CENTRE FOR
ECONOMIC
POLICY
RESEARCH

CEPR PRESS

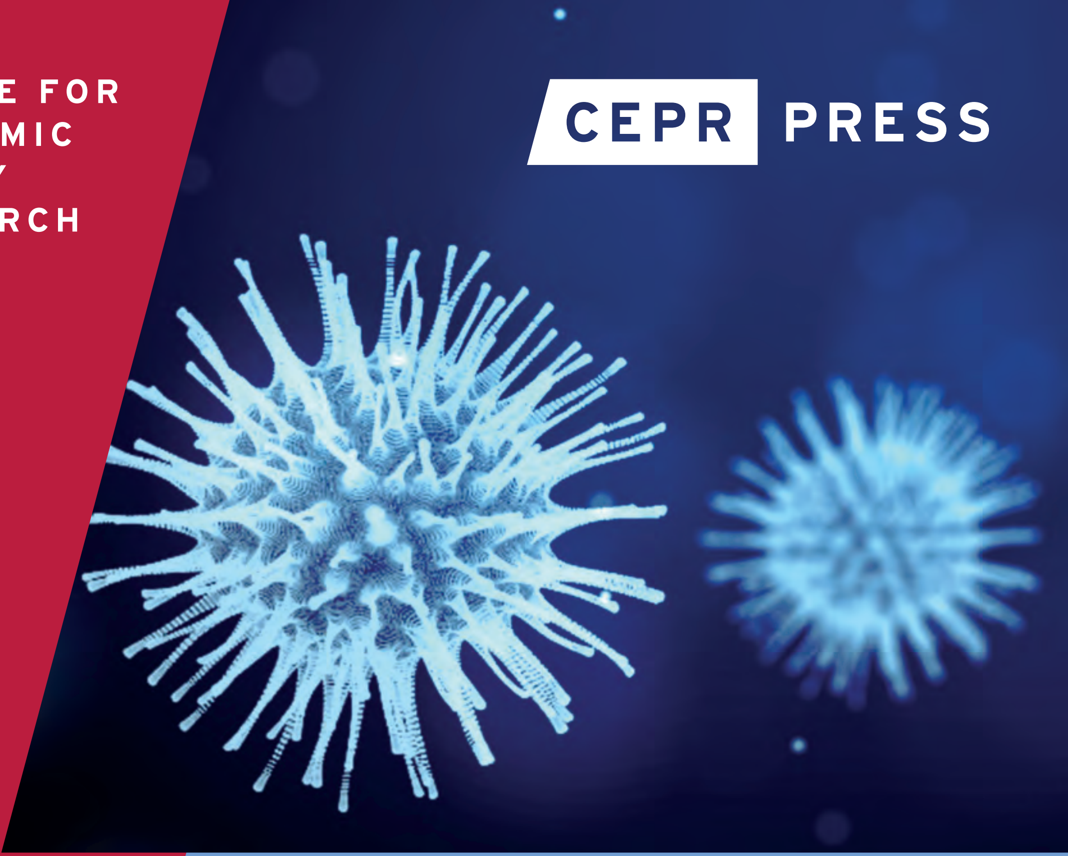

COVID ECONOMICS

VETTED AND REAL-TIME PAPERS

ISSUE 8
22 APRIL 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-real-time-journal-cej. 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.

Ethics

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

Editorial Board

Covid Economics
Vetted and Real-Time Papers

Issue 8, 22 April 2020

Contents

Scarring body and mind: The long-term belief-scarring effects of Covid-19[1]

Julian Kozlowski,[2] Laura Veldkamp[3] and Venky Venkateswaran[4]

Date submitted: 13 April 2020; **Date accepted:** 17 April 2020; **Date revised:** 8 September 2020

Covid Economics 8, 22 April 2020: 1-40

The largest economic cost of the COVID-19 pandemic could arise from changes in behavior long after the immediate health crisis is resolved. A potential source of such a long-lived change is scarring of beliefs, a persistent change in the perceived probability of an extreme, negative shock in the future. We show how to quantify the extent of such belief changes and determine their impact on future economic outcomes. We find that the long-run costs for the U.S. economy from this channel is many times higher than the estimates of the short-run losses in output. This suggests that, even if a vaccine cures everyone in a year, the COVID-19 crisis will leave its mark on the US economy for many years to come.

1 The views expressed are those of the authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors. We thank Kenneth Rogoff for a generous and insightful discussion, as well as Dean Corbae and Pablo D'Erasmo for sharing data on corporate defaults and Marco Del Negro and Andrea Tambalotti for liquidity and interest rate data.

2 Senior Economist, Federal Reserve Bank of St. Louis.
3 Professor of Finance, Columbia University and CEPR Research Fellow.
4 Associate Professor of Economics, Stern School of Business, New York University.

Covid Economics 8, 22 April 2020: 1-40

One of the most pressing questions of the day is the economic costs of the COVID-19 pandemic. While the virus will eventually pass, vaccines will be developed, and workers will return to work, an event of this magnitude could leave lasting effects on the nature of economic activity. Economists are actively debating whether the recovery will be V-shaped, U-shaped or L-shaped.[1] Much of this discussion revolves around confidence, fear and the ability of firms and consumers to rebound to their old investment and spending patterns. Our goal is to formalize this discussion and quantify these effects, both in the short- and long-run. To explore these conjectures about the extent to which the economy will rebound from this COVID-induced downturn, we use a standard economic and epidemiology framework, with one novel channel: a "scarring effect." Scarring is a persistent change in beliefs about the probability of an extreme, negative shock to the economy. We use a version of Kozlowski et al. (2020), to formalize this scarring effect and quantify its long-run economic consequences, under different scenarios for the dynamics of the crisis.

We start from a simple premise: No one knows the true distribution of shocks in the economy. Consciously or not, we all estimate the distribution using past events, like an econometrician would. Tail events are those for which we have little data. Scarce data makes new tail event observations particularly informative. Therefore, tail events trigger larger belief revisions. Furthermore, because it will take many more observations of non-tail events to convince someone that the tail event really is unlikely, changes in tail risk beliefs are particularly persistent.

We have seen the scarring effect in action before. Before 2008, few people entertained the possibility of a financial crisis in the US. Today, more than a decade after the Global Financial Crisis, the possibility of another run on the financial sector is raised frequently, even though the system today is probably much safer (Baker et al., 2019). Likewise, businesses will make future decisions with the risk of another pandemic in mind. Observing the pandemic has taught us that the risks were greater than we thought. It is this new-found knowledge that has long-lived effects on economic choices.

To explore tail risk in a meaningful way, we need to use an estimation procedure that does not constrain the shape of the distribution's tail. Therefore, we allow our agents to learn about the distribution of aggregate shocks non-parametrically. Each period, agents observe one more piece of data and update their estimates of the distribution. Section I shows how this process leads to long-lived responses of beliefs to transitory events, especially extreme, unlikely ones. The mathematical foundation for such persistence is the martingale property of beliefs. The logic is that once observed, the event remains in agents' data set. Long after the direct effect of the shock has passed, the knowledge of that tail event affects beliefs and therefore, continues to restrain economic activity.

[1] See e.g., Summers (FT, 2020), Krugman (2020), Reinhart and Rogoff (2020) and Cochrane (2020).

Covid Economics 8, 22 April 2020: 1-40

To illustrate the economic importance of these belief dynamics, Section II embeds our belief updating tool in a macroeconomic model with an epidemiology event that erodes the value of capital. This framework is designed to link tail events like the current crisis to macro outcomes in a quantitatively plausible way and has been used – e.g. by Gourio (2012) and Kozlowski et al. (2020) – to study the 2008-09 Great Recession. It features, among other elements, bankruptcy risk and elevated capital depreciation from social distancing, which separates labor from capital. Section III describes the data we feed into the model to discipline our belief estimates. Section IV combines model and data and uses the resulting predictions to show how belief updating can generate large, persistent losses. We compare our results to those from the same economic model, but with agents who have full knowledge of the distribution, to pinpoint belief updating as the source of the persistence.

We model the economic effects of the COVID-19 crisis as a combination of a productivity decline and accelerated capital obsolescence. We use the well-known SEIR (susceptible-exposed-infected-recovered) framework from the epidemiology literature to model the disease spread. But, it is the response to the disease that is the source of the adverse economic shock in our model. Our structure is capable of generating large asset price fluctuations, of the order observed at the onset of the pandemic, and provides a simple mapping from social distancing policies and other mitigation behavior to economic costs. It also allows us to connect to existing studies on tail risk in macroeconomics and finance. We present results for different scenarios, reflecting the considerable uncertainty about outcomes even in the short-run. Our point is not to make a forecast of the coming year's events but that that whatever you think will happen over the next year, the ultimate costs of this pandemic are much larger than your short-run calculations suggest.

In the first scenario, GDP drops by about 9% in 2020, recovers gradually but does not go back to its previous trajectory. It persistently remains about 4% below the previous pre-COVID steady state. The discounted value of the lost output is almost 10 times the 2020 drop and belief revisions account for bulk of the losses (almost 6 times the short-run effect). Greater tail risk makes investing less attractive, reducing the stock of productive capital and (and therefore, labor input demand) persistently. In the second scenario, which captures a milder mitigation response to the spread of the disease, both short- and long-run economic costs are longer, but the relative importance of belief revisions remains the same.

The model also makes a number of predictions about asset prices. Interestingly, after an initial shock, credit spreads and equity valuations are predicted to roughly return to their original level. This is because firms respond to this increase in riskiness by cutting back on debt. The effects of scarring are more clearly noticeable in options prices. In scenario 1, for example, the option-implied third moment in the risk-neutral distribution of equity returns becomes significantly more negative.

