those working from home to save valuable survey time. In future waves it is hoped this question is asked of all respondents. Most importantly, perhaps, it should be remembered that the scale is ordinal. As with all similar Likert-type scales, however, it is anchored with a natural reference point at 3, and responses above or below can be considered as improvements or declines compared to the pre-Covid period. In this paper we typically use simple means, effectively re-interpreting the scale as cardinal. For much of our analyses, however, we provide parallel results using ordered probit models in Appendix B, where we show that marginal effects computed this way are nearly identical.

We make use of much auxiliary information contained in the surveys. Of particular interest, all respondents were asked to report their baseline earnings and place of work just before the pandemic, in 'January/February'. The survey elicits industry of work both in the baseline period and currently. Unfortunately, the Covid survey does not elicit information on occupation. For this we use occupational information from wave 9, which relates to the job performed in either 2017 or 2018, whenever that wave's interview was performed. For occupation we make additional use of metrics obtained elsewhere in the literature which have typically been collected using the classification used in the US-based O*NET. We therefore typically convert our occupational information to this alternative using our own cross-walk. Our procedure is described in Appendix A. Finally, we also use productivity data from the UK Office for National Statistics; see Appendix B for a further discussion.

For mental well-being we use a Likert well-being index derived from the 12 questions of the General Health Questionnaire (GHQ-12). The GHQ battery asks questions regarding, for example, the ability to concentrate, loss of sleep and enjoyment of day-to-day activities. The GHQ questionnaire has been administered in all waves of UKHLS in exactly the same form, allowing us to examine changes in

Covid Economics 52, 15 October 2020: 118-151

Covid Economics 52, 15 October 2020: 118-151

well-being from a base period. We use a standardized and inverted index so that higher scores indicate higher well-being. Here we exactly follow the procedures in Etheridge and Spantig (2020); see that paper for further details.

Our total number of adjusted interviews in the June module is $11,496$. Of these interviews $6,504$ individuals were in work and reported information about working location. Of these the number who answered the question about productivity changes was $3,411$.

3 Results

3.1 Patterns of Working from Home

The largest change in working conditions during the pandemic has been the increased prevalence of working from home. We accordingly show patterns of home work over time in Table 1. To show some of the wide variation during the pandemic, we show a breakdown by industry. This variation has also been documented by Felstead and Reuschke (2020), among others. The first column reports baseline home work patterns in January/February, before the pandemic, and documents the proportion of workers who worked at home at least some of the time. The second column shows the proportion of workers in this category in April, at the height of the lockdown period. It shows a very large increase in the proportion working from home across almost all industries. There are, however, some exceptions: in 'Accommodation and Food Service', for example, the effect of the lockdown was seen not so much in an increase in home work, but rather widespread job losses. The third column then records the change in proportion of home workers from April to June. It shows there was very little change in working patterns by this metric even as the lockdown eased.

The final two columns of Table 1 show the proportion of respondents *always* working from home. Here we don't show results for the baseline because, in most industries, the numbers were small. The fourth column shows that in some sectors, such as 'Information and Communication', a large proportion of workers relocated to home permanently in April. By June, the proportion of workers always at home had declined slightly (Felstead and Reuschke, 2020). This is only slightly evident in the industry breakdown shown in column 5. However, one example stands out: a noticeably higher fraction of teachers worked away from home at least some of the time as schools partially reopened before the summer vacation. Table B.1 shows proportions of working from home by occupation, using reported occupation from wave 9 of the main survey. Similar patterns are seen as with industry, with the major change from spring to summer in occupations relating to teaching.

Covid Economics 52, 15 October 2020: 118-151

Table 1: Proportions of Working from Home: By Industry

	At least some of the time			Always	
	Jan/Feb	April	April to June	April	April to June
Agriculture, Forestry and Fishing	0.23	0.29	-0.03	0.15	-0.05
Mining and Quarrying	0.21	0.54	-0.06	0.47	0.00
Manufacturing	0.16	0.36	-0.05	0.23	-0.05***
Electricity and Gas	0.36	0.54	0.05*	0.48	-0.03
Water Supply and Sewerage	0.30	0.70	-0.02	0.45	0.01
Construction	0.24	0.37	-0.03	0.24	-0.03
Wholesale and Retail Trade	0.13	0.19	-0.02	0.10	-0.01
Repair of Motor Vehicles and Motorcycles	0.25	0.21	0.01	0.16	-0.08
Transportation and Storage	0.11	0.17	0.01	0.11	-0.00
Accommodation and Food Service	0.11	0.10	0.02	0.06	-0.00
Information and Communication	0.62	0.86	-0.01	0.75	-0.01
Financial and Insurance	0.48	0.86	-0.00	0.73	0.04
Real Estate Activities	0.45	0.71	0.04	0.40	-0.02
Professional and Technical	0.56	0.82	-0.02	0.67	-0.04
Administrative and Support Service	0.31	0.62	-0.04	0.47	0.00
Public Administration and Defence	0.38	0.67	0.01	0.49	0.02
Education	0.31	0.72	-0.01	0.44	-0.10***
Human Health and Social Work	0.25	0.39	-0.01	0.18	-0.02
Arts, Entertainment and Recreation	0.55	0.65	-0.04	0.51	-0.05
Other Service Activities	0.32	0.46	-0.03*	0.32	-0.02
Activities of HHs as Employers	0.18	0.22	0.04	0.12	-0.02
Observations	5601	5486	5475	5486	5475
Adjusted R^2	0.369	0.605	0.002	0.461	0.014

Source: UKHLS Covid module

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, used in third and fifth columns only.

Note: This table reports proportions of respondents who WFH for at least some time and always, respectively for each industry in the United Kingdom in 2020. The first two columns illustrate the proportions who spent at least some of the time WFH in January/February, and in April, respectively. The third column reports changes in proportions of WFH at least for some time from April to June. The last two columns demonstrate the proportion of always WFH in April and change in the proportion of always WFH from April to June, respectively. Standard errors omitted.

CEPR PRESS

Covid Economics 52, 15 October 2020: 118-151

3.2 Changes in Productivity by Basic Characteristics

We now document the changes in productivity reported in the June survey module, and for those working at home at least some of the time. We first document average changes according to characteristics of the worker. Our evidence is presented in Table 2. The table's first column examines the relationship between productivity changes and earnings, with workers split into terciles according to take home pay across the whole labour force in the baseline period. It seems the lowest earning group faced the worse decline in productivity on average, while productivity of top earners has been boosted significantly. As discussed in Section 2, the data here come from an ordinal Likert scale. In Table 2, as in the rest of the analysis, we construe responses as cardinal and interpret marginal effects in terms of standard deviations away from no productivity change. We provide robustness to these results in Appendix B where we perform the same analysis using ordered probits, with near identical results.

Despite the gradient by earnings, column two of Table 2 shows that on average productivity changes are not significantly dependent on degree holding itself. Although not shown here, productivity is also not noticeably different across age. The third column then illustrates a gender gap: on average females experienced a significant productivity fall, whereas males were not noticeably impacted. A possible cause for this is the unequal burden of home work, childcare and other distractions (Andrew et al., 2020). However, in terms of preliminary evidence here, the fourth column shows that productivity is not noticeably affected by the presence of children, at least not across the population as a whole. The final column shows that the self-employed group experiences a significantly worse productivity loss than employees. One important reason is that many self-employed were already in their ideal working environment before the pandemic, so they endured the negative effect of Covid, but did not feel the positive effect of relocating to a more productive space. For example, in January 2020, already 24.2% of self-employed worked at home. Though the fraction increased to 36.4% in April 2020, the increase is much smaller than that of the employed — from 3.8% to 34.5%.

To explore the gender divide in reported productivity in further detail, we present results broken down by gender together with other characteristics in Table 3. Now columns 1 and 2 do indeed show an effect of the presence of children: females with childcare duty suffer a significant loss in productivity, while males are not so affected. This analysis demonstrates one of the strengths of our metric over and above those used elsewhere in the literature, which typically focus on properties of the job specifically: our results indicate an important role for the circumstances of the individual over and above the pure effect of the job they are matched to. Turning to skill level, columns 3 and 4 again show differential effects across gender: on average, the productivity of females in the bottom earnings tercile fell significantly, whereas the productivity of males with the high (top) level of earnings increased noticeably.

CEPR PRESS

Covid Economics 52, 15 October 2020: 118-151

Table 2: Productivity Changes During Covid19: By Characteristics

	(1)	(2)	(3)	(4)	(5)
Earnings: Bottom	-0.29***				
	(0.06)				
Middle	-0.03				
	(0.05)				
Top	0.07**				
	(0.04)				
Education: No Degree		-0.04			
		(0.04)			
Degree		-0.01			
		(0.03)			
Gender: Male			0.05		
			(0.04)		
Female			-0.09**		
			(0.03)		
Children: None				-0.01	
				(0.03)	
At least one				-0.09	
				(0.07)	
Employment: Employed					0.02
					(0.03)
Self-employed					-0.31***
					(0.07)
Observations	2912	3254	3034	3254	3067
Adjusted R^2	0.019	0.000	0.004	0.000	0.011

Source: Covid module of UKHLS
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Note: Standard errors in parentheses. Table displays group means of variable of interest, which is semi-standardized productivity change in June 2020 compared to Jan/Feb 2020. The first column reports the changes in productivity for respondents grouped into tertiles of earnings reported for Jan/Feb. The fourth column is an indicator for the presence of children in the house. See text for more details.

CEPR PRESS

Covid Economics 52, 15 October 2020: 118-151

Although this analysis is very broad brush, it indicates, over and above the results for the presence of children, an important role for the different types of jobs that males and females are matched to across the earnings distribution. As the literature has emphasized, therefore, it is important to examine the characteristics of jobs themselves.

Table 3: Productivity Changes by Gender and Other Characteristics

	Children: Male	Children: Female	Earnings: Male	Earnings: Female
Children: None	0.05	-0.07*		
	(0.04)	(0.04)		
At least one	0.09	-0.22**		
	(0.10)	(0.09)		
Earnings: Bottom			-0.07	-0.37***
			(0.15)	(0.07)
Middle			-0.06	-0.01
			(0.08)	(0.06)
Top			0.10**	0.04
			(0.05)	(0.06)
Observations	1244	1790	1102	1619
Adjusted R^2	-0.001	0.001	0.005	0.025

Source: Wave 9 and Covid module of UKHLS.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Note: Standard errors in parentheses. This table reports grouped mean of variable of interest, which is semi-standardized productivity change over Jan/Feb to June 2020. See text for more details. Children is an indicator for the presence of children in the house. Last two columns report changes for individuals grouped into tertiles according to earnings reported in Jan/Feb.

3.3 Productivity Changes by Job Characteristics

We now examine reported productivity changes, focusing on characteristics of the job. As above, we first examine differential performance across industries. Industry-specific policy has been exploited already in the pandemic, such as with the 'Eat Out to Help Out' policy instigated in the UK in August, targeted at the restaurant sector. More generally, commentators and researchers have observed the wide differential impacts by sector. Baqaee and Farhi (2020), for example, examine changes in hours by industry and show that such sector-specific supply shocks, together with demand shocks, are necessary for capturing the disaggregated data on GDP, inflation and unemployment. They further show how a multi-sector Keynesian framework can be used to design optimal monetary policies.

Covid Economics 52, 15 October 2020: 118-151

Average productivity changes by industry are shown in the left sub-plot of Figure 1, which plots the 21 industries recorded in the survey ranked by average performance. The figure shows that productivity declines are largest for those working in 'Repair of Motor Vehicles', at least for those individuals doing some work from home. The magnitude of the decline is large, averaging one standard deviation of the entire distribution of reported changes. Other industries which show a decline that is statistically significant include 'Education', which was transformed by the pandemic, and arts-related activities. This latter industry is an interesting case: While the realized productivity change in this industry is negative (as reported both in our household data and official aggregate productivity statistics), job characteristics themselves predict a large fraction of jobs in this industry can be done at home (see for example, Adams-Prassl et al., 2020b).

The left sub-plot of Figure 1 also shows industries for which workers report productivity increases. As one might expect, these include jobs in both the IT and finance sectors, which external metrics indicate require less face-to-face interaction. The other two sectors which report significant productivity increases are trade, and transport and storage. Although jobs in these sectors are less able to be performed at home than those in, say, IT, they do not require physical proximity to other individuals. These observations indicate that there are multiple reasons why productivity may change after work is re-arranged. Again we explore these points in further detail below.

The right sub-plot of Figure 1 shows average productivity changes by occupation. Here we take reported occupation stated in wave 9 as baseline and categorize workers using the 22 two-digit O*NET codes. As explained in Appendix A, the two-digit O*NET codes are derived by using a cross-walk to convert the 3-digit SOC 2000 codes contained in the UKHLS.[4] Looking at the top of the sub-plot, the occupation that shows the largest productivity increase, 'Computer/Mathematical', is similar to the IT industrial sector in requiring little face-to-face interaction. The next occupation, 'Management', is an interesting case, given that it is one of the job types requiring the most interaction on most measures.[5] That managers report productivity increases is possibly very dependent on the current state of information technology. Very likely, if the pandemic had occurred 10 or 20 years previously, the ranking of occupations would look different. Looking at the bottom of the sub-plot, again some expected occupations, such as 'Personal Care', and 'Education' show productivity declines.

[4] For practical survey reasons, occupational data were not collected in the Covid module.

[5] For 'Management', our results are consistent with Ozimek (2020), which shows more managers report productivity increases than report decreases, from the survey of Upwork Future Workforce.