For monetary policy makers, one of the most pressing questions is how belief scarring will affect the long-run natural rate of interest, often referred to as “r-star." Following the onset of COVID in the U.S., interest rates declined rapidly. A significant portion of that decline is related to demand for liquidity. In order to understand how much of that decline was temporary and how much permanent – and more broadly about the interaction of liquidity and scarring – we introduce a role for liquid assets in an extension of our baseline model in Section V. When most capital is only partially pledgeable, but riskless assets are fully pledgeable, riskless assets, of course have more value. But what we learn is that value is sensitive to tail risk. A persistent increase in perceived risk from COVID-19 depresses the long-run natural rate of interest by 67 basis points.

Our results also imply that a policy that prevents capital depreciation or obsolescence, even if it has only modest immediate effects on output, can have substantial long-run benefits, several times larger than the short-run considerations that often dominate policy discussion. Obviously, no policy can prevent people from believing that future pandemics are more likely than they originally thought, but policy can change how the ongoing crisis affects capital returns. By changing that mapping, the costs of belief scarring can be mitigated. For example, bankruptcies can lead to destruction of specific investments and a permanent erosion in the value of capital. Interventions which prevent widespread bankruptcies can thus limit the adverse effects of the crisis on returns and yield substantial long-run benefits. While the short-run gains from limiting bankruptcies is well-understood, our analysis shows that neglecting the effect on beliefs leads one to drastically underestimate the benefits of such policies.

Of course, future governments could also invest in public health to mitigate the cost of future pandemics. The ability of such an investment to heal beliefs depends on the nature of belief changes induced by this episode. If we onlu updated our beliefs about the ability of a particular type of communicable diseases to disrupt economic activity, then health investments will be highly effective. However, traumatic events often leave survivors with a more general sense that unexpected, disastrous events can arise without warning. This more amorphous fear will be much harder for policy to combat.

Comparison to the literature There are many new studies of the impact of the COVID-19 pandemic on the U.S. economy, both model-based and empirical. Alvarez et al. (2020), Eichenbaum et al. (2020) and Farboodi et al. (2020) use simple economic frameworks to analyze the costs of the disease and the associated mitigation strategies. Leibovici et al. (2020) use an input-output structure to investigate the extent to which a shock to contact-intensive industries can propagate to the rest of the economy. Koren and Pető (2020) build a detailed theory-

Covid Economics 8, 22 April 2020: 1-40

based measures of the reliance of U.S. businesses on human interaction. On the empirical side, Ludvigson et al. (2020) use VARs to estimate the cost of the pandemic over the next few months, while Carvalho et al. (2020) use high-frequency transaction data to track expenditure and behavior changes in real-time. We add to this discussion by focusing on the long-term effects from changes in behavior that persist long after the disease is gone.

Other papers share our focus on long-run effects. Jorda et al. (2020) study rates of return on assets using a data-set stretching back to the 14th century, focusing on 15 major pandemics (with more than 100,000 deaths). Their evidence suggests a sustained downward pressure on interest rates, decades after the pandemic, consistent with long-lasting macroeconomic after-effects. Reinhart and Rogoff (2009) examine long-lived effects of financial crises. Correia et al. (2020) find evidence of persistent declines in economic activity following the 1918 influenza pandemic. A few papers also use beliefs but rely on other mechanisms, such as financial frictions, for propagation. Elenev et al. (2020) and Krishnamurthy and Li (2020) propagate the shock primarily through financial balance sheet effects. In a more informal discussion, Cochrane (2020) explores whether the recovery from the COVID-shock will be V, U or L shaped. This work formalizes many of the ideas in that discussion.

Outside of economics, biologists and socio-biologists have noted long ago that epidemics change the behavior of both humans and animals. Loehle (1995) explore the social barriers to transmission in animals as a mode of defense against pathogen attack. Past disease events have effects on mating strategies, social avoidance, group size, group isolation, and other behaviors for generations. Gangestad and Buss (1993) find evidence of similar behavior among human communities.

In the economics realm, a small number of uncertainty-based theories of business cycles also deliver persistent effects from other sorts of transitory shocks. In Straub and Ulbricht (2013) and Van Nieuwerburgh and Veldkamp (2006), a negative shock to output raises uncertainty, which feeds back to lower output, which in turn creates more uncertainty. To get even more persistence, Fajgelbaum et al. (2017) combine this mechanism with an irreversible investment cost, a combination which can generate multiple steady-state investment levels. These uncertainty-based explanations are difficult to embed in quantitative DSGE models and to discipline with macro and financial data.

Our belief formation process is similar to the parameter learning models by Johannes et al. (2016), Cogley and Sargent (2005) and Kozeniauskas et al. (2014) and is similar to what is advocated by Hansen (2007). However, these papers focus on endowment economies and do not analyze the potential for persistent effects in a setting with production.[2] The most important

[2] Other learning papers in this vein include papers on news shocks, such as, Beaudry and Portier (2004), Lorenzoni (2009), Veldkamp and Wolfers (2007), uncertainty shocks, such as Jaimovich and Rebelo (2006),

Covid Economics 8, 22 April 2020: 1-40

difference is that our non-parametric approach allows us to incorporate beliefs about tail risk.

I Belief Formation

Before laying out the underlying economic environment, we begin by explaining how we formalize the notion of belief scarring, the non-standard, but most crucial part of our analysis. We then embed it in an economic environment and quantify the effect of belief changes from the COVID-19 pandemic on the US economy.

No one knows the true distribution of shocks to the economy. All of us – whether in our capacity as economic agents or modelers or econometricians – estimate such distributions, updating our beliefs as new data arrives. Our goal is to model this process in a reasonable and tractable fashion. The first step is to choose a particular estimation procedure. A common approach is to assume a normal or other parametric distribution and estimate its parameters. The normal distribution, with its thin tails, is unsuited to thinking about changes in tail risk. Other distributions raise obvious concerns about the sensitivity of results to the specific distributional assumption used. To minimize such concerns, we take a non-parametric approach and let the data inform the shape of the distribution.

Specifically, we employ a kernel density estimation procedure, one of most common approaches in non-parametric estimation. Essentially, it approximates the true distribution function with a smoothed version of a histogram constructed from the observed data. By using the widely-used normal kernel, we impose a lot of discipline on our learning problem but also allow for considerable flexibility. We also experimented with a handful of other kernels.

Consider a shock $\tilde{\phi}_t$ whose true density g is unknown to agents in the economy. The agents do know that the shock $\tilde{\phi}_t$ is i.i.d. Their information set at time t, denoted $\mathcal{I}_t$, includes the history of all shocks $\tilde{\phi}_t$ observed up to and including t. They use this available data to construct an estimate $\hat{g}_t$ of the true density g. Formally, at every date, agents construct the following normal kernel density estimator of the pdf g

$$\hat{g}_t\left(\tilde{\phi}\right) = \frac{1}{n_t \kappa_t} \sum_{s=0}^{n_t-1} \Omega\left(\frac{\tilde{\phi} - \tilde{\phi}_{t-s}}{\kappa_t}\right) \tag{1}$$

where $\Omega(\cdot)$ is the standard normal density function, κ_t is the smoothing or bandwidth parameter and n_t is the number of available observations of at date t. As new data arrives, agents add the new observation to their data set and update their estimates, generating a sequence of beliefs $\{\hat{g}_t\}$.

Bloom et al. (2018), Nimark (2014) and higher-order belief shocks, such as Angeletos and La'O (2013) or Huo and Takayama (2015).

Covid Economics 8, 22 April 2020: 1-40

Covid Economics 8, 22 April 2020: 1-40

The key mechanism in the paper is the persistence of belief changes induced by transitory $\tilde{\phi}_t$ shocks. This stems from the martingale property of beliefs – i.e. conditional on time-t information ($\mathcal{I}_t$), the estimated distribution is a martingale. Thus, on average, the agent expects her future belief to be the same as her current beliefs. This property holds exactly if the bandwidth parameter κ_t is set to zero and holds with tiny numerical error in our application.[3] In line with the literature on non-parametric assumption, we use the optimal bandwidth.[4] As a result, any changes in beliefs induced by new information are expected to be approximately permanent. This property plays a central role in generating long-lived effects from transitory shocks.

II Economic and Epidemiological Model

To gauge the magnitude of the scarring effect of the COVID-19 pandemic on long-run economic outcomes, we need to embed it in an economic model in which tail risk and belief changes can have meaningful effects. For this, a model needs two key features. First, it should have the potential for 'large' shocks, that have both transitory and lasting effects. The former would include lost productivity from stay-at-home orders preventing services from reaching consumers. But for this shock to look like the extreme event it is to investors, the model must also allow for the possibility of a more persistent loss of productive capital. This loss represents the interior of the restaurant that went bankrupt, or the unused capacity of the hotel that will not fill again for many years to come. When stay-at-home orders forced consumers to work and consume differently, it persistently altered tastes and habits, rendering some capital obsolete. One might think this is hard-wiring persistence in the model. Yet, as we will show, this loss of capital by itself has a short lived effect and typically triggers an investment boom, as the economy rebuilds capital better suited to the new consumption normal. We explore two possible scenarios that highlight the enormous importance of preventing capital obsolescence, because of the scarring of beliefs.

The second key feature is sufficient curvature in policy functions, which serves to make

[3]As $\kappa_t \to 0$, the CDF of the kernel converges to $\hat{G}_t^0\left(\tilde{\phi}\right) = \frac{1}{n_t}\sum_{s=0}^{n_t-1} \mathbf{1}\left\{\tilde{\phi}_{t-s} \leq \tilde{\phi}\right\}$. Then, for any $\tilde{\phi}$, and any $j \geq 1$

$$\mathbb{E}_t\left[\hat{G}_{t+j}^0\left(\tilde{\phi}\right)\Big|\mathcal{I}_t\right] = \mathbb{E}_t\left[\frac{1}{n_t+j}\sum_{s=0}^{n_t+j-1}\mathbf{1}\left\{\tilde{\phi}_{t+j-s} \leq \tilde{\phi}\right\}\Bigg|\mathcal{I}_t\right] = \frac{n_t}{n_t+j}\hat{G}_t^0\left(\tilde{\phi}\right) + \frac{j}{n_t+j}\mathbb{E}_t\left[\mathbf{1}\left\{\tilde{\phi}_{t+1} \leq \tilde{\phi}\right\}\Big|\mathcal{I}_t\right]$$

Thus, future beliefs are, in expectation, a weighted average of two terms - the current belief and the distribution from which the new draws are made. Since our best estimate for the latter is the current belief, the two terms are exactly equal, implying $\mathbb{E}_t\left[\hat{G}_{t+j}^0\left(\tilde{\phi}\right)\Big|\mathcal{I}_t\right] = \hat{G}_t^0\left(\tilde{\phi}\right)$.

[4]See Hansen (2015).

economic activity sensitive to the probability of extreme large shocks. Two ingredients – namely, Epstein-Zin preferences and costly bankruptcy – combine to generate significant non-linearity in policy functions.

It is important to note that none of these ingredients guarantee persistent effects. Absent belief revisions, shocks, no matter how large, do not change the long-run trajectory of the economy. Similarly, the non-linear responses induced by preferences and debt influence the size of the economic response, but by themselves do not generate any internal propagation. They simply govern the magnitude of the impact, both in the short and long run.

To this setting, we add belief scarring. We model beliefs using the non-parametric estimation described in the previous section and show how to discipline this procedure with observable macro data, avoiding free parameters. This belief updating piece is not there to generate the right size reaction to the initial shock. Instead, belief updating adds the persistence, which considerably inflates the cost.

II.A The Disease Environment

This block of the model serves to generate a time path for disruption to economic activity, which will then be mapped into transitory productivity shock and capital obsolescence . Of course, we could have directly created scenarios for the shocks and arrived at the same predictions. The explicit modeling of the spread of disease allows us to see how different social distancing policies map into shocks and ultimately into long-term economic costs from belief scarring. Given this motivation, we build on a very simple SEIR model, which is a discrete-time version of Atkeson (2020) or Stock (2020), who build on work in the spirit of Kermack and McKendrick (1927). To this model, we add two ingredients: 1) a behavioral/policy rule that imposes capital idling when the infection rate increase (for example, this rule could represent optimal behavior or government policy); and 2) a higher depreciation rate of unused capital. While we normally think of capital utilization depreciating capital, this is a different circumstance where habits, technologies and norms are changing more rapidly than normal. Unused capital may be restaurants whose customers find new favorites, old conferencing technologies replaced with new online technology or office space that will be replaced with home offices. This higher depreciation rate represents a speeding up of capital obsolescence.

Disease and shutdowns On January 20 2020, the first case of COVID was documented in the U.S. Therefore, we start our model on that day, with one infected person. Because we are examining persistence mechanisms, we want to impose a clear end date to the COVID shock. Therefore, we assume that COVID-19 will be over by the end of 2020. The SEIR model predicts the evolution of the pandemic. Our policy shutdown rule, maps the infection rate series into a

Covid Economics 8, 22 April 2020: 1-40

Covid Economics 8, 22 April 2020: 1-40

value for the aggregate shock to the US economy in 2020. From 2021 onwards, we assume that COVID-19 will be over. However, we explore scenarios where the economy may suffer other pandemics in the future.

Time is discrete and infinite. For the disease part of the model, we will count time in days, indexed by $\tilde{t}$. Later, to describe long-run effects, we will change the measure of time to t, which represents years. There are N agents in the economy. At date 1, the first person gets infected. Let S represent the number of people susceptible to the disease, but not currently exposed, infected, dead or recovered. At date 1, that susceptible number is $S(1) = N - 1$. Let E be the number of exposed persons and I be the number infected. We start with $E(1) = 0$ and $I(1) = 1$. Finally, D represents the number who are either recovered or dead, where $D(1) = 0$. The following four equations describe the dynamics of the disease.

$$S(\tilde{t}+1) = S(\tilde{t}) - \tilde{\beta}_{\tilde{t}} S(\tilde{t}) I(\tilde{t})/N \tag{2}$$

$$E(\tilde{t}+1) = E(\tilde{t}) + \tilde{\beta}_{\tilde{t}} S(\tilde{t}) I(\tilde{t})/N - \sigma_E E(\tilde{t}) \tag{3}$$

$$I(\tilde{t}+1) = I(\tilde{t}) + \sigma_E E(\tilde{t}) - \gamma_I I(\tilde{t}) \tag{4}$$

$$D(\tilde{t}+1) = D(\tilde{t}) + \gamma_I I(\tilde{t}) \tag{5}$$

The parameter γ_I is the rate at which people exit infection and become deceased or recovered. Thus, the expected duration of infection is approximately $1/\gamma_I$, and the number of contacts an infected person has with a susceptible person is $\tilde{\beta}$ times the fraction of the population that is susceptible $S(\tilde{t})/N$. The initial reproduction rate, often referred to as R_0 is therefore $\tilde{\beta}/\gamma_I$.

We put a t subscript on $\tilde{\beta}_{\tilde{t}}$ because behavior and policy can change it. When the infection rate rises, people reduce infection risk by staying home. This reduces the number of social contacts, reducing $\tilde{\beta}$. Lockdown policies also work by reducing $\tilde{\beta}$. We capture this relationship by assuming that $\tilde{\beta}$ can vary between a maximum of $\gamma_I R_0$ and a minimum of $\gamma_I R_{min}$. R_{min} is the estimated U.S. reproduction rate for regions under lockdown. Where on the spectrum the contact rate lies depends on the last 30-day change in infection rates, measured with a 15-day lag.[5] Let ΔI_t be the difference between the average 15-29 day past infections and the average of 30-44 day infections: $\Delta I_t = (1/15)\left(\sum_{\tau=15}^{29} I(t-\tau) - \sum_{\tau=30}^{44} I(t-\tau)\right)$. This captures the fact that most policy makers are basing policy on two-week changes in hospitalization rates, which are themselves observed with a 14-day lag. Then policy and individual behavior achieves a

[5]This is consistent with the U.S. official policy on re-opening (CDC, 2020). Note that individual optimal choice to social distance are also included in this "policy." These optimal choices look similar. See Kaplan et al. (2020).

frequency of social contact:

$$\tilde{\beta}_{\tilde{t}} = \gamma_I \times min(R_0, max(R_{min}, R_0 - \zeta * \Delta I_t)) \tag{6}$$

The key part of the epidemic from a belief-scarring perspective is that reducing the contact rate requires separating labor from capital. In other words, capital is idle. No capital is idled (full capacity) when no mitigation efforts are underway, i.e. when $\tilde{\beta}_{\tilde{t}} = \gamma_I R_0$. But as $\tilde{\beta}_{\tilde{t}}$ falls, capital idling (K^-) rises. We formalize that relationship as

$$K^-_{\tilde{t}} = \tilde{\theta} * (R_0 - \tilde{\beta}_{\tilde{t}}/\gamma_I). \tag{7}$$

Idle capital depreciates as a rate $\tilde{\delta}$. As mentioned before, this is not physical deterioration of the capital stock. Instead, it represent a loss of value from accelerated obsolescence due to changes in tastes, habits and technologies. It could also represent a loss in value because of persistent upstream or downstream supply chain constraints.

II.B The Economy

Preferences and technology: To describe long-term economic consequences, we switch from the daily time index $\tilde{t}$ to an annual time index t. An infinite horizon, discrete time economy has a representative household, with preferences over consumption (C_t) and labor supply (L_t):

$$U_t = \left[(1-\beta) \left(c_t^{\gamma} (1-l_t)^{1-\gamma} \right)^{1-\psi} + \beta E_t \left(U_{t+1}^{1-\eta} \right)^{\frac{1-\psi}{1-\eta}} \right]^{\frac{1}{1-\psi}} \tag{8}$$

where ψ is the inverse of the inter-temporal elasticity of substitution, η indexes risk-aversion, γ indexes the share of consumption in the period utility function, and β represents time preference.

The economy is also populated by a unit measure of firms, indexed by i and owned by the representative household. Firms produce output with capital and labor, according to a standard Cobb-Douglas production function $z_t k_{it}^{\alpha} l_{it}^{1-\alpha}$.

Aggregate uncertainty is captured by a single random variable, $\tilde{\phi}_t$, which is i.i.d. over time and drawn from a distribution $g(\cdot)$. The i.i.d. assumption is made in order to avoid an additional, exogenous, source of persistence.[6] The effect of this shock on economic activity depends on the realized default rate Def_t (the fraction of firms who default in t, characterized later in this section). Formally, it induces a capital obsolescence 'shock' $\phi_t \equiv \Phi(\tilde{\phi}_t, Def_t)$. The function $\Phi(\cdot)$ will made explicit later. This composite shock has both permanent and transitory

[6] The i.i.d. assumption also has empirical support. In the next section, we use macro data to construct a time series for $\tilde{\phi}_t$. We estimate a (statistically insignificant) autocorrelation of 0.15.

Covid Economics 8, 22 April 2020: 1-40

effects. The permanent component works as follows: a firm that enters the period t with capital $\hat{k}_{it}$ has effective capital $k_{it} = \phi_t \hat{k}_{it}$.