CEPR PRESS

Covid Economics 52, 15 October 2020: 118-151

Figure 1: Mean Productivity Change, by Industry and by Occupation

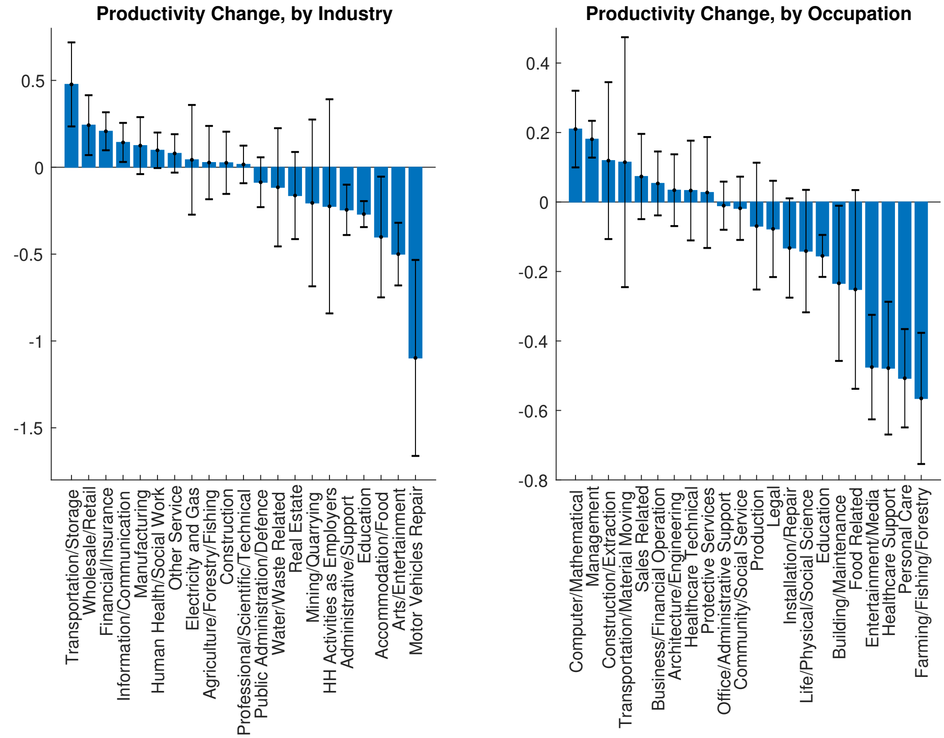

Source: Wave 9 and Covid module of UKHLS
Note: Bars represent the average semi-standardized productivity change. See text for more details. Lines illustrate 95% confidence intervals. Industry computed using current industry code. Occupation taken from occupation worked in wave 9. For consistency with other tables and the US-based literature, occupation is converted to a 2-digit O*NET classification.

We next examine how our self-reports of productivity changes relate to other measures of job performance examined in the literature, focusing on variation across occupations and industries. To this end, Figure 2 shows variation across the 21 industry codes, and according to three external measures. The top left subfigure plots our measure of productivity change against a measure of feasibility of working from home, taken from Adams-Prassl et al. (2020b). As discussed in the introduction, they obtain their measure by asking workers to report the fraction of job tasks that can be performed from home. As such, we would expect this feasibility measure to be a key input into observed productivity during the lockdown period. Here we take Adams-Prassl et al.'s industry averages. Indeed we find a positive, albeit weak correlation between this feasibility measure and reported productivity changes,

Covid Economics 52, 15 October 2020: 118-151

with an estimated coefficient, weighted by industry size, of 0.23.

Figure 2: Productivity Changes by Industry, and Industry Characteristics

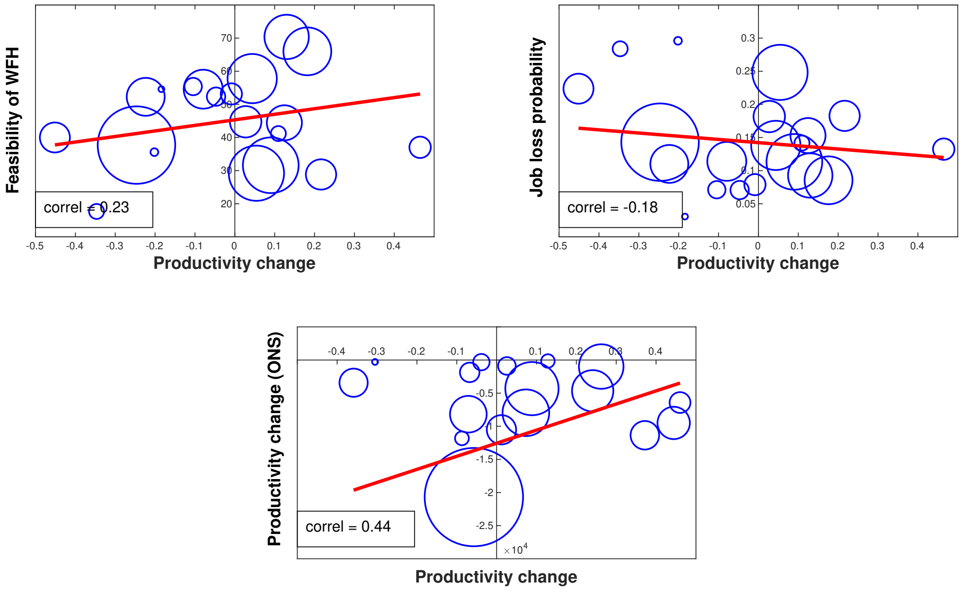

Source: Adams-Prassl et al. (2020a,b), Office of National Statistics (ONS) and wave 9 and Covid module of UKHLS

Note: Figure shows scatter plot of productivity changes against external measures, by industry. Bubble sizes are proportional to industry employment. Solid line is a line of (weighted) best fit. Top left plot uses the feasibility of home work measure of Adams-Prassl et al. (2020b). Top right plot uses industry-specific job loss in April 2020, from Adams-Prassl et al. (2020a). Bottom plot uses aggregate productivity change by industry from 2019Q4 to 2020Q2 from the UK ONS. See text for more details.

The top right sub-plot then shows a comparison of our productivity change measure with job loss by industry, taken from Adams-Prassl et al. (2020a). The definition of job loss used includes anyone detached from their previous job, not including those on furlough. Here, the relationship between our measure and the external indicator is not so clear cut. We would expect those industries where working is more difficult to show more job losses. On the other hand, and theoretically at least, heterogeneity might be important. Adams-Prassl et al. (2020b) report wide dispersion in feasibility of home work within industries, and varying degrees of this dispersion across industries. In some industries we might therefore expect job losses among those who cannot work productively in the new environment, but high productivity among those who stay. Moreover, industries with less labour hoarding should exhibit higher productivity. These latter two effects would induce a relationship between productivity and job

Covid Economics 52, 15 October 2020: 118-151

loss that is positive. Overall, however, we do see a negative correlation between job losses and reported productivity, albeit a weak one.

Finally, the bottom sub-plot of Figure 2 compares our self-reported productivity changes with official aggregate productivity by industry reported by the ONS.[6] For better comparison with the external measure, here we compute a measure of industry-level aggregate productivity change. We do this by weighting the reported changes by earnings reported in January. This weighted correlation coefficient is higher than those in the top two panels, at 0.44. Of course, the discrepancy between our measure and official productivity may be caused by any number of factors. These include biases in self-reporting, and the fact that our data omits those still always working outside the home.

We show variation by occupation in Figure 3.[7] Whereas our industry measure captures current work status, our measure of occupation is taken from wave 9, just before the pandemic. Nevertheless, this measure should capture baseline occupation well; the available evidence suggests there was little noticeable rise in occupational mobility during the Covid-19 period (Office for National Statistics, 2020). As discussed above and in Section 2, the occupational information in UKHLS is provided at the 3-digit SOC 2000 level. In order to compare to measures in the literature we convert to 2-digit O*NET occupations using the cross-walk described in Appendix A.

The top left panel again shows a comparison with feasibility of working from home, taken from Adams-Prassl et al. (2020b). Compared to the equivalent subpanel in Figure 2, the correlation between our measure and the external measure is now stronger, at 0.56. This is perhaps to be expected: feasibility of working from home presumably depends more on occupational rather than on industrial characteristics. The top right sub-plot of Figure 3 also shows the equivalent panel to that shown previously, plotting productivity change against job losses. Now the negative correlation with productivity changes is particularly strong, at -0.67. This indicates that it is at the occupation level that productivity changes determine job losses, rather than at the industry sector level.

The bottom sub-plot of Figure 3 now introduces another metric discussed in the literature that should be related to productivity. It compares our reported changes to a measure of need for physical proximity with others, derived by Mongey et al. (2020) using O*NET descriptors. Again these measures are reported by the authors at the 2-digit O*NET occupation level. Those occupations which are indicated to require close physical interaction between workers, such as ‘Personal Care’ and ‘Health-

[6]The ONS combines three industries, ‘Public Administration and Defense’, ‘Education’, and ‘Human Health and Social Work Activities’ into one category and also combines ‘Other Service Activities’ and ‘Activities of Households as Employers’ into another category. Therefore, for consistency, we combine our industry data similarly.

[7]In this figure, the category ‘Farming, Fishing, and Forestry Occupations’ is dropped, for comparability with Adams-Prassl et al. (2020b).

Covid Economics 52, 15 October 2020: 118-151

Figure 3: Productivity Changes by Occupation, and Occupation Characteristics

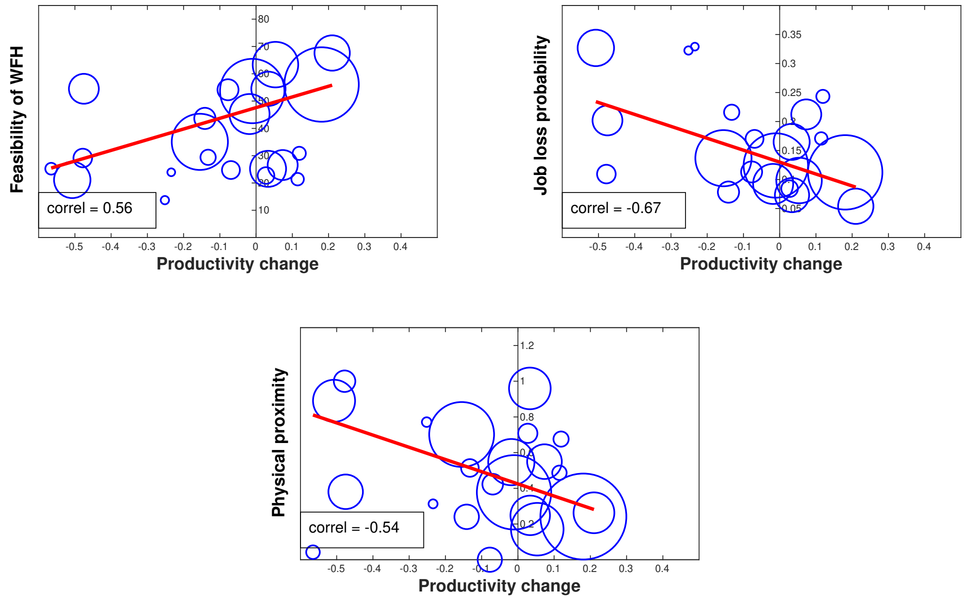

Source: Adams-Prassl et al. (2020a,b), Mongey et al. (2020) and wave 9 and Covid module of UKHLS
Note: Figure shows scatter plot of productivity changes against external measures, by occupation. See text and notes to Figure 2 for details on overall structure of top two plots. Occupation is from wave 9, converted to 2-digit O*NET code. See appendix for full discussion. Bottom plot uses measure of physical proximity in job, from Mongey et al. (2020).

care Support' show the largest productivity declines. In fact the correlation here is also strong, at -0.54, indicating that individual productivity is just as much affected by this factor as pure feasibility of home work.

We finish this subsection by exploring productivity changes by intensity of home working, with results reported in Table 4. In this table, the rows record the intensity of working from home in January/February, and the columns record status in June. Respondents are put into groups by homeworking intensity change.[8] The left panel of the table illustrates average productivity change for each group. The general pattern is the following: If there are large increases in homeworking intensity (from 'Never' or 'Often/Sometimes' to 'Always'), then workers typically report productivity increases; otherwise, i.e. there are little or no increases in the intensity, workers report productivity declines. It is consistent

[8]Note that, those never work from home in June are not asked about their productivity changes and thus 'Never' is omitted in the column dimension.

CEPR PRESS

Covid Economics 52, 15 October 2020: 118-151

with the evidence in Felstead and Reuschke (2020), which shows that 'new home-centred' workers report productivity gains whereas 'established home-centred' workers report productivity losses. The pattern in this table suggests that productivity changes are the net outcome of two off-setting effects. The implementation of nation-wide lockdown was a negative shock to productivity, but increasing homeworking intensity yields positive impacts.

Table 4: Productivity Changes and Feasibility of Working from Home, by WFH Intensity

		Working from home in June			
		Often/Sometimes	Always	Often/Sometimes	Always
		Average productivity change		Average feasibility	
WFH in Jan/Feb	Never	-0.24***	0.08**	40.21	47.48
		(0.05)	(0.04)	(0.35)	(0.27)
	Often/Sometimes	-0.10**	0.11***	42.69	50.97
		(0.05)	(0.04)	(0.38)	(0.22)
	Always	-0.15	-0.26***	42.36	47.24
		(0.16)	(0.06)	(1.52)	(0.50)

Source: Wave 9 and Covid module of UKHLS and Adams-Prassl et al. (2020b)
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$
Note: Standard errors in parentheses. The left panel reports average productivity change by intensity of WFH in Jan/Feb and June 2020. Right panel reports average feasibility of WFH, using data from Adams-Prassl et al. (2020b). This feasibility measure is computed using occupational data from wave 9, converted to 2-digit O*NET occupational categories. See text for more details. Asterisks for statistical significance are omitted in the right panel.

Of course this simple interpretation glosses over other possible explanations, such as that workers in each cell vary systematically by their own or their job characteristics. In this light, the right part of Table 4 reports average feasibility to work from home using the occupation-level measure from Adams-Prassl et al. (2020b). The table shows, as expected, that those working always at home in June are systematically in jobs that are better suited to home work. However, the combination of left and right panels yields an interesting finding: those who always worked from home both before and during the pandemic have experienced productivity declines despite a high score on home-working feasibility. Again this finding suggests a negative direct impact of the pandemic on worker productivities.

Covid Economics 52, 15 October 2020: 118-151

3.4 Productivity Changes and Well-Being

A well-documented feature of the lockdown period has been a noticeable decline in well-being (Banks and Xu, 2020; Etheridge and Spantig, 2020). An association of well-being with changing work patterns, and particularly home work, has been documented by Felstead and Reuschke (2020). Here we examine the association of well-being with productivity changes. Mental health problems are known to adversely affect productivity on the job (Greenberg et al., 2003). It is reasonable to hypothesize that difficulty in performing one's job is a stressor and cause of mental health problems likewise.