In addition to this permanent component, the shock ϕ_t also has a temporary effect, through the TFP term $z_t = \phi_t^{\nu}$. The parameter ν governs the relative strength of the transitory component. This specification allows us to capture both permanent and transitory disruptions with only one source of uncertainty. By varying ν, we can capture a range of scenarios without having to introduce additional shocks.

Firms are also subject to an idiosyncratic shock v_{it}. These shocks scale up and down the total resources available to each firm (after paying labor, but before paying debtholders' claims)

$$\Pi_{it} = v_{it} \left[z_t k_{it}^{\alpha} l_{it}^{1-\alpha} - W_t l_{it} + (1-\delta) k_{it} \right] \tag{9}$$

where δ is the ordinary rate of capital depreciation. The additional obsolescence from idle capital is already removed from k_{it}, via the shock ϕ_t. The shocks v_{it} are i.i.d. across time and firms and are drawn from a known distribution, F.[7] The mean of the idiosyncratic shock is normalized to be one: $\int v_{it} \quad di = 1$. The primary role of these shocks is to induce an interior default rate in equilibrium, allowing a more realistic calibration, particularly of credit spreads.

What is capital obsolescence? Capital obsolescence shock reflects a long-lasting change in the economic value of the average unit of capital. A realization of $\phi < 1$ captures the loss of specific investments or other forms of lasting damage from a prolonged shutdown. This could come from the lost value of cruise ships that will never sail again, businesses that do not re-open, loss of customer capital or just less intensive use of commercial space due to a persistent preference for more distance between other diners, travelers, spectators or shoppers. It could also represent permanent changes in health and safety regulations that make transactions safer, but less efficient from an economic standpoint.

An important question is whether future investment could be made in ways or in sectors that avoid these costs. Of course, such substitution is likely to happen to some extent. But, the fact that the patterns of investment were not chosen previously suggests that these adjustments are costly or less profitable. More importantly, we learned that the world is riskier and more unpredictable than we thought. The shocks that hit one sector (or type of capital) today may hit another tomorrow, in ways that are impossible to foresee.

Credit markets and default: Firms have access to a competitive non-contingent debt market, where lenders offer bond price (or equivalently, interest rate) schedules as a function of

[7]This is a natural assumption: with a continuum of firms and a stationary shock process, firms can learn the complete distribution of any idiosyncratic shocks after one period.

Covid Economics 8, 22 April 2020: 1-40

Covid Economics 8, 22 April 2020: 1-40

aggregate and idiosyncratic states, in the spirit of Eaton and Gersovitz (1981). A firm enters period $t+1$ with an obligation, b_{it+1} to bondholders The shocks are then realized and the firm's shareholders decide whether to repay their obligations or default. Default is optimal for shareholders if and only if

$$\Pi_{it+1} - b_{it+1} + \Gamma_{t+1} < 0$$

where Γ_{t+1} is the present value of continued operations. Thus, the default decision is a function of the resources available to the firm Π_{it+1} (output plus undepreciated capital less wages) and the obligations to bondholders b_{it+1}. Let $r_{it+1} \in \{0, 1\}$ denote the default policy of the firm.

In the event of default, equity holders get nothing. The productive resources of a defaulting firm are sold to an identical new firm at a discounted price, equal to a fraction $\theta < 1$ of the value of the defaulting firm. The proceeds are distributed *pro-rata* among the bondholders.[8]

Let q_{it} denote the bond price schedule faced by firm i in period t, i.e. the firm receives q_{it} in exchange for a promise to pay one unit of output at date $t+1$. Debt is assumed to carry a tax advantage, which creates incentives for firms to borrow. A firm which issues debt at price q_{it} and promises to repay b_{it+1} in the following period, receives a date-t payment of $\chi q_{it} b_{it+1}$, where $\chi > 1$. This subsidy to debt issuance, along with the cost of default, introduces a trade-off in the firm's capital structure decision, breaking the Modigliani-Miller theorem.[9]

For a firm that does not default, the dividend payout is its total available resources, minus its payments to debt and labor, minus the cost of building next period's capital stock (the undepreciated current capital stock is included in Π_{it}), plus the proceeds from issuing new debt, including its tax subsidy

$$d_{it} = \Pi_{it} - b_{it} - \hat{k}_{it+1} + \chi q_{it} b_{it+1}. \tag{10}$$

Importantly, we do not restrict dividends to be positive, with negative dividends interpreted as (costless) equity issuance. Thus, firms are not financially constrained, ruling out another potential source of persistence.

Bankruptcy and obsolescence: Next, we spell out the relationship between default and capital obsolescence, $\phi_t = \Phi(\tilde{\phi}_t, Def_t)$ where $Def_t \equiv \int r_{it} di$. This is meant to capture the idea that widespread bankruptcies can amplify the erosion in the economic value of capital arising from the primitive shock $\tilde{\phi}_t$. This might come from lost supply chain linkages, inter-

[8]In our baseline specification, default does not destroy resources - the penalty is purely private. This is not crucial - it is straightforward to relax this assumption by assuming that all or part of the cost of the default represents physical destruction of resources.

[9]The subsidy is assumed to be paid by a government that finances it through a lump-sum tax on the representative household.

Covid Economics 8, 22 April 2020: 1-40

firm relationships or other ways in which economic activity is inter-connected. For example, a retailer might ascribe a lower value to space in a mall if a number of other stores go out of business. Similarly, a manufacturer might need to undertake costly search or make adjustments to his factory in order to accommodate new suppliers. We capture these effects with a flexible functional form:

$$\ln \phi_t \;=\; \ln \Phi(\tilde{\phi}_t, Def_t) \;=\; \ln \tilde{\phi}_t - \mu \; Def_t^{1-\varpi}, \tag{11}$$

where μ and ϖ are parameters that govern the relationship between default and the loss of capital value.

Timing and value functions:

1. Firms enter the period with a capital stock $\hat{k}_{it}$ and outstanding debt b_{it}.
2. The aggregate capital obsolescence shocks are realized.[10] Labor choice is made and production takes place.
3. Firm-specific shocks v_{it} are realized. The firm decides whether to default or repay ($r_{it} \in \{0,1\}$) its debt claims and distribute any remaining dividends.
4. The firm makes capital $\hat{k}_{it+1}$ and debt b_{it+1} choices for the following period.

In recursive form, the problem of the firm is

$$V\left(\hat{k}_{it}, b_{it}, v_{it}, \mathcal{S}_t\right) = \max\left[0, \max_{d_{it}, l_{it}, \hat{k}_{it+1}, b_{it+1}} d_{it} + \mathbb{E}_t M_{t+1} V\left(\hat{k}_{it+1}, b_{it+1}, v_{it+1}, \mathcal{S}_{t+1}\right),\right] \tag{12}$$

where M_{t+1} is the representative households's stochastic discount factor, subject to

$$\begin{aligned}
&\text{Dividends:} && d_{it} \;\leq\; \Pi_{it} - b_{it} - \hat{k}_{it+1} + \chi q_{it} b_{it+1} && (13)\\
&\text{Resources:} && \Pi_{it} \;=\; v_{it}\left[z_t k_{it}^{\alpha} l_{it}^{1-\alpha} - W_t l_{it} + (1-\delta) k_{it}\right] && (14)\\
&\text{Bond price:} && q_{it} \;=\; \mathbb{E}_t M_{t+1}\left[r_{it+1} + (1 - r_{it+1}) \frac{\theta \tilde{V}_{it+1}}{b_{it+1}}\right] && (15)
\end{aligned}$$

Finally, firms hire labor in a competitive market at a wage W_t. We assume that this decision is made after observing the aggregate shock but before the idiosyncratic shocks are observed,

[10] To simulate the COVID-19 pandemic, we run the epidemiology model from Section II.A for one year and use the predicted capital obsolescence as the realized shock for 2020. For more details, see Section III.

i.e. labor choice is solves the following static problem:

$$\max_{l_{it}} \quad z_t(\phi_t \hat{k}_{it})^{\alpha} l_{it}^{1-\alpha} - W_t l_{it}$$

The first max operator in (12) captures the firm's option to default. The expectation $\mathbb{E}_t$ is taken over the idiosyncratic and aggregate shocks, given beliefs about the aggregate shock distribution. The value of a defaulting firm is simply the value of a firm with no external obligations, i.e. $\tilde{V}_{it} = V\left(\hat{k}_{it}, 0, v_{it}, \mathcal{S}_t\right)$.

The aggregate state $\mathcal{S}_t$ consists of $(\hat{K}_t, \tilde{\phi}_t, \mathcal{I}_t)$ where $\mathcal{I}_t$ is the economy-wide information set. Equation (15) reveals that bond prices are a function of the firm's capital $\hat{k}_{it+1}$ and debt b_{it+1}, as well as the aggregate state $\mathcal{S}_t$. The firm takes the aggregate state and the function $q_{it} = q\left(\hat{k}_{it+1}, b_{it+1}, \mathcal{S}_t\right)$ as given, while recognizing that its choices affect its bond price.

Information, beliefs and equilibrium The set $\mathcal{I}_t$ includes the history of all shocks $\tilde{\phi}_t$ observed up to and including time-t. The expectation operator $\mathbb{E}_t$ is defined with respect to this information set. Expectations are probability-weighted integrals, where the probability density is $\hat{g}(\tilde{\phi})$. The function $\hat{g}$ arises from using the kernel density estimation procedure in equation (1).

For a given belief $\hat{g}$, a recursive equilibrium is a set of functions for (i) aggregate consumption and labor that maximize (8) subject to a budget constraint, (ii) firm value and policies that solve (12) , taking as given the bond price function (15) and the stochastic discount factor, (iii) aggregate consumption and labor are consistent with individual choices and (iv) capital obsolescence is consistent with default rates according to (11).

II.C Characterization

The equilibrium of the economic model is a solution to the following set of non-linear equations. First, the fact that the constraint on dividends (13) will bind at the optimum can be used to substitute for d_{it} in the firm's problem (12). This leaves us with 2 inter-temporal choice variables $(\hat{k}_{it+1}, b_{it+1})$ and a default decision. The latter is described by a threshold rule in the idiosyncratic output shock v_{it}:

$$r_{it} = \begin{cases} 0 & \text{if} \quad v_{it} < \underline{v}_t \\ 1 & \text{if} \quad v_{it} \geq \underline{v}_t \end{cases}$$

which implies that the default rate $Def_t = F(\underline{v}_t)$. It turns out to be more convenient to redefine variables and cast the problem as a choice of $\hat{k}_{it+1}$ and leverage, $lev_{it+1} \equiv \frac{b_{it+1}}{\hat{k}_{it+1}}$. The full characterization to the Appendix. Since all firms make symmetric choices for these objects,

Covid Economics 8, 22 April 2020: 1-40

in what follows, we suppress the i subscript. The optimality condition for $\hat{k}_{t+1}$ is:

$$1 = \mathbb{E}[M_{t+1}R^k_{t+1}] + (\chi - 1)lev_{t+1}q_t - (1-\theta)\mathbb{E}[M_{t+1}R^k_{t+1}h(\underline{v}_{t+1})] \tag{16}$$

$$\text{where} \qquad R^k_{t+1} = \frac{\phi^{\alpha+\nu}_{t+1}\hat{k}^{\alpha}_{t+1}l^{1-\alpha}_{t+1} - W_{t+1}l_{t+1} + (1-\delta)\,\phi_{t+1}\hat{k}_{t+1}}{\hat{k}_{t+1}} \tag{17}$$

The object R^k_{t+1} is the *ex-post* per-unit, post-wage return on capital, which is obviously a function of the obsolescence shock ϕ_t. The default threshold is given by $\underline{v}_{t+1} = \frac{lev_{t+1}}{R^k_{t+1}}$ while $h(\underline{v}) \equiv \int_{-\infty}^{\underline{v}} vf(v)dv$ is the default-weighted expected value of the idiosyncratic shock.

The first term on the right hand side of (16) is the usual expected direct return from investing, weighted by the stochastic discount factor. The other two terms are related to debt. The second term reflects the indirect benefit to investing arising from the tax advantage of debt - for each unit of capital, the firm raises $\frac{b_{t+1}}{\hat{k}_{t+1}}q_t$ from the bond market and earns a subsidy of $\chi - 1$ on the proceeds. The last term is the cost of this strategy - default-related losses, equal to a fraction $1-\theta$ of available resources.

Note that the default threshold is a function of ϕ_t, which in turn is affected by default, through (11). Therefore, the threshold equation $\underline{v}_{t+1} = \frac{lev_{t+1}}{R^k_{t+1}}$ implicitly defines a fixed-point relationship:

$$\underline{v}_{t+1} = \frac{lev_{t+1}}{R^k_{t+1}} = \frac{lev_{t+1}}{\phi^{\alpha+\nu}_{t+1}\hat{k}^{\alpha-1}_{t+1}l^{1-\alpha}_{t+1} - W_{t+1}\frac{l_{t+1}}{\hat{k}_{t+1}} + (1-\delta)\,\phi_{t+1}} \tag{18}$$

Next, the firm's optimal choice of leverage, lev_{t+1} is

$$(1-\theta)\,\mathbb{E}_t\left[M_{t+1}\frac{lev_{t+1}}{R^k_{t+1}}f\left(\frac{lev_{t+1}}{R^k_{t+1}}\right)\right] = \left(\frac{\chi-1}{\chi}\right)\mathbb{E}_t\left[M_{t+1}\left(1-F\left(\frac{lev_{t+1}}{R^k_{t+1}}\right)\right)\right]. \tag{19}$$

The left hand side is the marginal cost of increasing leverage - it raises the expected losses from the default penalty (a fraction $1-\theta$ of the firm's value). The right hand side is the marginal benefit - the tax advantage times the value of debt issued.

Finally, firm and household optimality implies that labor solves the intra-temporal condition:

$$\frac{(1-\alpha)y_t}{l_t} = W_t = \frac{1-\gamma}{\gamma}\frac{c_t}{1-l_t} \tag{20}$$

The optimality conditions, (16) - (20), along with those from the household side, form the system of equations we solve numerically.

Covid Economics 8, 22 April 2020: 1-40

Covid Economics 8, 22 April 2020: 1-40

III Measurement, Calibration and Solution Method

This section describes how we use macro data to estimate beliefs and parameterize the model, as well as our computational approach. A strength of our theory is that we can use observable data to estimate beliefs at each date.

Measuring past shocks Of course, we have not seen a health event like COVID in the last 95-100 years. However, from an economic point of view, COVID is one of many past shocks to returns that happens to be larger. When we think about COVID changing our beliefs, or our perceived probability distribution of outcomes, those outcomes are realized returns on capital. Therefore, to estimate the pre-COVID and post-COVID probability distributions, we first set out to measure past capital returns that map neatly into our model.

A helpful feature of capital obsolescence shocks, like the ones in our model, is that their mapping to available data is straightforward. A unit of capital installed in period $t-1$ (i.e. as part of $\hat{K}_t$) is, in effective terms, worth ϕ_t units of consumption goods in period t. Thus, the change in its market value from $t-1$ to t is simply ϕ_t.

We apply this measurement strategy to annual data on commercial capital held by US corporates. Specifically, we use two time series Non-residential assets from the Flow of Funds, one evaluated at market value and the second, at historical cost.[11] We denote the two series by NFA_t^{MV} and NFA_t^{HC} respectively. To see how these two series yield a time series for ϕ_t, note that, in line with the reasoning above, NFA_t^{MV} maps directly to effective capital in the model. Formally, letting P_t^k the nominal price of capital goods in t, we have $P_t^k K_t = NFA_t^{MV}$. Investment X_t can be recovered from the historical series, $P_{t-1}^k X_t = NFA_t^{HC} - (1-\delta)\, NFA_{t-1}^{HC}$. Combining, we can construct a series for $P_{t-1}^k \hat{K}_t$:

$$\begin{aligned} P_{t-1}^k \hat{K}_t &= (1-\delta) P_{t-1}^k K_{t-1} + P_{t-1}^k X_t \\ &= (1-\delta) NFA_{t-1}^{MV} + NFA_t^{HC} - (1-\delta)\, NFA_{t-1}^{HC} \end{aligned}$$

Finally, in order to obtain $\phi_t = \frac{K_t}{\hat{K}_t}$, we need to control for nominal price changes. To do this, we proxy changes in P_t^k using the price index for non-residential investment from the National

[11] These are series FL102010005 and FL102010115 from Flow of Funds.

Income and Product Accounts (denoted $PINDX_t$).[12] This yields:

$$\begin{aligned}\phi_t = \frac{K_t}{\hat{K}_t} &= \left(\frac{P_t^k K_t}{P_{t-1}^k \hat{K}_t}\right)\left(\frac{PINDX_{t-1}^k}{PINDX_t^k}\right) \\ &= \left[\frac{NFA_t^{MV}}{(1-\delta)NFA_{t-1}^{MV} + NFA_t^{HC} - (1-\delta)\, NFA_{t-1}^{HC}}\right]\left(\frac{PINDX_{t-1}^k}{PINDX_t^k}\right) \end{aligned} \tag{21}$$

Using the measurement equation (21), we construct an annual time series for capital depreciation shocks for the US economy since 1950. The mean and standard deviation of the series over the entire sample are 1 and 0.03 respectively. The autocorrelation is statistically insignificant at 0.15.

Next, we recover the primitive shock $\tilde{\phi}_t$ from the time series ϕ_t. To do this, we use (11), along with data on historical default rates from Moody's Investors Service (2015)[13] and values for the feedback parameters (μ, ϖ) as described below. The first panel of Figure 2 shows the estimated $\tilde{\phi}$.

Parameterization A period t is interpreted as a year. We choose the discount factor $\beta = 0.95$, depreciation $\delta = 0.06$, and the share of capital in the production, α, is 0.40. The recovery rate upon default, θ, is set to 0.70, following Gourio (2013). The distribution for the idiosyncratic shocks, v_{it} is assumed to be lognormal, i.e. $\ln v_{it} \sim N\left(-\frac{\hat{\sigma}^2}{2}, \hat{\sigma}^2\right)$ with $\hat{\sigma}^2$ chosen to target a default rate of 0.02.[14] The share of consumption in the period utility function, γ, is set to 0.4.

For the parameters governing risk aversion and intertemporal elasticity of substitution, we use standard values from the asset pricing literature and set $\psi = 0.5$ (or equivalently, an IES of 2) and $\eta = 10$. The tax advantage parameter χ is chosen to match a leverage target of 0.50, the ratio of external debt to capital in the US data – from Gourio (2013). Finally, we set the parameters of the default-obsolescence feedback function, namely μ and ϖ. Ideally, these parameters would be calibrated to match the variability of default and its covariance with the observed ϕ_t shock. Unfortunately, our one-shock model fails to generate enough volatility in default rates and therefore, struggles to match these moments. Fixing this would almost certainly require a richer model with multiple shocks and more involved financial frictions. We

[12] Our results are robust to alternative measures of nominal price changes, e.g. computed from the price index for GDP or Personal Consumption Expenditure.

[13] The Moody's data are for rated firms and shows a historical average default rate of 1% (our calibration implies a default rate of 2%), probably reflecting selection. Accordingly, we scaled the Moody's estimates by a factor of 2 while performing this calculation. We also used estimates of exit and bankruptcy rates from Corbae and D'Erasmo (2017) and found broadly similar results.