Table 5: Reasons for Declines in Productivity: By Gender

	Gender		
Fall in productivity, working at home	male	female	Total
I have had less work to do	30.98	28.99	29.79
I have had to provide childcare/home schooling and/or care for others	22.59	33.37	29.03
The equipment, software and/or internet connection limits what I can do	11.56	12.30	12.00
Lack of motivation/focus/concentration	6.85	6.88	6.92
I have been interrupted by noise made by others	8.51	5.23	6.55
Lack of contact/interaction with colleagues	5.58	1.38	2.99
I have had to share space and equipment	2.92	2.37	2.59
Distractions at home	3.76	1.47	2.35
Need to be at workplace for full role	0.58	3.05	2.13
Changes in how work organised because of Covid-19 restrictions	3.91	1.00	2.11
Tired, ill, other health issues	0.85	1.53	1.28
More work, longer hours	0.89	0.56	0.69
Furloughed	0.76	0.52	0.61
Different/new job	0.27	0.60	0.48
Maternity/paternity leave	0.00	0.76	0.48
Sample size N =	390	686	1076

Source: UKHLS Covid module, June wave
Note: Table shows proportions of stated reasons for productivity declines. Reasons only elicited for those who reported a decline compared to pre-Covid. Survey weights used. See text for more details.

Before examining these associations in detail, we document responses to a question asking for the main reason for productivity declines. This question was asked of anyone responding 'I get much less done' or '...a little less done'. The responses are tabulated in Table 5, and split by gender. A multitude of responses are given, indicating the varied reasons why productivity has declined. The most common reasons relate to childcare and to a lack of available work. Lack of work is evidence

Covid Economics 52, 15 October 2020: 118-151

of labour hoarding, or perhaps inefficient allocation of work across co-workers. While lack of work is reported with similar frequency across gender, the presence of children is cited as a reason by far more females than males. This latter result is consistent with widespread evidence discussed above finding that the bulk of childcare and homeschooling during lockdown was performed by females (see for example, Andrew et al., 2020). Beyond these causes the next most frequent response relates to lack of adequate equipment or software at home. Further down the list, the only reasons quoted by a non-negligible fraction of respondents are a lack of motivation, and noise distractions by others, presumably other than children. The former reason most directly indicates a causal effect from mental health declines. Reports of noise distractions further indicate the stressful situations under which some workers were required to perform their jobs.

In Table 6 we show the relationship between changes in productivity, working from home and mental well-being. Our measure of mental well-being is the individual change since wave 9 in standardized inverted Likert score. Accordingly, we are associating individual-level changes in well-being with reported changes in various factors. In the first column we regress the change in well-being on dummies for each of the productivity change indicators. Here, a report of 'I get about the same done' is the base category, with its effect on well-being captured by the coefficient on the constant. Relative to this base, those who report getting much less done also report substantially lower well-being. The coefficient of -0.54 standard deviations is large, and roughly in line with what is typically observed during a spell of unemployment. At the other end of the scale, those who report getting much more done report substantially higher well-being. In the second column we perform the same regression, but including controls for gender, age, degree-holding status and industry, with almost identical results.

In the third column we look at the relationship between changes in well-being and working-from-home status during lockdown. These regressions include all workers; those who never work from home are now the omitted category. The relationship between these variables has been explored in a similar way, and using these data, by Felstead and Reuschke (2020), who find that during the early part of the lockdown, workers, who worked at home sometimes, often, or always, all experienced a significant and similar level of decline in mental well-being. The evidence here indicates that by June, well-being has little noticeable relationship with location of work itself. This result pertains with or without basic controls.

As discussed, the strong association between change in productivity and mental well-being likely reflects causal relationships in both directions. In the final column we provide some preliminary evidence for an effect from productivity to mental well-being by instrumenting productivity changes using information available elsewhere in the survey. Here a variety of instruments could be considered. Given the proceeding discussion in this section, obvious candidates are industry or occupation of work.

CEPR PRESS

Covid Economics 52, 15 October 2020: 118-151

Table 6: Productivity Changes, Working from Home, and Mental Well-Being

	Change in well-being				
	OLS	OLS	OLS	OLS	IV
Prod. change (index)					0.24***
					(0.07)
Prod: much less done	-0.54***	-0.58***			
	(0.12)	(0.10)			
Get little less done	-0.25***	-0.24***			
	(0.07)	(0.07)			
Get little more done	0.02	0.05			
	(0.07)	(0.07)			
Get much more done	0.30***	0.30***			
	(0.08)	(0.08)			
Working from Home: Always			-0.03	0.01	
			(0.04)	(0.05)	
Often			-0.04	-0.01	
			(0.07)	(0.07)	
Sometimes			-0.06	-0.06	
			(0.09)	(0.09)	
Constant	-0.20***	-0.16	-0.23***	1.64***	-0.08
	(0.04)	(0.19)	(0.03)	(0.11)	(0.22)
Observations	3190	2957	6024	5513	2957
Controls		✓		✓	✓
Adjusted R^2	0.043	0.084	-0.000	0.024	0.081

Source: Wave 9 and Covid module of UKHLS

* $p < 0.10$,** $p < 0.05$, *** $p < 0.01$

Note: Standard errors in parentheses. In column 5, productivity change is instrumented with report of having ineffective equipment. In columns 2, 4 and 5, regressions include following controls: respondent's gender, age, education level (degree holding) and job industry. In column 3 and 4, we report relationship between mental wellbeing and WFH intensity using June wave of UKHLS Covid module.

Covid Economics 52, 15 October 2020: 118-151

However, it could be argued that industry or occupation affect well-being not only through job efficacy but also through other channels, such as differential social interaction, and differential exposure to Covid-related anxieties. For these reasons we favour an alternative approach. Here we use the reasons for productivity declines stated in Table 5. Specifically we instrument productivity changes with an indicator for whether the individual reports having inadequate equipment or software at home. Our maintained hypothesis is that lack of equipment only affects change in well-being through its effect on productivity. A possible criticism of this approach is that, given that the reasons for productivity changes are never elicited from those who report productivity increases, then the 'first-stage' regression is ensured by construction. Nevertheless we feel it is realistic to assume that those who experience equipment problems do find it detrimental on average. Turning to results, we find (but do not show) that those reporting inadequate equipment indeed suffer declines in well-being. Accordingly, and in terms of an IV regression, the fifth column of Table 6 shows that the effect of productivity changes on well-being is strong.

4 Conclusion

The Covid-19 pandemic has caused widespread disruption to working practices. The most noticeable change has been the vast increase in working from home. While some changes to working practices are probably temporary, many could very likely be persistent. Even after the pandemic ends, home working in particular is expected to be much more prevalent than previously.

In this paper we use the Covid-19 module from a household panel in the UK, which provides representative data on home workers' self-reported productivity towards the end of the lock-down period, in June 2020. In this survey, anyone working from home (WFH) at least some of the time was asked about changes in their productivity since before the pandemic period, at the beginning of the year. These data allow us to examine how productivity changes vary across both job and worker types.

We find that workers in industries and occupations that are less suitable for working from home report lower productivities than before the pandemic. Consistently with this, and with the literature, females and low earners also report lower productivity at home on average. For females, this lower productivity is not only due to the average characteristics of their jobs, but also because they are disproportionately effected by the presence of children. When examining workers based on changes in their WFH intensity, the evidence suggests that working from home itself has largely been beneficial, and has offset the other negative effects of the pandemic on productivity. Finally we produce evidence suggesting that difficulty in performing one's job causes lower mental well-being.

The evidence provided in this paper is relevant for policy in several ways. Most importantly it

CEPR PRESS

contributes to our understanding of the sector-specific impacts of the pandemic. This in turn helps inform policy-makers of the likely efficacy of targeted policies. It also informs quantitative analyses involving models of sector-specific supply shocks.

Covid Economics 52, 15 October 2020: 118-151

References

Adams-Prassl, A., Boneva, T., Golin, M., and Rauh, C. (2020a). Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys. *Journal of Public Economics*, 189.

Adams-Prassl, A., Boneva, T., Golin, M., and Rauh, C. (2020b). Work That Can Be Done from Home: Evidence on Variation within and across Occupations and Industries.

Alon, T., Doepke, M., Olmstead-Rumsey, J., and Tertilt, M. (2020). The Impact of COVID-19 on Gender Equality. *Covid Economics*, (4):62–85.

Andrew, A., Cattan, S., Costa Dias, M., Farquharson, C., Kraftman, L., Krutikova, S., Phimister, A., and Sevilla, A. (2020). How Are Mothers and Fathers Balancing Work and Family under Lockdown? Technical report, The Institute for Fiscal Studies.

Banks, J. and Xu, X. (2020). The Mental Health Effects of the First Two Months of Lockdown and Social Distancing during the Covid-19 Pandemic in the UK. *Covid Economics*, (28):91–118.

Baqaee, D. R. and Farhi, E. (2020). Supply and Demand in Disaggregated Keynesian Economies with an Application to the COVID-19 Crisis.

Bartik, A. W., Cullen, Z. B., Glaeser, E. L., and Stanton, C. T. (2020). What Jobs are Being Done at Home during the Covid-19 Crisis ? Evidence from Firm-Level Surveys.

Battiston, D., Blanes i Vidal, J., and Kirchmaier, T. (2017). Face-to-Face Communication in Organisations.

Bloom, B. N. (2020). How Working from Home Works out. Technical report.

Bloom, N., Liang, J., Roberts, J., and Ying, Z. J. (2015). Does Working from Home Work? Evidence from a Chinese Experiment. *The Quarterly Journal of Economics*, 130(1):165–218.

Blundell, R. and Etheridge, B. (2010). Consumption, Income and Earnings Inequality in Britain. *Review of Economic Dynamics*, 13(1):76–102.

Blundell, R., Green, D. A., and Jin, W. (2016). The UK Wage Premium Puzzle: How Did a Large Increase in University Graduates Leave the Education Premium Unchanged?

Bonadio, B., Huo, Z., Levchenko, A., and Pandalai-Nayar, N. (2020). Global Supply Chains in the Pandemic. *NBER Working Paper*.

Covid Economics 52, 15 October 2020: 118-151

Brodeur, A., Clark, A., Fleche, S., and Powdthavee, N. (2020). Assessing the Impact of the Coronavirus Lockdown on Unhappiness, Loneliness, and Boredom Using Google Trends.

Brooks, S. K., Webster, R., Smith, L. E., Woodland, L., Wessely, S., Greenberg, N., and Rubin, G. J. (2020). The Psychological Impact of Quarantine and How to Reduce It: Rapid Review of the Evidence. *The Lancet*, 395(10227):912–920.

Brynjolfsson, E., Horton, J. J., Ozimek, A., Rock, D., Sharma, G., Tuye, H.-Y., and Upwork, A. O. (2020). COVID-19 and Remote Work: An Early Look at US Data.

Chang, S. S., Stuckler, D., Yip, P., and Gunnell, D. (2013). Impact of 2008 Global Economic Crisis on Suicide: Time Trend Study in 54 Countries. *BMJ*, 347:1–15.

Dagher, R. K., Chen, J., and Thomas, S. B. (2015). Gender Differences in Mental Health Outcomes before, during, and after the Great Recession. *PLoS ONE*, 10(5):1–16.

Dingel, J. I. and Neiman, B. (2020). How Many Jobs Can Be Done at Home? *Journal of Public Economics*, 189.

Dutcher, G. (2012). The Effects of Telecommuting on Productivity: An Experimental Examination. The Role of Dull and Creative Tasks. *Journal of Economic Behavior and Organization*, 84(1):355–363.

Etheridge, B. and Spantig, L. (2020). The Gender Gap in Mental Well-Being during the COVID-19 Outbreak: Evidence from the UK. *Covid Economics*, (33):46–72.

Eurofound (2020). Living, Working and COVID-19 - First Findings. Technical report.

Felstead, A. and Reuschke, D. (2020). Homeworking in the UK: Before and during the 2020 Lockdown. Technical report.

Fetzer, T., Hensel, L., Hermle, J., and Roth, C. (2020). Coronavirus Perceptions and Economic Anxiety. *Review of Economics and Statistics*.

Greenberg, P. E., Kessler, R. C., Birnbaum, H. G., Leong, S. A., Lowe, S. W., Berglund, P. A., and Corey-Lisle, P. K. (2003). The Economic Burden of Depression in the United States: How Did It Change between 1990 and 2000? *Journal of Clinical Psychiatry*, 64(12):1465–1475.

Heathcote, J., Perri, F., and Violante, G. L. (2010). Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006. *Review of Economic Dynamics*, 13(1):15–51.

Covid Economics 52, 15 October 2020: 118-151

Holmes, E. A., O'Connor, R. C., Perry, V. H., Tracey, I., Wessely, S., Arseneault, L., Ballard, C., Christensen, H., Cohen Silver, R., Everall, I., Ford, T., John, A., Kabir, T., King, K., Madan, I., Michie, S., Przybylski, A. K., Shafran, R., Sweeney, A., Worthman, C. M., Yardley, L., Cowan, K., Cope, C., Hotopf, M., and Bullmore, E. (2020). Multidisciplinary Research Priorities for the COVID-19 Pandemic: A Call for Action for Mental Health Science. *The Lancet Psychiatry*, 7:547–560.

Institute for Social and Economic Research (2020). Understanding Society: COVID-19 Study. pages SN: 8644, 10.5255/UKDA–SN–8644–1.

Janke, K., Lee, K., Propper, C., Shields, K., and Shields, M. A. (2020). Macroeconomic Conditions and Health in Britain: Aggregation, Dynamics and Local Area Heterogeneity.

Knipe, D., Evans, H., Marchant, A., Gunnell, D., and John, A. (2020). Mapping Population Mental Health Concerns Related to COVID-19 and the Consequences of Physical Distancing: A Google Trends Analysis. *Wellcome Open Research*, 5.

Masayuki, M. (2020). Productivity of Working from Home during the COVID-19 Pandemic: Evidence from an Employee Survey. *Covid Economics*, (49):132–147.

Mongey, S., Pilossoph, L., and Weinberg, A. (2020). Which Workers Bear the Burden of Social Distancing Policies? *Covid Economics*, (12):69–86.

Office for National Statistics (2020). Coronavirus and Occupational Switching: January to June 2020. Technical report.

Ozimek, A. (2020). The Future of Remote Work.

Perri, F. and Steinberg, J. (2012). Inequality and Redistribution during the Great Recession.