[14] This is in line with the target in Khan et al. (2017), though a bit higher than the one in Gourio (2013). We verified that our quantitative results are not sensitive to this target.

Covid Economics 8, 22 April 2020: 1-40

take a simpler way out here and target a relatively modest feedback with values of $\mu = 0.2$ and $\varpi = 0.5$. These values imply roughly an amplification 3% at a baseline default rate of 2%, rising to 5% for a 6% default.[15]

Epidemiology parameters. A major hurdle to quantifying the long-run effects is the lack of data and uncertainty surrounding estimates of the short-run impact. While this will surely be sorted out in the months to come, for now, with the crisis still raging and policy still being set, the impact is uncertain. More importantly for us, the nature of the economic shock is uncertain. It may be a temporary closure with furloughs, or it could involve widespread bankruptcies and changes in habits that permanently separate workers from capital or make the existing stock of capital ill-suited to the new consumption demands. Since it is too early to know this, we present two possible scenarios, chosen to illustrate the interaction between learning and the type of shock we experience. All involve significant losses in the short term but their long-term effects on the economy are drastically different.

We begin by describing parameter choices that are fixed across the scenarios. Following Wang et al. (2020)'s study of infection in Hubei, China, we calibrate $\sigma_E = 1/5.2$ and $\gamma_I = 1/18$ to the average duration of exposure (5.2 days) and infection (18 days). We use an initial reproduction number of $R_0 = 3.5$, based on more recent estimates of higher antibody prevalence and more asymptomatic infection than originally thought and $R_{min} = 0.8$ based on the estimates of the spread in New York, at the peak of the lockdown (Center for Disease Control, 2020). This implies that the initial number of contacts per period must be $\tilde{\beta} = \gamma_I R_0$.

The extent to which capital idling reduces contact rates is set to $\tilde{\theta} = 1/3$. This implies that a lockdown which reduces the reproduction number to 0.8 is associated with 50% capital idling. This is broadly consistent with the 25% drop in output, estimated during the lockdown period in Hubei province, China. The rate of excess depreciation of idle capital at the rate of 6.5% per month or $\tilde{\delta} = 0.065/30$ daily. As we will see, this implies a 10% erosion of the value of capital in our first scenario, which lines up with the drop in commercial real estate prices since the pandemic started – see CPPI (2020).

The two scenarios, which differ in the sensitivity of lockdown policy to observed infection increases, i.e. the parameter ζ_I. In scenario 1, we set $\zeta_I = 300$, which generates an initial lockdown that lasts for 2 months. This version of the model predicts waves of re-infection and new lockdowns in the months to come, echoing predictions by the Center for Disease Control. Scenario 2, which considers a much less aggressive response by setting $\zeta_I = 50$, has only one lockdown episode.

[15] Section VI studies a version without default amplification and finds that it generates similar patterns, albeit with slightly smaller magnitudes, as our benchmark economy.

Covid Economics 8, 22 April 2020: 1-40

Table 1 summarizes the resulting parameter choices.

Parameter	Value	Description
Preferences:		
β	0.95	Discount factor
η	10	Risk aversion
ψ	0.50	1/Intertemporal elasticity of substitution
γ	0.40	Share of consumption in the period utility function
Technology:		
α	0.40	Capital share
δ	0.06	Depreciation rate
$\hat{\sigma}$	0.28	Idiosyncratic volatility
Debt:		
χ	1.06	Tax advantage of debt
θ	0.70	Recovery rate
μ	0.2	Default-obsolescence feedback
ϖ	0.5	Default-obsolescence elasticity
Disease / Policy:		
R_0	3.5	Initial disease reproduction rate
R_{min}	0.8	Minimum U.S. disease reproduction rate
σ_E	1/52	Exposure to infection transition rate
γ_I	1/18	Recovery / death rate
ζ_I	300 (50)	Lockdown policy sensitive to past infections
$\tilde{\theta}$	0.19	Capital idling required to reduce transmission
$\tilde{\delta}$	0.002	Excess depreciation (daily) of idle capital

Table 1: **Parameters** Number in parentheses is used in scenario 2.

Numerical solution method Since curvature in policy functions is an important feature of the economic environment, our algorithm solves equations (20) – (19) with a non-linear collocation method. Appendix A.B describes the iterative procedure. In order to keep the computation tractable, we need one more approximation. The reason is that date-t decisions (policy functions) depend on the current estimated distribution ($\hat{g}_t(\tilde{\phi})$) and the probability distribution h over next-period estimates, $\hat{g}_{t+1}(\tilde{\phi})$. Keeping track of $h(\hat{g}_{t+1}(\tilde{\phi}))$, (a compound lottery) makes a function a state variable, which renders the analysis intractable. However, the approximate martingale property of $\hat{g}_t$ discussed in Section I offers an accurate and computationally efficient approximation to this problem. The martingale property implies that the average of the compound lottery is $E_t[\hat{g}_{t+1}(\tilde{\phi})] \approx \hat{g}_t(\tilde{\phi})$, $\forall\tilde{\phi}$. Therefore, when computing policy functions, we approximate the compound distribution $h(\hat{g}_{t+1}(\tilde{\phi}))$ with the simple lottery $\hat{g}_t(\tilde{\phi})$, which is today's estimate of the probability distribution. We use a numerical experiment to show that this approximation is quite accurate. The reason for the small approximation error

Covid Economics 8, 22 April 2020: 1-40

is that $h(\hat{g}_{t+1})$ results in distributions centered around $\hat{g}_t(\tilde{\phi})$, with a small standard deviation. Even 30 periods out, $\hat{g}_{t+30}(\tilde{\phi})$ is still quite close to its mean $\hat{g}_t(\tilde{\phi})$. For 1-10 years ahead, where most of the utility weight is, this standard error is tiny.

To compute our benchmark results, we begin by estimating $\hat{g}_{2019}$ using the data on $\tilde{\phi}_t$ described above. Given this $\hat{g}_{2019}$, we compute the stochastic steady state by simulating the model for 5000 periods, discarding the first 500 observations and time-averaging across the remaining periods. This steady state forms the starting point for our results. Subsequent results are in log deviations from this steady state level. Then, we subject the model economy to two possible additional adverse realizations for 2020, one at a time. Using the one additional data point for each scenario, we re-estimate the distribution, to get $\hat{g}_{2020}$. To see how persistent economic responses are, we need a long future time series. We don't know what distribution future shocks will be drawn from. Given all the data available to us, our best estimate is also $\hat{g}_{2020}$. Therefore, we simulate future paths by drawing many sequences of future $\tilde{\phi}$ shocks from the $\hat{g}_{2020}$ distribution. In the results that follow, we plot the mean future path of various aggregate variables.

IV Main Results

Our goal in this paper is to quantify the long run effect of the COVID crises, stemming from the belief scarring effect, i.e. from learning that pandemics are more likely than we thought. We formalize and quantify the effect on beliefs, using the assumption that people do not know the true distribution of aggregate economic shocks and learn about it statistically. This is the source of the long-run economic effects. Comparing the resulting outcomes to ones from the same model under the assumption of full knowledge of the distribution (no learning) reveals the extent to which beliefs matter.

But first, we briefly describe the disease spread, the policy reaction and the economic shocks these policies generate.

Epidemiology and economic shutdown. Figure 1 illustrates the spread of disease, in both scenarios, as well as the response, which results in capital idling. Recall that Scenario 2 has $\zeta_I = 50$, i.e. a policy that is six times less responsive to changes in the infection rate than the $\zeta_I = 300$ policy in scenario 1. As a result, it also has significantly less idle capital and a faster spike in infection rates.

For our purposes, the sufficient statistic in each scenario is the realization for $\tilde{\phi}_{2020}$. In scenario 1, the COVID-19 shock implies $\tilde{\phi}_t = 0.9$, i.e. the loss of value due to obsolescence is equal to 10% of the capital stock. In scenario 2, only 5% of capital is lost to obsolescence:

CEPR PRESS

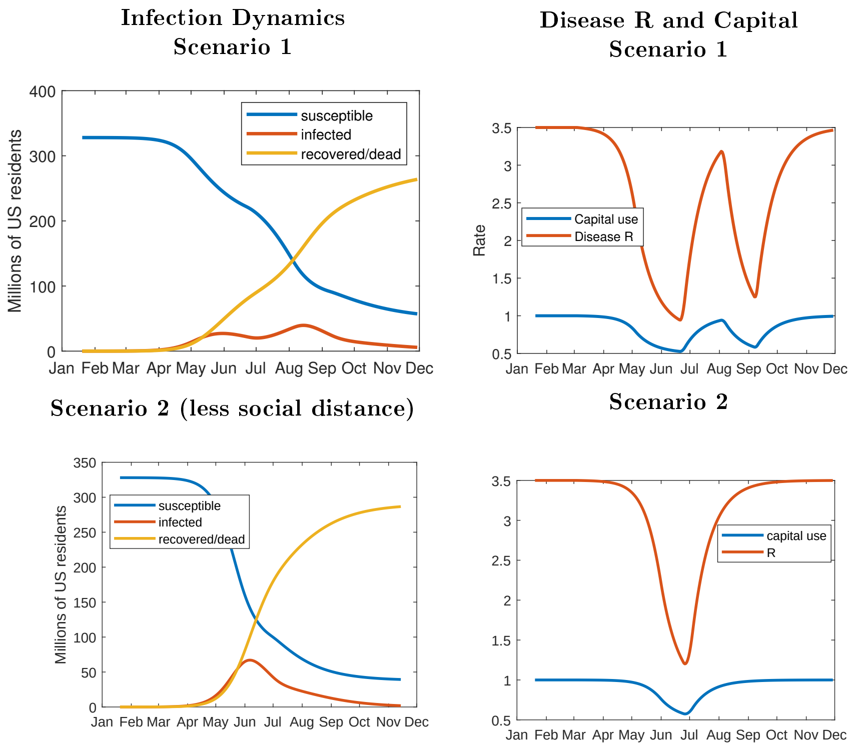

Figure 1: **Disease spread and capital dynamics.**
Parameters listed in Table 1. Scenario 1 uses an aggressive lockdown policy $\zeta_I = 300$*, while scenario 2 uses a more relaxed policy of* $\zeta_I = 50$.

$\tilde{\phi}_t = 0.95$. The target for the initial, transitory impact is line with most forecasts for 2020: a 9% (or 6%) annual decline in GDP. This is likely a conservative estimate for Q2 2020, but more extreme than some forecasts for the entire year.

How much belief scarring? We apply our kernel density estimation procedure to the capital return time series and our two scenarios to construct a sequence of beliefs. In other words, for each t, we construct $\{\hat{g}_t\}$ using the available time series until that point. The resulting estimates for 2019 and 2020 are shown in Figure 2. The differences are subtle. Spotting them requires close inspection where the dotted and solid lines diverge, around 0.90 and 0.95, in scenarios 1, and 2 respectively. They show that the COVID-19 pandemic induces an increase in the perceived likelihood of extreme negative shocks. In scenario 1, the estimated density for 2019 implies near zero (less than $10^{-5}\%$) chance of a $\tilde{\phi} = 0.90$ shock; the 2020 density attaches

Covid Economics 8, 22 April 2020: 1-40

Covid Economics 8, 22 April 2020: 1-40

a 1-in-70 or 1.4% probability to a similar event recurring.

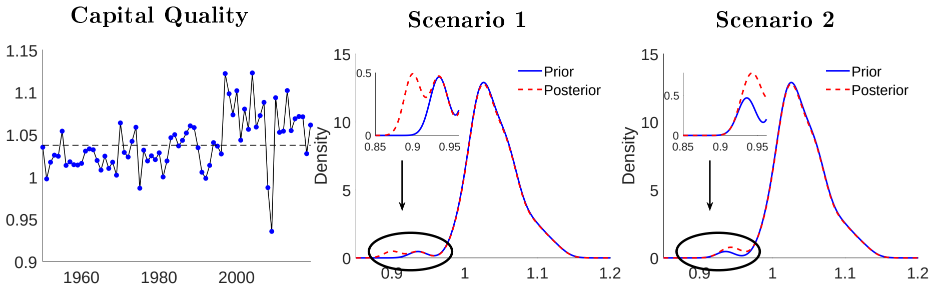

Figure 2: **Beliefs about the probability distribution of outcomes, plotted before and during the COVID-19 crisis.**
The first panel shows the realizations of $\tilde{\phi}$. The second and third panels show the estimated kernel densities for 2019 (solid line) and 2020 (dashed line) for the two scenarios. The subtle changes in the left tail represent the scarring effect of COVID-19.

As the graph shows, for most of the sample period, the shock realizations are in a relatively tight range around 1, but we saw a large adverse realizations during the Great Recession of 0.93 in 2009. This reflects the large drops in the market value of non-residential capital stock. The COVID shock is now a second extreme realization of negative capital returns in the last 20 years. It makes such an event appear much more likely.

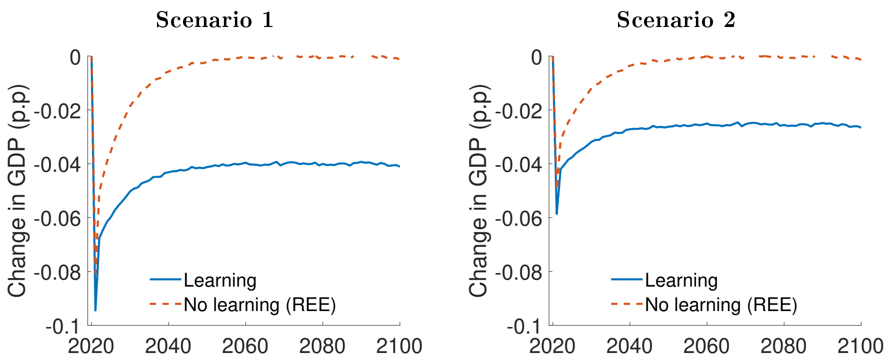

Figure 3: **Output with scarring of beliefs (solid line) and without (dashed line).**
Units are percentage changes, relative to the pre-crisis steady-state, with 0 being equal to steady state and -0.1 *meaning 10% below steady state. Common parameters listed in Table 1. Scenario-specific parameters are: Scenario 1:* $\tilde{\phi}_{2020} = 0.90$ *Scenario 2:* $\tilde{\phi}_{2020} = 0.95$.

Effect on GDP Observing a tail event like the COVID-19 pandemic changes output in a persistent way. Figure 3 compares the predictions of our model for total output (GDP) to an identical model without learning. The units are log changes, relative to the pre-crisis steady-state. In the model without learning, agents are assumed to know the true probability of pandemics. As a result, when they see the COVID crisis, they do not update the distribution. This corresponds to the canonical "rational expectations" assumption in macroeconomics. The model with learning, which uses our real-time kernel density estimation to inform beliefs, generates similar short-term reactions, but worse long-term effects. The post-2020 paths are simulated as follows: each economy is assumed to be at its stochastic steady state in 2019 and is subjected to the same 2020 $\tilde{\phi}$ shock; subsequently, sequences of shocks drawn from the estimated 2020 distribution.

The scenarios under learning correspond to what one might call a V-shaped or tilted-V recession: the recovery after the shock has passed is significant but not complete. Note that the drop in GDP on impact is a calibration target – what we are interested in its persistence, which arguably matters more for welfare. The graph shows that, in Scenario 1, learning induces a long-run drop in GDP of about 4%. The right panel shows a similar pattern but the magnitudes are smaller. Of course, agents also learn from smaller capital obsolescence shocks. These also scar their beliefs going forward. But the scarring is much less, producing only a 3% loss in long-run annual output.

Higher tail risk (i.e. greater likelihood of obsolescence going forwards) increases the risk premium required on capital investments, leading to lower capital accumulation. It is important that these shocks make capital obsolete, rather than just reduce productivity, because obsolescence has a much bigger effect on capital returns than lower productivity does. Labor also contracts, but that is a reaction to the loss of available capital that can be paired with labor. When a chunk of capital becomes mal-adapted and worthless, that is an order of magnitude more costly to the investor than the temporary decline in capital productivity. Since most of the economic effect works through capital risk deterring investment, that lower return is important to get the economic magnitudes right.

Turning off belief updating When agents do not learn, both scenarios exhibit quick and complete recoveries, even with a large initial impact. Without the scarring of beliefs, facilities are re-fitted, workers find new jobs, and while the transition is painful, the economy returns to its pre-crisis trajectory relatively quickly. In other words, without belief revisions, the negative shock leads to an investment boom, as the economy replenishes the lost effective capital. While the curvature in utility moderates the speed of this transition to an extent, the overall pattern of a steady recovery back to the original steady state is clear. This is in sharp contrast to the version with learning. Note that since the no-learning economy is endowed with the same end-of-sample beliefs as the learning model, they both ultimately converge to the same *levels*.

Covid Economics 8, 22 April 2020: 1-40

CEPR PRESS

But, they start at different steady states (normalized to 0 for each series). This shows that learning is what generates long-lived reductions in economic activity.

Decomposing long-run losses. Next, we perform a simple calculation to put the size of the long-run loss in perspective. Specifically, we use the stochastic discount factor implied by the model to calculate the expected discounted value of the reduction in GDP. These estimates, reported in Table 2, imply that the representative agent in this economy values the cumulative losses between 57% and 90% of the pre-COVID GDP. Most of this comes from the belief scarring mechanism.

Scenario	2020 GDP drop	NPV(Belief Scarring)	NPV(Obsolete capital)
I. Tough	-9%	-52%	-38%
II. Light	- 6%	-33%	-24%

Table 2: New present value costs in percentages of 2019 GDP.

Note that the 1-year loss during the pandemic is 6-9% of GDP. The cost of belief scarring is five to six times as large, in both cases. The cost of obsolete capital is about four times as large as the damage done during the pandemic. Figure 4 illustrates the losses each year from the capital obsolescence and belief changes. The area of each of these regions, discounted as one moves to the right in time, is the NPV calculation in the table above. The one-year cost is a tiny fraction of this total area.

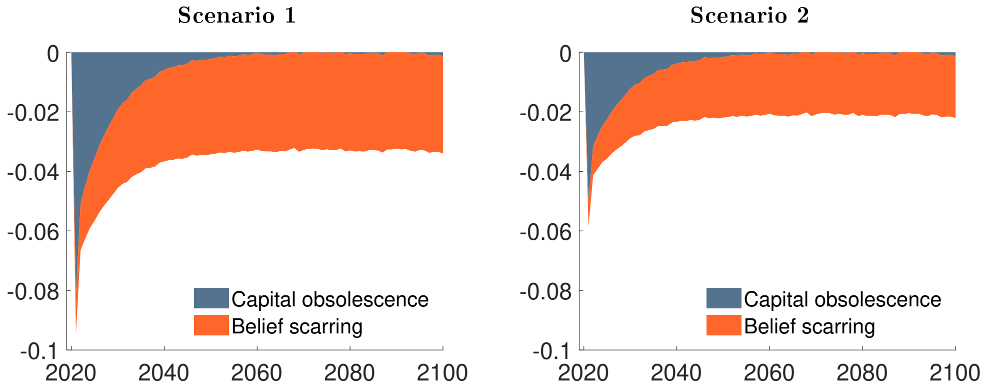

Figure 4: Long-term costs of the pandemic.