Rubin, O., Nikolaeva, A., Nello-Deakin, S., and te Brommelstroet, M. (2020). What Can We Learn from the COVID-19 Pandemic about How People Experience Working from Home and Commuting? Technical report.

Tubadji, A., Boy, F., and Webber, D. J. (2020). Narrative Economics, Public Policy and Mental Health. *Covid Economics*, (20):109–131.

World Health Organization (2020). Mental Health and Psychosocial Considerations during COVID-19 Outbreak. Technical report.

Appendix

A Cross-walk between SOC 2000 and O*NET Occupation

Table A.1 shows the cross-walk this paper adopts to convert the Standard Occupational Classification (SOC) 2000 to the Occupational Information Network (O*NET) codes. Specifically, we assign each 3-digit SOC (sub-major occupation groups) into 2-digit O*NET codes (major occupation groups) by looking into the 4-digit SOC (sub-sub-major occupation groups) under each 3-digit SOC classification and matching them (4-digit SOC) with the most appropriate 2-digit O*NET category. Then, we assign each 3-digit SOC, based on the matching outcomes of 4-digit SOC to 2-digit O*NET code using an employment-weighted majority rule. Further, we also utilize industry information to split occupations under 3-digit SOC. Specifically, under SOC 922 '*Elementary Personal Services Occupations',* several food preparation related occupations are listed, such as '*Kitchen and catering assistants',* '*Waiters and Waitresses'*. These occupations belong to the industry related to food. Therefore, we move those respondents whose 3-digit SOC is 922 and industry related to Food into O*NET 35 '*Food Preparation and Serving Related Occupations'*.

Although in most cases the overwhelming majority of 4-digit SOC codes are assigned to the same 2-digit O*NET code, this is not always the case. As a result, some matches between SOC 2000 and O*NET codes are necessarily imprecise. For instance, SOC 231 '*Teaching Professionals'* is classified into O*NET 25 **'***Education, Training, and Library Occupations',* yet under it, SOC 2317 '*Registrars and senior administrators of educational establishments*' is more appropriate to be put into 2-digit O*NET 11 '*Management Occupations*', according to O*NET description. Due to the unavailability of 4-digit SOC information, we are unable to specifically subtract sub-sub-major occupation group SOC 2317 from sub-major occupation group SOC 231. Similarly, we cannot move SOC 5241 '*Electricians'* out of O*NET 49 '*Installation, Maintenance, and Repair Occupations'* and into O*NET 47 '*Construction and Extraction Occupations*'.

Figure A.1 plots occupation distributions of respondents from wave 9 and the Covid module of UK Household Longitudinal Study (UKHLS) and national employment statistics from 2019 US Bureau of Labor Statistics (BLS).

In the figure, white columns represent occupation percentages in UK-HLS and grey columns rep-

Covid Economics 52, 15 October 2020: 118-151

Table A.1: Cross-walk between 3-digit SOC 2000 to 2-digit O*NET Classification

3-digit SOC	SOC title	2-digit O*NET	O*NET title
111	Corporate managers and senior officials	11	Management
112	Production managers	11	Management
113	Functional managers	11	Management
114	Quality and customer care managers	11	Management
115	Financial institution and office managers	11	Management
116	Managers in distribution, storage and retailing	11	Management
117	Protective service officers	11	Management
118	Health and social services managers	11	Management
121	Managers in farming, horticulture, forestry and fishing	11	Management
122	Managers and proprietors in hospitality and leisure services	11	Management
123	Managers and proprietors in other service industries	11	Management
211	Science professionals	19	Life, Physical, and Social Science
212	Engineering professionals	17	Architecture and Engineering
213	Information and communication technology professionals	15	Computer and Mathematical
221	Health professionals	29	Healthcare Practitioners and Technical
231	Teaching professionals	25	Education, Training, and Library
232	Research professionals	19	Life, Physical, and Social Science
241	Legal professionals	23	Legal
242	Business and statistical professionals	13	Business and Financial Operations
243	Architects, town planners, surveyors	17	Architecture and Engineering
244	Public service professionals	21	Community and Social Service
245	Librarians and related professionals	25	Education, Training, and Library
311	Science and engineering technicians	17	Architecture and Engineering
312	Draughtspersons and building inspectors	17	Architecture and Engineering
313	IT service delivery occupations	15	Computer and Mathematical
321	Health associate professionals	29	Healthcare Practitioners and Technical
322	Therapists	29	Healthcare Practitioners and Technical
323	Social welfare associate professionals	21	Community and Social Service
331	Protective service occupations	33	Protective Service
341	Artistic and literary occupations	27	Arts, Design, Entertainment, Sports, and Media
342	Design associate professionals	27	Arts, Design, Entertainment, Sports, and Media
343	Media associate professionals	27	Arts, Design, Entertainment, Sports, and Media
344	Sports and fitness occupations	27	Arts, Design, Entertainment, Sports, and Media
351	Transport associate professionals	53	Transportation and Material Moving
352	Legal associate professionals	23	Legal
353	Business and finance associate professionals	13	Business and Financial Operations
354	Sales and related associate professionals	41	Sales and Related
355	Conservation associate professionals	45	Farming, Fishing, and Forestry
356	Public service and other associate professionals	21	Community and Social Service
411	Administrative occupations: Government and related	43	Office and Administrative Support
412	Administrative occupations: Finance	43	Office and Administrative Support
413	Administrative occupations: Records	43	Office and Administrative Support
414	Administrative occupations: Communications	43	Office and Administrative Support
415	Administrative occupations: General	43	Office and Administrative Support
421	Secretarial and related occupations	43	Office and Administrative Support
511	Agricultural trades	45	Farming, Fishing, and Forestry
521	Metal forming, welding and related trades	47	Construction and Extraction
522	Metal machining, fitting and instrument making trades	51	Production
523	Vehicle trades	49	Installation, Maintenance, and Repair
524	Electrical trades	49	Installation, Maintenance, and Repair
531	Construction trades	47	Construction and Extraction
532	Building trades	47	Construction and Extraction
541	Textiles and garments trades	51	Production
542	Printing trades	51	Production

Covid Economics 52, 15 October 2020: 118-151

Covid Economics 52, 15 October 2020: 118-151

Table A.1 (Continue): Cross-walk between 3-digit SOC 2000 to 2-digit O*NET Classification

3-digit SOC	SOC title	2-digit O*NET	O*NET title
543*	Food preparation trades	35	Food Preparation and Serving Related
549	Skilled trades	51	Production
611	Healthcare and related personal services	31	Healthcare Support
612	Childcare and related personal services	39	Personal Care and Service
613	Animal care services	39	Personal Care and Service
621	Leisure and travel service occupations	39	Personal Care and Service
622	Hairdressers and related occupations	39	Personal Care and Service
623	Housekeeping occupations	37	Building and Grounds Cleaning and Maintenance
629	Personal services occupations N.E.C.	39	Personal Care and Service
711	Sales assistants and retail cashiers	41	Sales and Related
712	Sales related occupations	41	Sales and Related
721	Customer service occupations	43	Office and Administrative Support
811	Process operatives	51	Production
812	Plant and machine operatives	51	Production
813	Assemblers and routine operatives	51	Production
814	Construction operatives	47	Construction and Extraction
821	Transport drivers and operatives	53	Transportation and Material Moving
822	Mobile Machine Drivers And Operatives	53	Transportation and Material Moving
911	Elementary Agricultural Occupations	45	Farming, Fishing, and Forestry
912	Elementary construction occupations	47	Construction and Extraction
913	Elementary process plant occupations	51	Production
914	Elementary goods storage occupations	53	Transportation and Material Moving
921	Elementary administration occupations	43	Office and Administrative Support
922	Elementary personal services occupations	39	Personal Care and Service
923	Elementary cleaning occupations	37	Building and Grounds Cleaning and Maintenance
924	Elementary security occupations	33	Protective Service
925	Elementary sales occupations	41	Sales and Related

Note: Part of occupation 922 is allocated to O*NET occupation 35 Food Preparation and Serving Related. See text for more details.

Figure A.1: Occupation Percentage Distributions, UK-Household Longitudinal Study (HLS) and US-Bureau of Labor Statistics (BLS)

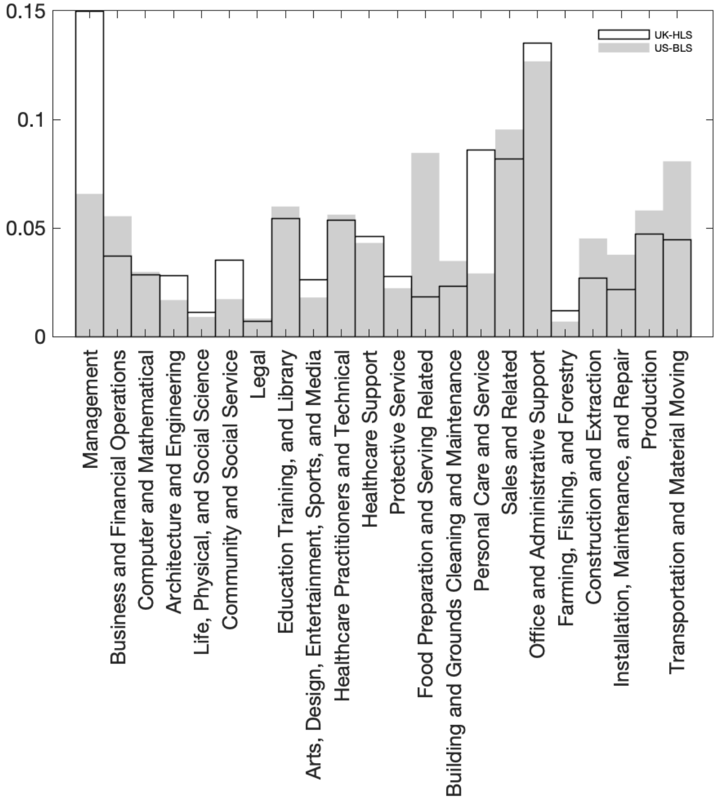

Source: Wave 9 and Covid module of UKHLS and BLS 2019 statistics

Covid Economics 52, 15 October 2020: 118-151

CEPR PRESS

Covid Economics 52, 15 October 2020: 118-151

resent occupation percentages in US-BLS. The correlation coefficient between both is around 0.7. The occupation categories showing largest differences are Management and Food Preparation and Serving Related. The sign of these differences is, at least, very likely genuine. The UK is reported to be particularly intensive in managers (Blundell et al., 2016). Similarly, the US is more intensive in Food Serving (waitering). If we exclude these occupations, the correlation coefficient between UK and US occupation percentage rises to around 0.8.

B Additional Information and Results

B.1 Productivity Data from the ONS

We also utilize the productivity statistics reported by Office for National Statistics (ONS) in each industry in UK. The productivity measures cover from 1997 Q2 to 2020 Q2 for UK main industries. Three seasonally adjusted statistics related to industry-level productivity are reported: gross value added (GVA), hours worked and output per hour. Both GVA and output per hour are measured by 2016 GBP. The relationship between these three statistics is: GVA equals the product of hours worked and output per hour. Further, we derive the industry-level productivity changes by calculating the difference of GVA between 2020 Q2 and 2019 Q4 for each industry. Note that, since for Manufacture industry, 13 sub-industry statistics are reported separately, e.g. Manufacture of food products, beverages and tobacco and Manufacture of textiles, wearing apparel and leather products, we obtain Manufacture-level GVA by aggregating sub-industries through calculating the product of hours worked and output per hour for all 13 individual sub-industries and then summing the 13 products up. The industry-level productivity change for Manufacture is derived by calculating the difference between the derived 2020 Q2 GVA and the 2019 Q4 GVA.

Moreover, in reporting, ONS combines three industries, Public Administration and Defense, Education and Human Health and Social Work Activities into one category and also combines the other three industries, Other Service Activities, Activities of Households as Employers and Activities of Extraterritorial Organizations and Bodies, into one category. For consistency, when plotting ONS measures of productivity change against our productivity change measures, we combine our statistics in the same way as ONS.

B.2 Additional Tables Mentioned in the Text

Covid Economics 52, 15 October 2020: 118-151

Table B.1: Proportions of Working from Home, by Occupation

	At least some of the time			Always	
	Jan/Feb	April	April to June	April	April to June
Management	0.49	0.66	-0.00	0.46	-0.03
Business and Financial Operations	0.57	0.92	-0.02	0.79	-0.03
Computer and Mathematical	0.60	0.89	0.01	0.76	-0.07*
Architecture and Engineering	0.34	0.68	-0.03*	0.50	-0.05*
Life, Physical, and Social Science	0.33	0.71	-0.09	0.54	-0.08
Community and Social Service	0.52	0.79	-0.07***	0.56	-0.07*
Legal	0.47	0.84	0.02	0.79	-0.01
Education, Training, and Library	0.49	0.89	-0.02	0.55	-0.15***
Arts, Design, Entertainment	0.64	0.78	-0.07	0.57	-0.05
Healthcare Practitioners and Technical	0.24	0.38	-0.01	0.11	0.01
Healthcare Support	0.11	0.16	0.02	0.07	0.01
Protective Service	0.12	0.22	-0.01	0.08	-0.01
Food Preparation and Serving Related	0.05	0.05	0.04*	0.01	0.01
Building Cleaning and Maintenance	0.11	0.09	-0.04	0.01	-0.01
Personal Care and Service	0.18	0.33	0.04	0.16	-0.01
Sales and Related	0.17	0.25	-0.02	0.18	-0.02
Office and Administrative Support	0.24	0.55	-0.01	0.40	0.00
Farming, Fishing, and Forestry	0.16	0.23	-0.08	0.16	-0.05
Construction and Extraction	0.11	0.18	-0.07	0.06	-0.04
Installation, Maintenance, and Repair	0.13	0.34	-0.14	0.19	-0.09
Production	0.14	0.21	-0.00	0.12	-0.01
Transportation and Material Moving	0.07	0.10	-0.01	0.04	0.01
Observations	6010	6743	5070	6743	5070
Adjusted R^2	0.402	0.622	0.008	0.475	0.016

Source: Wave 9 and Covid module of UKHLS

$^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

Note: This table reports proportions of respondents who WFH for at least some time and always, respectively for each occupation in the United Kingdom in 2020. The classification of occupations is converted from the UK Standard Occupational Classification (SOC) system to US O*NET system. The first two columns illustrate the proportions who spent at least some of the time WFH in January/February, and in April, respectively. The third column reports changes in proportions of WFH at least for some time from April to June. The last two columns demonstrate the proportion of always WFH in April and change in the proportion of always WFH from April to June, respectively.

Table B.2: Productivity Changes during Working from Home: By Characteristics, Ordered Probit vs OLS

	Probit	OLS	Probit	OLS	Probit	OLS	Probit	OLS	Probit	OLS	Probit	OLS
Earnings: Bottom	-0.27***	-0.29***										
	(0.08)	(0.09)										
Top	0.11*	0.12*										
	(0.06)	(0.07)										
Age: 16 to 29			0.01	0.01								
			(0.09)	(0.10)								
over 50			0.09	0.09								
			(0.05)	(0.06)								
Education: No Degree					-0.03	-0.03						
					(0.05)	(0.06)						
Gender: Female							-0.15***	-0.16***				
							(0.05)	(0.06)				
Children: None									0.08	0.08		
									(0.08)	(0.08)		
Employment: Self-employed											-0.34***	-0.36***
											(0.09)	(0.09)
Constant		2.96***		2.94***		2.99***		3.06***		2.90***		3.38***
		(0.05)		(0.04)		(0.04)		(0.04)		(0.08)		(0.10)
Observations	2912	2912	3254	3254	3254	3254	3034	3034	3254	3254	3067	3067

Source: Wave 9 and Covid module of UKHLS

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Standard errors in parentheses. This table reports estimation results for productivity change during WFH by respondent's characteristics, i.e. earnings, age, education, gender, children and employment type, using ordered probit and OLS, respectively. Omitted groups are base groups.

CEPR PRESS

152

Covid Economics Issue 52, 15 October 2020

The heterogeneous effects of COVID-19 on labor market flows: Evidence from administrative data[1]

Alessandra Casarico[2] and Salvatore Lattanzio[3]

Date submitted: 5 October 2020; Date accepted: 6 October 2020

Covid Economics 52, 15 October 2020: 152-174

We investigate the short term labor market response to the pandemic in Italy and provide a first evaluation of the policies put in place to shield workers from the disruption of economic activity. Using administrative data on a sample of contracts active in the first quarter of 2020, we show that, before the pandemic, workers employed in non-essential activities were in majority men, younger than 35 years old, located in the North of the country and with lower levels of education. When looking at the change in hirings and separations and decomposing it by age, gender, region, type of contract (open-ended or temporary), education level, and sector (essential vs non-essential activities), we find that from the nineth week of the year, there was a pronounced drop in hirings and terminations. On the contrary, firings and quits spiked right after the nineth week, and then dropped significantly, reflecting the effects of the firing freeze and the easing of access to STW compensation schemes. We further explore separations by examining which factors predict the probability of job loss. We find that those workers that were already suffering the consequences of the previous recession (young, temporary, low-skill workers) are those at higher risk of losing their job because of COVID-19. Gender, instead, is a non-significant predictor of job loss in the aggregate.

1 We would like to thank Bruno Anastasia, Francesco D'Amuri and Andrea Garnero for useful comments and discussion.
2 Bocconi University, Cesifo, Dondena.
3 University of Cambridge.

CEPR PRESS

1 Introduction

The COVID-19 pandemic is having dramatic consequences on society. In order to contain the spread of the virus, many governments around the world adopted unprecedented interventions that in most cases resulted in lockdowns of entire regions or countries. The suspension of economic activities had severe repercussions on employment and earnings of individuals and on profits of firms. As a consequence, global GDP growth is projected at −4.9 percent (IMF, 2020), with considerable heterogeneity between advanced (−8 percent) and emerging economies (−3 percent). Governments responded to the economic downturn with encompassing packages of fiscal measures, ranging from transfers, loans, postponements of tax dues, to facilitating liquidity and access to credit for firms. Preventing or reducing the disruption of the labor market was among the main goals of government intervention, and the specific instruments adopted varied across countries, also in light of pre-existing labor market institutions. The implemented policy measures and pre-existing labor market conditions and institutions mediate the impact of the pandemic on jobs. For example, Adams-Prassl et al. (2020) compare the United Kingdom, the United States and Germany and show that the job losses were higher in the first two countries, which are characterized by more flexible labor markets.

In this paper, we investigate the labor market response to the pandemic in Italy and provide a first evaluation of the policies put in place to shield workers from the disruption of economic activity associated with the pandemic. Italy was the first country in Europe to be hit by COVID-19 and the first to implement a national lockdown, which was then adopted in most European countries. The lockdown was shortly after followed by two further policy measures relevant for labor market dynamics: a firing freeze and an ease of the requirements to access short-time work (STW) compensation schemes. While the former is unique to Italy for its breadth, the latter is common to most European countries (see OECD, 2020, for details on government policy responses across OECD countries).

Using administrative data on a sample of contracts active in the first quarter of 2020, we look at the ex-ante characteristics of the labor market before the start of the pandemic and at its ex-post short-run impact on the labor market, analyzing how government policies mediated the effects of the spread of the virus. First, we provide descriptive evidence on the personal and job characteristics of workers employed in essential and non-essential activities, as of January 2020. In the wake of the pandemic, the government decided to shut down entire sectors, which were deemed as non-essential. They were mainly concentrated in services, such as restaurants, bars, hotels, and some categories of wholesale and retail shops, in line with government decisions in other countries. We show that workers employed in non-essential

Covid Economics 52, 15 October 2020: 152-174

activities were in majority men, younger than 35 years old, located in the North of the country and with lower levels of education. Second, we analyze the change in hirings and separations – distinguishing between firings, terminations and quits – in each week of the first quarter of 2020 relative to the average of 2017-19. For each labor market flow, we provide graphical evidence of the weekly change and its decomposition into subgroups based on age, gender, region, type of contract (open-ended or temporary), education level, and sector (essential vs non-essential activities). The descriptive evidence shows that, before the pandemic, hirings were slightly above their average in the previous three years, similarly to terminations and quits, whereas firings were in line with their levels in the past. When COVID-19 spread quickly around the country, starting from the nineth week of the year, there was a pronounced drop in hirings and terminations. On the contrary, firings and quits spiked right after the nineth week, and then dropped significantly. The evolution of firings reflects the policy introduced on 17 March, that explicitly forbids firms from firing workers and, at the same time, eases the requirements to have access to STW compensation schemes. Absent the policy, firings were rising with respect to the past. Moreover, the firing freeze may also have contributed to the decreasing dynamics of hirings, as the higher employment protection for workers may have decreased turnover. Third, we further explore separations by examining which factors predict the probability of job loss. We find that a younger age, being on a temporary or part-time contract, working in the Centre or the South relative to the North, having less than upper secondary education are all significant predictors of the separation probability: in other words, those workers that were already suffering the consequences of the previous recession (young, temporary, low-skill workers) are those at higher risk of losing their job because of COVID-19. In this light, the firing freeze guaranteed protection to the most vulnerable groups in the labor market. We also explore the difference in the job loss probability between essential and non-essential activities: we find that the same job and personal characteristics of workers are associated with the separation margin, but in non-essential activities coefficients are generally larger in magnitude. Finally, we explore differences by gender, which in our data is a non-significant predictor of job loss, contrarily to the evidence shown for example in Adams-Prassl et al. (2020). Given the higher concentration of women in temporary contracts and part-time positions, coupled with the nation-wide school closures, one could expect a harsher impact of the crisis on women, as highlighted by Alon et al. (2020). We do not find a higher separation probability for women relative to men in the aggregate, but we do find a significantly higher separation margin for female workers with upper secondary education and female domestic workers (while the opposite holds true for female farming workers). The null effect of gender in the aggregate may be due to different factors. First, we are focusing on the short-term effect of the

Covid Economics 52, 15 October 2020: 152-174

pandemic recession, but differences by gender might take more time to materialize. Second, the policies put in place by the government have been effective in protecting more vulnerable categories and contracts, among which women are more represented. Third, differently from the evidence provided by Blundell et al. (2020) for the UK, we find that female presence in non-essential activities is lower relative to men, hence women's jobs may have had better chances to survive because less likely to be in shut-down sectors. Finally, if women are employed in jobs whose tasks can be more easily performed from home, their employment is relatively more protected. While we can only discuss the extensive margin of adjustment – whether a worker separates from her job or not – clearly, the adjustment may happen also on the intensive margin, if women had to adjust their work hours in response to the pandemic. This is an important element we cannot directly address with the data at hand.

Our analysis contributes to the recent and growing literature on the effects of the pandemic recession on economic activity (e.g. Carvalho et al., 2020; Chetty et al., 2020; Baker et al., 2020) and, specifically, on the labor market and the policy responses put in place by governments. Evidence using real-time survey data (Bick and Blandin, 2020; Adams-Prassl et al., 2020; von Gaudecker et al., 2020), administrative data (Cajner et al., 2020) and a combination of both (Forsythe et al., 2020) highlights the severe and unequal consequences of the pandemic recession on the labor market. A strand of this literature specifically focuses on how different categories of workers were affected by the pandemic (Blundell et al., 2020; Crossley et al., 2020), with particular focus on age (Belot et al., 2020) and gender (Alon et al., 2020; Hupkau and Petrongolo, 2020; Farré et al., 2020). We provide new evidence based on detailed administrative data on a sample of active, new and terminated contracts, coming from the *Comunicazioni Obbligatorie*, i.e. the compulsory information firms need to provide on their workforce. These data are highly reliable and less subject to measurement errors with respect to survey data. We can explore many dimensions of heterogeneity and provide an exhaustive picture of the unequal impact of COVID-19. We also assess the short run impact of a government policy that explicitly forbids dismissals and extends the generosity of STW compensation schemes. We show they were successful in taming firings – as expected –, but may also have reduced hirings. This lays the groundwork for a medium term assessment of their impact on labor market dynamics. Finally, by showing how workers on different types of contracts and different degrees of employment protection are affected by the pandemic recession, we contribute to the literature that analyzes the margins of adjustment in the labor market in the presence of negative shocks (Izquierdo et al., 2017; Garin and Silvério, 2019; Adamopoulou et al., 2020).

The remainder of the paper is organized as follows. Section 2 describes the data and gives details about the evolution of the pandemic in Italy and the policy response by the

CEPR PRESS

government. Section 3 shows the distribution of workers in essential and non-essential activities before the pandemic. Section 4 analyzes the changes in hirings and separations between 2020 and previous years, whereas section 5 focuses on a formal analysis of the determinants of the job loss probability. Finally, section 6 concludes.

2 Data and Institutional Context

2.1 Data and Descriptive Statistics

We use data from a random sample of mandatory notifications (*Campione Integrato delle Comunicazioni Obbligatorie*, CICO) that firms submit to relevant public agencies in Italy and to the Ministry of Labor and Social Policy. The data collects information on a sample of contracts activated and terminated between 2009 and the first quarter of 2020 for public- and private-sector workers, farming and domestic workers.[1] For each contract we have information on the exact start date and, if the contract is terminated, on the end date and the reason for its termination (mainly, firings, termination of temporary contracts, voluntary quits).[2] Furthermore, we have information on the type of contract (open-ended or temporary, full-time or part-time), detailed occupational and sectoral codes (6-digit Isco and Ateco 2007, respectively) and individual characteristics of workers, such as gender, the year of birth, the region of domicile and work, and the education level. Table 1 reports descriptive statistics on the contracts – and on the individual characteristics of the workers holding them – and compares them with the population of workers from the national statistical institute (Istat). Our data over-samples contracts held by female workers, workers in the age group 15-34, workers with high-school diploma and under-samples contracts of workers on open-ended and full-time positions. The bottom part of the table, column (1), reports the sample size of CICO, distinguishing total contracts, total workers (as workers can hold multiple contracts), employment contracts and employees (i.e. the subset of workers holding employment contracts, therefore excluding domestic and farming workers and collaboration contracts) by the end of the sample period.[3] Column (2) reports the number of workers/employees from Istat.

[1]The sampling strategy is based on the day of birth: workers born on the 1st, 9th, 10th and 11th day of each month and year in the full administrative records are included in the sample. CICO contains information on contracts that have been activated, transformed or ended starting from 2009. Hence, the data contains information on new contracts from 2009 and on contracts that have been established before 2009 but that were either terminated or transformed in subsequent years. Therefore, the data do not contain information on contracts that have been stipulated before 2009 and that have not been modified since then.

[2]We exclude from the sample terminations due to retirement, death or modification of the end date of the contract, as there is no further information on whether the end date is anticipated or postponed.

[3]Throughout the rest of the paper the unit of observation will be the single contract. Hence, it might be that workers holding multiple contracts appear more than once in our data. This choice does not affect the

Covid Economics 52, 15 October 2020: 152-174

Table 1: Descriptive statistics

	(1) CICO	(2) Istat
Female	0.46	0.42
Age 15-34	0.31	0.22
Age 35-54	0.53	0.56
Age 55+	0.16	0.22
North	0.53	0.52
Centre	0.22	0.21
South	0.25	0.26
High-school diploma	0.84	0.76
University degree	0.16	0.24
Open-ended contract	0.62	0.83
Full-time contract	0.62	0.79
Industry	0.23	0.26
Total contracts	2,314,429	-
Total workers	1,951,450	23,383,281
Total employment contracts	1,352,872	-
Total employees	1,164,297	18,096,880

Notes. The table reports the share of contracts in each group from the sample of *Comunicazioni Obbligatorie* (CICO) and the share of workers from official statistics provided by the National Statistical Institute (Istat). The last four rows of the table report the total number of contracts and workers and the total number of employment contracts and employees (as a worker/employee may have multiple contracts) present in CICO and the total number of workers and employees in Istat.

Overall, our sample represents approximately 8.3% of the population of workers in Italy and 6.4% of employees. The fact that younger and female workers are over-represented, whereas more stable contractual arrangements – such as, open-ended and full-time contracts – are under-represented comes as no surprise given the sample selection described above. The data over-samples contracts stipulated in the last decade, which capture the first contract of new workers, who are therefore more likely to be young and on temporary/part-time positions. Women may be over-represented in light of the progression in female labor force participation in recent years. However, although not representative of the population of workers at a given point in time, the data allows us to compare flows between different years (e.g. the change in hirings or separations over time) and also compare the distribution of workers in the subgroups of essential and non-essential activities, as one can believe the sampling bias

composition of our sample: if we use only one observation per worker and replicate Table 1 we get almost identical sample shares.

Covid Economics 52, 15 October 2020: 152-174

Covid Economics 52, 15 October 2020: 152-174

would not be different in the two subgroups.

2.2 COVID-19 in Italy and Public Policy

The first cases of COVID-19 in Italy date back to 31 January 2020, but the disease began to spread exponentially in the second half of February. At the beginning, the virus spread predominantly in Northern regions and the first COVID-related death was registered in Lombardy on 21 February. Following the diffusion of the virus in the North, two "red zones" were implemented, involving 11 municipalities in Lombardy and Veneto. At the same time, many Northern regions opted to close schools, a measure that extended to the whole nation on 4 March. On 10 March the whole country went into lockdown. The decree establishing the nationwide lockdown also specified the activities that were deemed as essential and could continue to operate and those that were classified as non-essential and were forced to shut down: the former mainly include agriculture, some manufacturing, energy and water supply, transports and logistics, ICT, banking and insurance, professional and scientific activities, public administration, education, healthcare and some service activities; shutdown sectors include most of manufacturing activities, wholesale and retail trade, hotels, restaurants and bars, entertainment and sport activities. In light of these closures the government adopted on 17 March a Decree Law that considerably increased worker's employment protection. Two main labor market policies were adopted:

(1) A special COVID-related STW compensation scheme of the duration of 9 weeks, that could apply retroactively starting from 23 February. This measure aimed at preserving employment relationships and allowing firms to cut labor costs during the lockdown period, by reducing hours of work thanks to a wage subsidy granted by the government. The measure extended the regular STW by allowing firms with less than 15 employees and firms that were already using the extra-ordinary STW (one of the sub-species of STW granted by the Italian employment protection legislation) to use it. Moreover, firms using the COVID-related STW could renew temporary contracts, waiving to the norms of the standard regulation.

(2) A firing freeze that halted firings for 60 days from 17 March and that could be applied retroactively to pending firings (i.e. those that were yet to be validated) from 23 February.

In the rest of the paper we will highlight how our results may be affected by the implementation of such policies.

Covid Economics 52, 15 October 2020: 152-174

3 Before the Pandemic: the Distribution of Workers in Essential and Non-Essential Activities

Using data from CICO up to January 2020, we show the distribution of workers in essential and non-essential activities (i.e. between open and shutdown sectors) at the onset of the pandemic. Figure 1, panels (a)-(d), shows the distribution of workers by gender, age, region of work and education level. Panel (a) shows that women are over-represented in essential activities (67.1%) relative to men (57.8%): this result is in contrast with the evidence provided, for example, by Blundell et al. (2020) for the UK, where more women than men were employed in shutdown sectors before the pandemic. This may be explained by low female labor force participation and by the strong positive selection of women in the Italian labor market.[4]

Panel (b) shows the distribution by age, distinguishing workers in age groups 15-34, 35-54 and 55 or older. The figure shows that, while young workers are almost equally distributed between essential and non-essential activities, middle-aged and older workers are more present in essential activities. Hence, the closure of non-essential sectors has a stronger impact on young workers, 48.2% of whom are employed in shutdown sectors.

Panel (c) reports the distribution by region of work. Differences between the North, the Centre and the South are small and, if anything, more workers are employed in shutdown sectors in the North, relative to the rest of the country. This may seem counter-intuitive, considering that tourism and connected services are some of the strengths of Southern Italy. This distribution may also be correlated with the presence of the informal economy, which is higher in the South, as documented, for example, in Boeri et al. (2019), and particularly relevant for workers in accommodation, tourism and restaurants—sectors belonging to non-essential activities.

Panel (d) shows the distribution by education level. While 41.6% and 41.7% of workers with lower and upper secondary education are in shutdown sectors, only 18.5% of individuals with university degree work in non-essential activities, suggesting a disproportionate impact of the pandemic on workers with lower levels of education.

This analysis takes a snapshot of the Italian labor market at the onset of the pandemic. We now turn to the inspection of the impact of the crisis on hirings and separations in the first quarter of 2020.

[4]This result is not driven by the fact that our data over-samples women relative to the population. Aggregate data from the National Statistical Institute (Istat, 2020), based on the labor force survey, confirm that women are over-represented in essential activities (72.8%). If anything, our sample share is a lower bound to the population share.

Covid Economics 52, 15 October 2020: 152-174

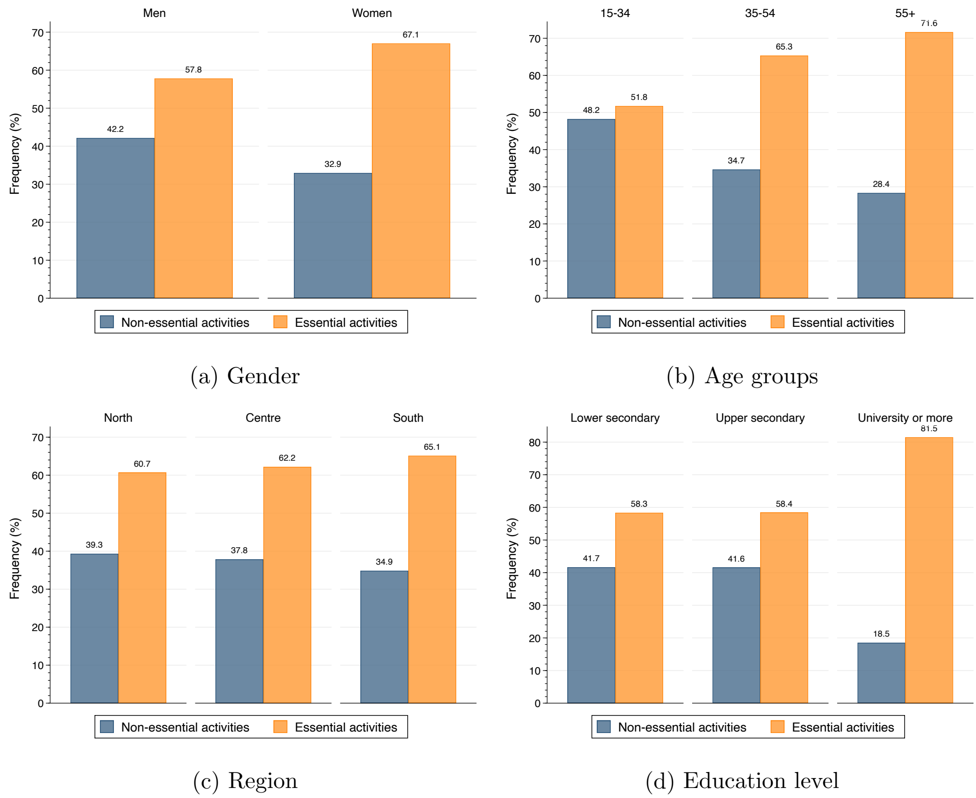

(a) Gender

(b) Age groups

(c) Region

(d) Education level

Figure 1: Distribution of workers in essential and non-essential activities as of January 2020

4 After the Pandemic: the Short-Run Impact of the Recession on Hirings and Separations

In this section, we analyze the dynamics of hirings (H) and separations – distinguished in firings (F), terminations (T) and quits (Q) – in the first quarter of 2020. Specifically, we compare the weekly change in hirings and separations between 2020 and the average of 2017-19, decomposing such change into the contribution of various subgroups based on the following categories G: age, gender, region of work, type of contract (permanent or temporary), education level and sector (essential or non-essential). In other words, for each week t we compute the percent change in $Y_t = \{H_t, F_t, T_t, Q_t\}$ as:

$$\triangle Y_t = \frac{Y_{t,2020} - \overline{Y}_{t,2017-19}}{\overline{Y}_{t,2017-19}}.$$

Covid Economics 52, 15 October 2020: 152-174

We present the weekly change $\triangle Y_t$ graphically, together with its decomposition into subgroups $g \in G$, $\triangle Y_t^g$:

$$\triangle Y_t = \sum_{g \in G} \triangle Y_t^g = \frac{\sum_{g \in G} Y_{t,2020}^g - \overline{Y}_{t,2017-19}^g}{\overline{Y}_{t,2017-19}}.$$

Hirings Figure 2, panels (a)-(f), shows the weekly change in hirings in 2020 relative to 2017-19. In each graph, the black line is the total change. The figures show that 2020 had more hirings than previous years in the first 8 weeks (that is, until the week ending on 25 February[5]). On average, in the first 8 weeks of 2020, weekly hirings have been 10.6% higher than previous years. Starting from week 9, though, weekly hirings experience a sharp drop, which becomes even worse from the 12th week of the year. The first drop in hirings is attributable to the nationwide lockdown and the closure of activities which effectively froze the labor market. The second drop in hirings happens after the decision to freeze layoffs. Although it is impossible to separately measure the impact of the lockdown and that of the firing freeze, the evidence is in line with the latter policy having a negative effect on hirings.[6]

Panel (a) decomposes the total weekly change in hirings into the contribution of the age groups 15-34, 35-54 and above 55. It is clear that most of the decrease in hirings comes from a decline in new hires of younger and middle-aged workers. For example, in week 13 the total weekly change in hirings amounts to −67%, −31% due to the drop in hirings of workers of age 15-34, −30% due to the drop for workers in the age group 35-54 and only −6% attributable to older workers. When looking at differences by gender, in panel (b), we do not see significant differences between men and women and, if anything, the drop in hirings was slightly more pronounced for male workers: on average, the drop in weekly hirings after week 8 amounts to −37%, −20% for men and −17% for women. Panel (c) shows the decomposition based on the region of work. In weeks 9 and 10, the reduction in hirings was stronger in the North, as it went into partial lockdown before the rest of the country. Starting from week 12, when the lockdown extended to the whole country, the contribution of the South to the drop in hirings increased substantially. On average, however, half of the drop in hirings is concentrated in the North (−18%). Remarkably, the drop in hirings is almost entirely concentrated among temporary contracts (panel d) and workers with low levels of education (panel e): workers on open-ended contracts and with a university degree contribute only for −1.6% and −3.3%, respectively, to the weekly change of −37% on average for weeks 9-13. Finally, panel (f) shows that workers employed in non-essential activities contributed slightly

[5]The first COVID-related death was notified in that week, on 21 February.

[6]See, e.g., Kugler and Pica (2008) for evidence on the impact of increased employment protection legislation on worker flows.

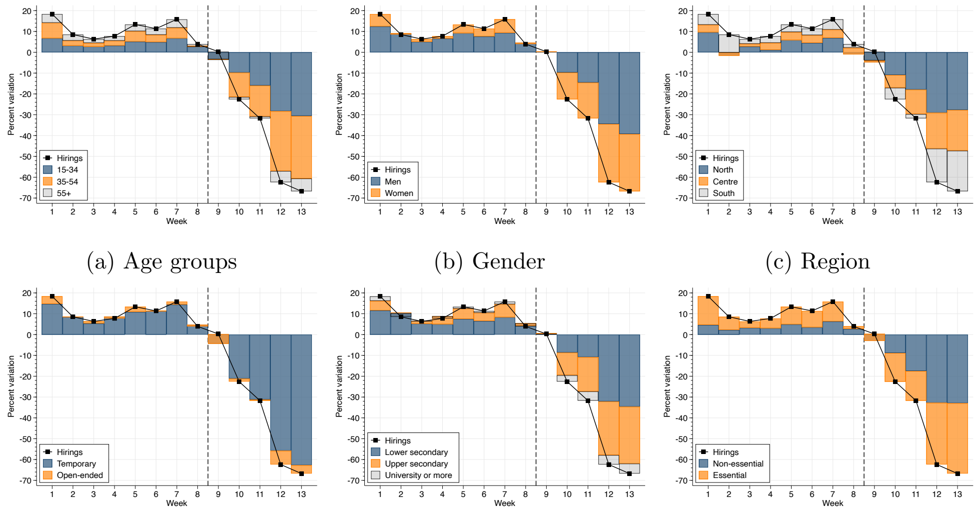

(a) Age groups (b) Gender (c) Region

(d) Type of contract (e) Education level (f) Sectors

Figure 2: Weekly change in hirings between first quarter of 2020 relative to average 2017-2019

more to the drop in hirings (−19%) relative to essential activities (−18%).

Firings Figure 3 reports the weekly change in firings and its decomposition into different subgroups. Until the beginning of the pandemic recession, weekly firings were in line with those registered in the past (on average −1.8% in the first 8 weeks). After the onset of the pandemic, we observe a sharp increase in firings, which was particularly evident in weeks 10 and 11, when firings were 70% and 33% higher than their level in previous years, respectively. In week 12, the firing freeze together with special STW compensation scheme came into force and we observe a sharp drop in firings, as expected. Overall, this evidence suggests that, absent the policy, firms would have resorted to firings to cut labor costs, although it is difficult to separate the impact of the firing freeze from that of STW. It is unclear, however, whether the firings we observe in the data in weeks 10 and 11 have been validated or not: in fact, the firing freeze, although introduced on 17 March, had retroactive effect until 23 February. Hence, in principle, those workers could have been reinstated. This does not change the main conclusion that we draw from this analysis, i.e. that the firing freeze effectively stopped firms from laying off workers, that traditional instruments (e.g. the regular STW scheme) would have not protected.

Which categories of workers were being fired and which benefited the most from the firing freeze? Decomposition results are similar to those outlined for hirings. Panel (a) shows that

Covid Economics 52, 15 October 2020: 152-174

Covid Economics 52, 15 October 2020: 152-174

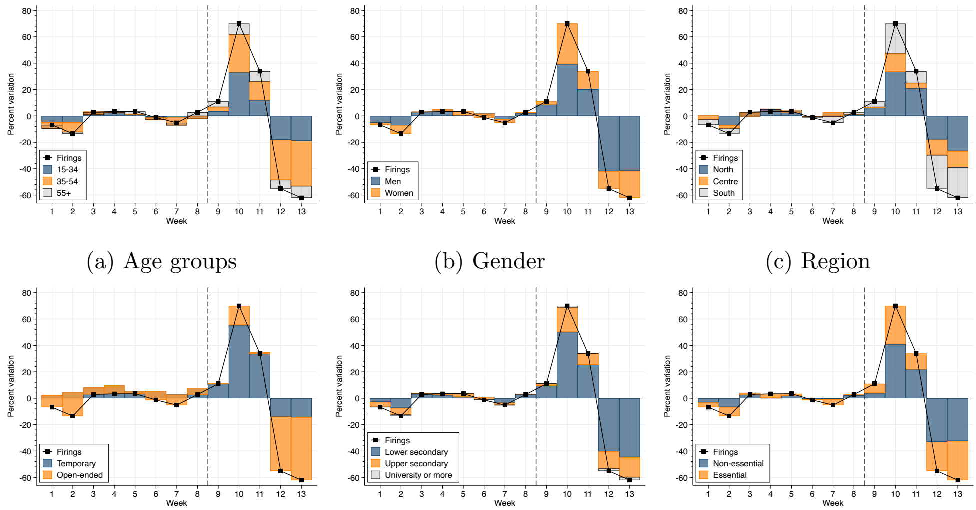

(a) Age groups (b) Gender (c) Region

(d) Type of contract (e) Education level (f) Sectors

Figure 3: Weekly change in firings between first quarter of 2020 relative to average 2017-2019

workers being fired in weeks 10 and 11 are mainly young and middle aged workers, which were also the most protected categories in weeks 12-13. Panel (b) displays differences by gender. The rise in firings hit men more: the average rise in firings of 52% was due to a rise of male firings of 30% and of female hirings of 22%. After the freeze, men benefited more from the protection of the policy and the reduction in their firings constitute most of the decline in weeks 12 and 13 (-42% over an average total change of -59%). When we look at geographical differences, we find that the initial rise in firings was concentrated in the North (half of the total change), but a significant portion of firings happened in the South, too, though the consequences of the pandemic hit it at least two weeks later than the North. This may be due to anticipation effects, disruptions of supply chains and trade relationship with the North or lower travels to the South, which may have impacted the accommodation sector. Panel (d) shows that temporary contracts were bearing most of the burden of firings in weeks 10-11, but when the firing freeze takes effect workers on open-ended contracts benefited more. Panel (e) shows that workers with low levels of education were being fired and then shielded by the firing freeze. Finally, panel (f) suggests that the majority of the surge in firings is concentrated in shutdown non-essential activities (-31% versus -21% in essential activities on average in weeks 10 and 11). The reduction in firings was also higher in non-essential activities after the imposition of the firing freeze.

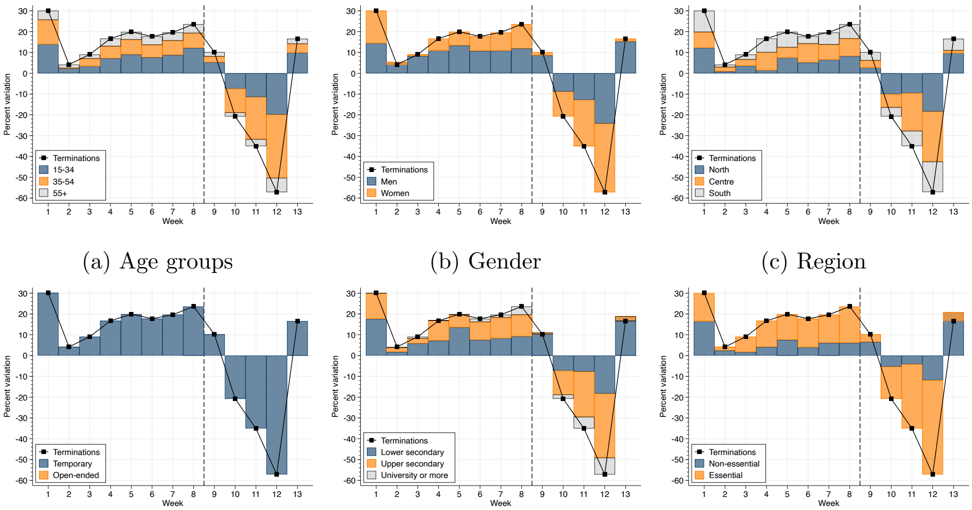

(a) Age groups (b) Gender (c) Region

(d) Type of contract (e) Education level (f) Sectors

Figure 4: Weekly change in terminations between first quarter of 2020 relative to average 2017-2019

Terminations Figure 4 reports the evolution of terminations of temporary contracts. In the first 8 weeks of 2020, terminations have been 18% higher on average relative to the period 2017-19.[7] After the start of the pandemic, there was a drop, particularly evident in weeks 10-12, when weekly terminations were on average 38% lower than in previous years. The lockdown may have played a role in determining the drop in terminations, as a lower number of temporary contracts implies also a lower number of terminations. Moreover, some temporary contracts have been suspended, by delaying the termination date to when businesses re-opened.[8] We find that most of the reduction in terminations involves younger workers (panel a), women (panel b), workers in the North and Centre (panel c), workers with low education levels (panel e) and workers in non-shutdown essential sectors (panel f).[9]

[7]The fact that terminations were higher relative to the past in the first two months of the year may be related to the effects of the so-called Dignity Decree, introduced in the summer of 2018. The Decree attempted to limit the use of temporary contracts by firms, by decreasing the maximum length of contracts from 36 to 24 months and the maximum number of renewals from five to four, requiring employers to specify the causes for renewals after the first 12 months.

[8]It is also important to highlight that the special COVID-related STW compensation scheme allowed firms requesting it to roll-over existing temporary contracts and hire temporary workers, contrarily to regular STW which forbids it.

[9]Reassuringly for the quality of the data, panel (d) confirms that separations due to terminations involve temporary contracts only.

Covid Economics 52, 15 October 2020: 152-174

Covid Economics 52, 15 October 2020: 152-174

Quits Figure 5 reports the evolution of quits. Voluntary quits were on average 15% higher in the first 8 weeks of 2020 than what registered in the period 2017-19. The first week of the pandemic recession – week 9 – sees a spike in quits, which reach a level about 40% higher than the same week in 2017-19. This spike in quits merits attention. On the one hand, the spike can be the consequence of school closures, which started during that week and became a nationwide decision the week after. With school closures, in the absence of alternative options, some workers may have decided to quit their jobs to take care of their children: in fact, panel (b), suggests that most of the increase in quits in week 9 is due to an increase in female quits, which before the pandemic were contributing only for one third of total quits whilst in that week they contributed for more than half. On the other hand, given the uncertainty determined by the pandemic, quits may have also been a way of anticipating retirement for older workers (or, again, a way for parents to help their sons and daughters with family duties, once schools closed): this is consistent with the evidence presented in panel (a) that shows how the majority of quits came from middle aged and old workers. Surprisingly, quits were also more present in the South in week 9 rather than the North (panel c), which was the first part of the country to be hit by the pandemic. Therefore, firms may have used quits as an alternative to firings, either by bargaining with the worker in order to reduce labor costs or by forcing workers to blank resignations.[10]

After week 9 we start to observe a downward trend in quits which become negative relative to the past in weeks 11-13, as with the other separation flows. In particular, quits have declined substantially for younger workers, males, for workers in the North of the country, with open-ended contracts, low education levels and working in non-essential activities (see panels a-f of Figure 5).

5 Job Loss Probability

We focus on the job loss probability, by analyzing what categories of workers are more likely to lose their job during the recession. To this end, we identify all contracts active between 24 February 2020 and 31 March 2020 (when our data ends) and all contracts that have ceased over the same period. We select a total of 2.3 million contracts. We then estimate the following cross-sectional linear probability model:

$$s_i = \alpha + \beta X_i + \delta_{s(i)} + \phi_{o(i)} + \epsilon_i, \tag{1}$$

[10] The latter possibility is however much more difficult to materialize, as the Jobs Act (Legislative Decree 151/2015) changed the procedure to communicate quits – which have to be done online through specific online forms provided by the Ministry of Labor – and explicitly forbid employers to make any changes to those forms.

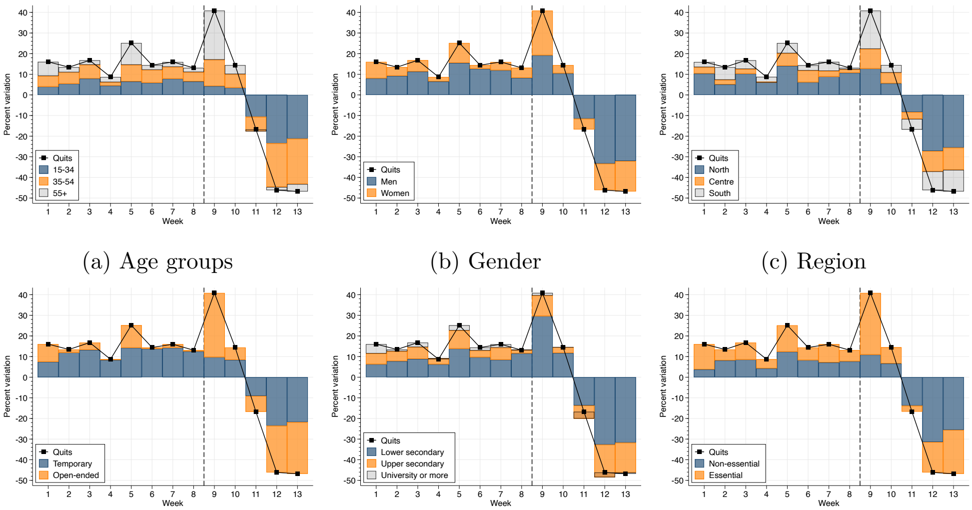

(a) Age groups (b) Gender (c) Region

(d) Type of contract (e) Education level (f) Sectors

Figure 5: Weekly change in quits between first quarter of 2020 relative to average 2017-2019

where s_i is a dummy equal to one if contract i ceases in the period under analysis (because of termination, firing or quit). α is a constant. X_i is a vector of observables that includes a dummy for female workers, the type of contract (temporary or open-ended), the geographical area of work (North, Centre or South), the level of education (lower secondary, upper secondary, university), age and the type of worker (employee, farmer, domestic worker or "not linked"[11]). Finally, $\delta_{s(i)}$ and $\phi_{o(i)}$ are sector and occupation fixed effects (both at a 6-digit level). We are interested in the vector of coefficients β, which measures the correlation between the vector of characteristics $\mathbf{X}$ and the separation probability.

Table 2 reports the results of the estimation of equation (1). Columns (1) and (2) use the whole sample of employees (linked with social security records and not linked), farming and domestic workers. Columns (3) and (4) restrict the sample to employees only. Columns (2) and (4) further restrict the sample excluding separations due to firings, since this was the separation margin most affected by the firing freeze policy. The estimates show that being on a temporary contract implies a higher separation probability of approximately 8 p.p. (ranging from 7.8 to 8.4 p.p depending on the sample restriction). Moreover, older workers are less likely to separate (being 10 years older implies a 0.4-0.7 p.p. lower separation probability). Working in the Centre or South implies a higher separation probability relative

[11]This category contains workers not linked with administrative social security records: mainly, workers on seasonal contracts or collaboration contracts.

Covid Economics 52, 15 October 2020: 152-174

to the North, whereas higher levels of education shield against job loss relative to workers with lower secondary education. Moreover, if we focus on columns (1) and (2) only, employees, farmers and domestic workers are less likely to separate from their jobs relative to the residual category of workers on less stable contracts. Finally, gender is not a significant predictor of the job loss probability. Although the magnitude of coefficients changes according to the sample used, the results are consistent across samples.

Table 3 uses the full sample and reports results from the estimation of equation (1) separately for workers employed in essential activities (columns 1 and 2) and non-essential ones (columns 3 and 4) and again restricting the sample to exclude firings in columns (2) and (4). The sign of the coefficients is the same across both subgroups, but estimates are higher in magnitude in the subgroup of non-essential activities. For example, workers on temporary contracts are 9.3-9.5 p.p. more likely to lose their job relative to permanent workers in non-essential activities, as opposed to a point estimate of 7.5 p.p. in essential activities (although there is a wide overlap in confidence intervals between the two subgroups). Similarly, the coefficients on age, on the dummy for full-time workers and workers in the South, and on the indicator for farming workers are larger in non-essential activities. In other words, being employed in non-essential shutdown sectors is more penalizing in terms of job loss probability for these categories of workers relative to those employed in essential activities. On the contrary, the gap in the job loss probability between workers with upper secondary education and lower secondary education is smaller in non-essential activities than in essential activities and the coefficients for workers with university degree are similar across subgroups. Finally, we find that women are slightly less likely to lose their jobs in non-essential activities relative to men. We further explore differences by gender in the next paragraph.

Differences by gender The evidence presented so far is in line with gender not being a significant explanatory variable for the job loss probability. However, a number of papers highlight how this recession may be particularly harmful for female employment, because it has hit the service sector more, and the presence of female workers (Alon et al., 2020; Adams-Prassl et al., 2020) is higher in that sector. In addition, school closures and working from home burdened women with additional time devoted to childcare and household chores (Del Boca et al., 2020). We further investigate gender differences in job loss probability, by estimating equation (1), gradually including controls and fixed effects and focusing on the coefficient on gender only, to understand whether the rich set of controls included explain the insignificance of the gender dummy. The results are presented in Figure 6, panel (a), which reports estimates from four different specifications. The first one (blue circle) includes only the dummy for female workers as explanatory variable. The second one (red diamond)

Covid Economics 52, 15 October 2020: 152-174

Table 2: Determinants of job loss probability

	Full sample		Employees only	
	(1)	(2)	(3)	(4)
	All	No firings	All	No firings
Woman	-0.25	-0.23	-0.18	-0.12
	(0.17)	(0.16)	(0.10)	(0.09)
Full-time	-0.72*	-0.78**	-0.86**	-0.89**
	(0.30)	(0.27)	(0.19)	(0.17)
Age	-0.04**	-0.04**	-0.07**	-0.07**
	(0.01)	(0.01)	(0.01)	(0.01)
Temporary contract	8.26**	8.38**	7.83**	7.96**
	(1.09)	(1.07)	(1.10)	(1.08)
Apprenticeship	0.78	1.10	-1.50	-1.12
	(1.13)	(1.07)	(0.94)	(0.84)
Centre	0.66*	0.63*	0.50**	0.48**
	(0.28)	(0.28)	(0.18)	(0.17)
South	1.89**	1.65**	1.94**	1.66**
	(0.36)	(0.36)	(0.34)	(0.31)
Upper secondary education	-0.55**	-0.38*	-0.80**	-0.63**
	(0.18)	(0.17)	(0.12)	(0.11)
University or more	-0.97**	-0.73*	-1.27**	-1.00**
	(0.29)	(0.28)	(0.26)	(0.25)
Employee	-2.29**	-2.34**		
	(0.48)	(0.48)		
Farming worker	-1.90**	-1.43**		
	(0.46)	(0.43)		
Domestic worker	-1.63	-0.97		
	(1.03)	(0.61)		
Constant	6.04**	5.37**	5.79**	5.02**
	(0.67)	(0.63)	(0.54)	(0.57)
Occupation fixed effects	Yes	Yes	Yes	Yes
Sector fixed effects	Yes	Yes	Yes	Yes
R^2	0.083	0.089	0.062	0.063
Observations	2,332,888	2,319,918	1,370,157	1,362,254

Notes. The table reports estimates of a linear probability model where the dependent variable is a dummy equal to 1 for contracts ended between 24 February and 31 March 2020. Columns (1) and (2) report results for the full sample of workers. Columns (3) and (4) report results for employees only. Columns (2) and (4) further exclude firings from the sample. Coefficients are multiplied by 100. Standard errors, robust to clustering within 3-digit sectors and occupations, are reported in parentheses. Significance levels: * $p < 0.05$, ** $p < 0.01$.

Table 3: Determinants of job loss probability in essential and non-essential activities

	Essential activities		Non-essential activities	
	(1) All	(2) No firings	(3) All	(4) No firings
Woman	-0.22	-0.24	-0.33**	-0.25**
	(0.27)	(0.24)	(0.09)	(0.08)
Full-time	-0.58	-0.64*	-1.02**	-1.04**
	(0.33)	(0.32)	(0.29)	(0.24)
Age	-0.03**	-0.03**	-0.06**	-0.06**
	(0.01)	(0.01)	(0.02)	(0.02)
Temporary contract	7.46**	7.53**	9.33**	9.52**
	(0.97)	(0.94)	(1.58)	(1.57)
Apprenticeship	1.85	2.14	-0.18	0.16
	(1.56)	(1.48)	(0.75)	(0.74)
Centre	0.80	0.77	0.44*	0.42*
	(0.41)	(0.41)	(0.18)	(0.18)
South	1.53**	1.36**	2.54**	2.17**
	(0.46)	(0.48)	(0.48)	(0.43)
Upper secondary education	-0.72*	-0.51	-0.33*	-0.22
	(0.33)	(0.32)	(0.14)	(0.13)
University or more	-1.01*	-0.76	-1.04**	-0.80**
	(0.44)	(0.43)	(0.25)	(0.25)
Employee	-2.69**	-2.73**	-1.73**	-1.80**
	(0.69)	(0.70)	(0.40)	(0.38)
Farming worker	-1.66**	-1.13**	-2.61*	-2.54*
	(0.41)	(0.38)	(1.29)	(1.19)
Domestic worker	-2.08**	-1.24**	0.26*	0.16*
	(0.43)	(0.19)	(0.11)	(0.07)
Constant	5.98**	5.30**	6.25**	5.58**
	(0.71)	(0.64)	(1.07)	(1.04)
Occupation fixed effects	Yes	Yes	Yes	Yes
Sector fixed effects	Yes	Yes	Yes	Yes
R^2	0.099	0.106	0.064	0.068
Observations	1,447,031	1,439,678	885,855	880,238

Notes. The table reports estimates of a linear probability model where the dependent variable is a dummy equal to 1 for contracts ended between 24 February and 31 March 2020. Columns (1) and (2) report results for workers in essential activities. Columns (3) and (4) report results for workers in non-essential activities. See text for a definition of essential and non-essential activities. Columns (2) and (4) further exclude firings from the sample. Coefficients are multiplied by 100. Standard errors, robust to clustering within 3-digit sectors and occupations, are reported in parentheses. Significance levels: * $p < 0.05$, ** $p < 0.01$.

Covid Economics 52, 15 October 2020: 152-174

Covid Economics 52, 15 October 2020: 152-174

further controls for all variables included in **X**, besides gender (therefore, dummies for full-time, temporary and apprentice contracts, geographical dummies, education dummies and type of worker dummies). The third one (green square) adds sector fixed effects. The fourth one (yellow triangle) adds occupation fixed effects (which is equivalent to the estimate presented in column 1 of Table 2). In all specifications, the coefficient on the gender dummy is not statistically significant at 95% confidence level and the point estimate becomes closer to zero as more explanatory variables are added in the model. Hence, we conclude that – in the Italian labor market at the onset of the pandemic recession – there is no evidence of a different separation probability by gender.

The null effect of gender in the aggregate may be due to different factors. First, we are focusing on the short-time effect of the pandemic recession, since our data covers the first quarter of 2020 only. Differences by gender may take more time to materialize. Second, the policies put in place by the government were intended to protect more vulnerable workers, among which women are more represented. Third, given their presence in non-essential activities is lower relative to men (see section 3), their jobs may have had better chances to survive because less likely to be in shut-down sectors. Finally, if women are employed in tasks that can be more easily performed from home, they are less likely to be separated during the pandemic. Note that the absence of a gender-differentiated effect of the pandemic on hirings and separations does not exclude that adjustments are present on the intensive margin, and that men and women may change differently their number of hours worked which, however, we cannot observe.

Since the null effect on gender may mask heterogeneous effects, we estimate a linear probability model interacting the female dummy with the explanatory variables included in **X** (excluding gender) and we report the coefficients of the interactions. Figure 6, panel (b), shows the estimates. Generally, the null result is confirmed also in subgroups determined by observable worker characteristics. However, we do find a significant effect for women with upper secondary education, who are 0.6 p.p. more likely to lose their job relative to men with the same education level. Furthermore, we find that women employed in the farming sector are 2.2 p.p. less likely to lose their job relative to men, and female domestic workers are 1.8 p.p. more likely to lose their job relative to male domestic workers. However, these two subgroups represent only 3.4% and 6.5% of workers included in the sample and they display a clear gender imbalance, with more men employed as farmers and more women employed as domestic workers. Hence, part of the gender differences in job loss probability may be explained by selection. We therefore confirm there are no strong indications of an overall gender difference in the separation probability. This result is in contrast with the evidence for the US and the UK, but in line with that on Germany (Adams-Prassl et al., 2020), a

Covid Economics 52, 15 October 2020: 152-174

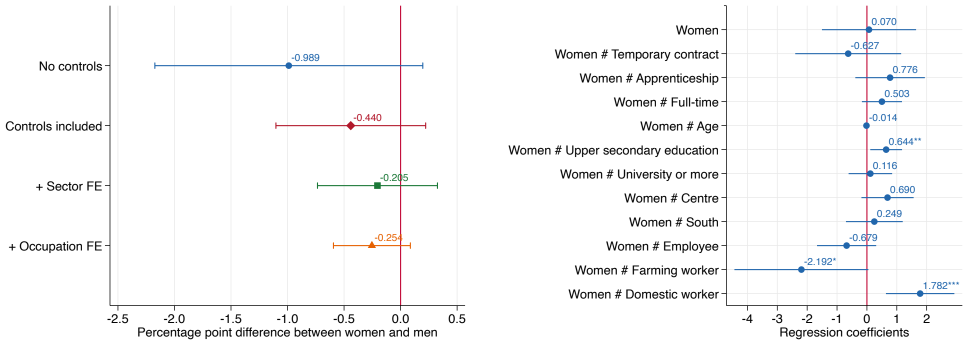

(a) Regression models with different controls (b) Interaction effects

Figure 6: Impact of COVID-19 on job loss probability by gender

country with labor market institutions and policies closer to the Italian ones.

6 Conclusion

This paper explores the short-run heterogeneous effects of COVID-19 on labor market flows in Italy and how policy enacted to reduce the spread of the virus and the disruption of economic activity mediated them.

We show that, before the pandemic, workers employed in non-essential activities shut-down by the government were in majority men, younger than 35 years old, located in the North and with lower levels of education. When looking at the change in hirings and separations and decomposing it by age, gender, region, type of contract (open-ended or temporary), education level, and sector (essential vs non-essential activities), we find that from the nineth week of the year – when the virus started to spread exponentially across the country –, there was a pronounced drop in hirings and terminations. On the contrary, firings and quits spiked right after the nineth week, and then dropped significantly, reflecting the effects of the firing freeze and the easing of access to STW compensation schemes. The firing freeze may also have contributed to the decreasing dynamics of hirings, as the higher employment protection for workers may have decreased turnover. We further explore separations by examining which factors predict the probability of job loss. We find that those workers that were already suffering the consequences of the previous recession (young, temporary, low-skill workers) are those at higher risk of losing their job because of COVID-19. Gender, instead, is a non-significant predictor of job loss in the aggregate, but we do find a significantly higher separation rate for female workers with upper secondary education and female domestic workers.

While we focus on short-term outcomes and cannot account for changes in hours worked, our evidence contributes to the understanding of labor market and policy responses in the wake of the pandemic. The use of detailed administrative data allows us to separately analyze how hirings and separations – distinguishing between firings, terminations and quits – have evolved relative to normal times and how different categories of workers have been affected. Given the critical importance of the firing freeze and the special STW compensation scheme in affecting labor market flows, it is important to monitor the labor market transitions if these policies will be lifted, since they have protected vulnerable workers the most.

Covid Economics 52, 15 October 2020: 152-174

References

Adamopoulou, E., De Philippis, M., Sette, E., and Viviano, E. (2020). The Long Run Earnings Effects of a Credit Market Disruption. IZA Discussion Papers 13185, Institute of Labor Economics (IZA).

Adams-Prassl, A., Boneva, T., Golin, M., and Rauh, C. (2020). Inequality in the impact of the coronavirus shock: Evidence from real time surveys. *Journal of Public Economics*, 189:104245.

Alon, T., Doepke, M., Olmstead-Rumsey, J., and Tertilt, M. (2020). The Impact of COVID-19 on Gender Equality. *Covid Economics: Vetted and Real-Time Papers*, 4:62–85.

Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., and Yannelis, C. (2020). How does household spending respond to an epidemic? Consumption during the 2020 COVID-19 pandemic. *Covid Economics: Vetted and Real-Time Papers*, 18:73–108.

Belot, M., Choi, S., Jamison, J. C., Papageorge, N. W., Tripodi, E., and Van den Broek-Altenburg, E. (2020). Unequal Consequences of Covid-19 across Age and Income: Representative Evidence from Six Countrie. *Covid Economics: Vetted and Real-Time Papers*, 38:196–217.

Bick, A. and Blandin, A. (2020). Real time labor market estimates during the 2020 coronavirus outbreak. *Unpublished Manuscript, Arizona State University.*

Blundell, R., Costa Dias, M., Joyce, R., and Xu, X. (2020). COVID-19 and Inequalities. *Fiscal Studies*, 41(2):291–319.

Boeri, T., Ichino, A., Moretti, E., and Posch, J. (2019). Wage Equalization and Regional Misallocation: Evidence from Italian and German Provinces. NBER Working Papers 25612, National Bureau of Economic Research, Inc.

Cajner, T., Crane, L. D., Decker, R. A., Grigsby, J., Hamins-Puertolas, A., Hurst, E., Kurz, C., and Yildirmaz, A. (2020). The US labor market during the beginning of the pandemic recession. Technical report, National Bureau of Economic Research.

Carvalho, V. M., Hansen, S., Ortiz, A., Ramón García, J., Rodrigo, T., Rodriguez Mora, S., and Ruiz, J. (2020). Tracking the COVID-19 Crisis with High-Resolution Transaction Data. CEPR Discussion Papers 14642, C.E.P.R. Discussion Papers.

Chetty, R., Friedman, J. N., Hendren, N., Stepner, M., et al. (2020). How did COVID-19 and stabilization policies affect spending and employment? A new real-time economic tracker based on private sector data. Technical report, National Bureau of Economic Research.

Crossley, T. F., Fisher, P., and Low, H. (2020). The Heterogeneous and Regressive Consequences of COVID-19: Evidence from High Quality Panel Data. Department of Economics Discussion Paper Series 919, Department of Economics, University of Oxford.

Del Boca, D., Oggero, N., Profeta, P., and Rossi, M. (2020). Women's and men's work, housework and childcare, before and during COVID-19. *Review of Economics of the Household*.

Farré, L., Fawaz, Y., Gonzalez, L., and Graves, J. (2020). How the COVID-19 Lockdown Affected Gender Inequality in Paid and Unpaid Work in Spain. IZA Discussion Papers 13434, Institute of Labor Economics (IZA).

Forsythe, E., Kahn, L. B., Lange, F., and Wiczer, D. (2020). Labor demand in the time of COVID-19: Evidence from vacancy postings and UI claims. *Journal of Public Economics*, 189:104238.

Garin, A. and Silvério, F. (2019). How Responsive are Wages to Demand within the Firm? Evidence from Idiosyncratic Export Demand Shocks. Technical report.

Hupkau, C. and Petrongolo, B. (2020). Work, care and gender during the Covid-19 crisis. CEP Covid-19 Briefings cepcovid-19-002, Centre for Economic Performance, LSE.

IMF (2020). World Economic Outlook Update. Report, Internation Monetary Fund.

Istat (2020). Esame del disegno di legge A.S. 1766 Conversione in legge del decreto-legge 17 marzo 2020, n. 18. Report, Istat.

Covid Economics 52, 15 October 2020: 152-174

Izquierdo, M., Jimeno, J. F., Kosma, T., Lamo, A., Millard, S., Rõõm, T., and Viviano, E. (2017). Labour market adjustment in Europe during the crisis: microeconomic evidence from the Wage Dynamics Network survey. *Banco de Espana Occasional Paper*, (1704).

Kugler, A. and Pica, G. (2008). Effects of employment protection on worker and job flows: Evidence from the 1990 Italian reform. *Labour Economics*, 15(1):78 – 95.

OECD (2020). *OECD Employment Outlook 2020*.

von Gaudecker, H.-M., Holler, R., Janys, L., Siflinger, B., and Zimpelmann, C. (2020). Labour Supply in the Early Stages of the Covid-19 Pandemic: Empirical Evidence on Hours, Home Office, and Expectations.

Covid Economics 52, 15 October 2020: 152-174