Of course, that calculation misses an important aspect of what we've learned – that pandemics will recur. Since our agents have 70 years of data, during which they've seen one pandemic, they assess the future risk of pandemics to be 1-in-70 initially. That probability

Covid Economics 8, 22 April 2020: 1-40

CEPR PRESS

Covid Economics 8, 22 April 2020: 1-40

declines slowly as time goes on and other pandemics are not observed. But there is also the risk there will be more pandemics, like this. This is not really a result of this pandemic. But that risk of future pandemics is what we should consider if we think about the benefits of public health investments. The pandemic cost going forward, in a world where a pandemic has a 1/70th probability of occurring each year, is given in Figure 5. Note that the risk of future pandemics costs the economy about 7-12% of GDP. This is similar to the one-year cost during the COVID crisis.

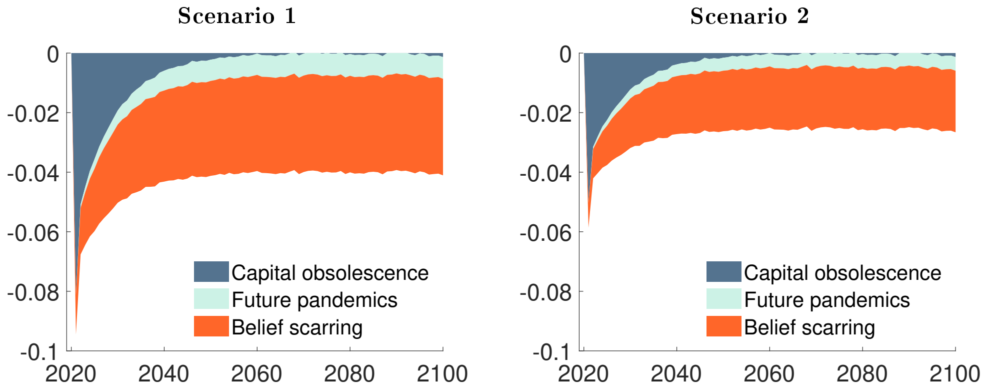

Figure 5: Long-term costs of with future pandemics.

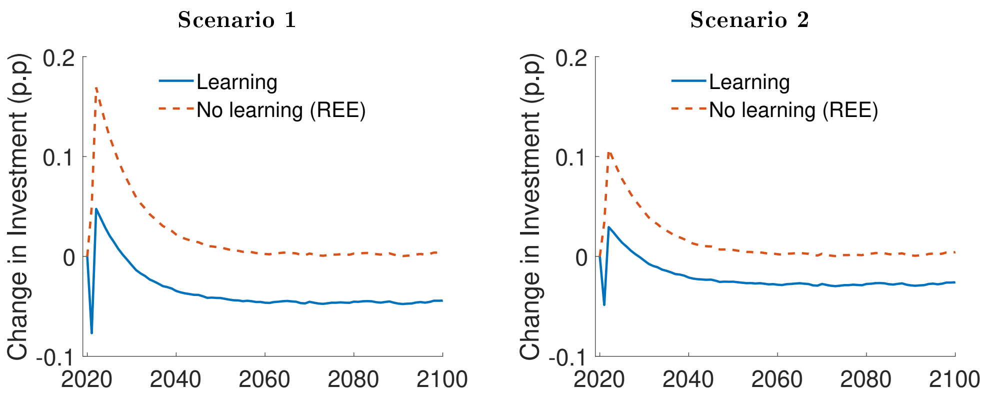

Figure 6: **Without belief scarring, investment surges.**
Results show average aggregate investment, with scarring of beliefs (solid line) and without (dashed line). Common parameters listed in Table 1. Scenario-specific parameters are: Scenario 1: $\tilde{\phi}_{2020} = 0.90$ *Scenario 2:* $\tilde{\phi}_{2020} = 0.95$.

Investment and Labor. Figure 6 shows the effect of belief changes on investment. When agents do not learn, investment surges immediately (as the economy replenishes the obsolete capital). With learning, investment shows a much smaller surge (starting in 2021), but eventually falls below the pre-COVID levels.

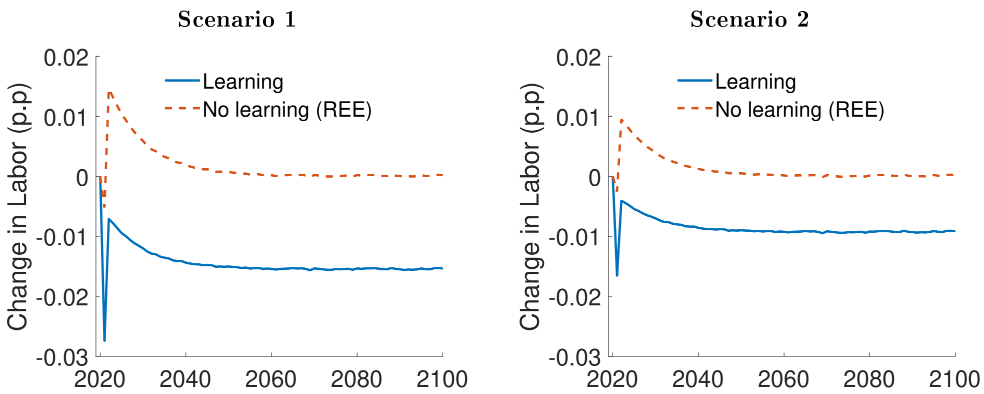

Figure 7: **Labor with scarring of beliefs (solid line) and without (dashed line).** *Common parameters listed in Table 1. Scenario-specific parameters are: Scenario 1:* $\tilde{\phi}_{2020} = 0.90$ *Scenario 2:* $\tilde{\phi}_{2020} = 0.95$.

In figure 7, we see that the initial reaction of labor is milder than for investment, but the bigger differences arise from 2021 onwards. When the transitory shock passes, investment surges, to higher than its initial level, to compensate for the obsolescence shock. But labor remains below the pre-COVID levels, reflecting the effect of the scarring effect on the stock of capital and through that on the demand for labor.

Defaults, Riskless Rates and Credit Spreads. The scenarios differ in their short-term implications for default as well. Default spikes only in 2020, the period of impact, returning to average from 2021 onwards. But, the higher default rate in scenario 1 (6% relative to 4% in Scenario 2) contributes to greater scarring (since default amplifies obsolescence). This result suggests a role for policy: preventing default/bankruptcy can lead to long-lasting benefits. In Section VI, we present a quantitative analysis of such a policy.

Nearly immediately, after the pandemic passes, default rates in both scenarios return to their original level. While defaults leave permanent scars on beliefs, the defaults themselves are not permanently higher. It is the memory of a transitory event that is persistent.

Because defaults were elevated, the pandemic had a large, immediate impact on credit spreads. However, these high spreads were quickly reversed. Some authors have argued that heightened tail risk should inflate risk premia, as well as credit spreads (Hall, 2016). While

Covid Economics 8, 22 April 2020: 1-40

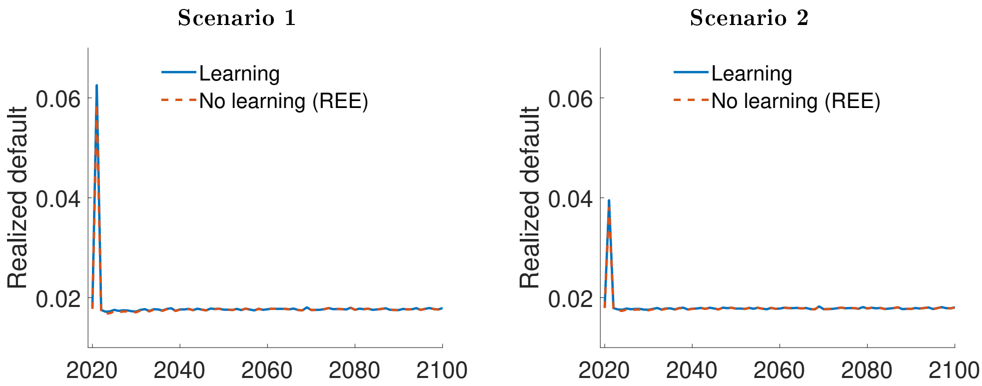

Figure 8: **Realized default does not respond much to beliefs.**
Results show with scarring of beliefs (solid line) and without (dashed line), often with the two lines on top of each other. Common parameters listed in Table 1. Scenario-specific parameters are: Scenario 1: $\tilde{\phi}_{2020} = 0.90$ *Scenario 2:* $\tilde{\phi}_{2020} = 0.95$

the argument is intuitive, it ignores any endogenous response of discounting, investment or borrowing. A surge in risk triggers disinvestment and de-leveraging. Because firms borrow less, this lowers default rates back down, which offsets the increase in the credit spread. We can see this channel at work in the drop in debt and the lack of change in long-run defaults (Table 3). The credit spread is the implied interest on risky debt, $1/q_t$ less the risk-free rate r^f. The credit spread in the stochastic steady state under the 2019 belief is less than a basis point higher is the post-pandemic long-run. Thus, belief revisions can have significant and long-lived real effects, even if the long-run change in credit spreads is very small.

Equity markets and implied skewness. One might think the recent recovery in equity prices appears inconsistent with a persistent rise in tail risk. The model teaches us why this logic is incomplete. When firms face higher tail risk, they also reduce debt, which pushes in the opposite direction as the rise in risk. Furthermore, when firms reduce investment and capital stocks decline, the marginal value of capital rises. Finally, when interest rates and thus future discount rates decline, future equity payments are worth more in present value terms. These competing effects cancel each other out. In our model, the market value of a dividend claim associated with a unit of capital is nearly identical under the post-COVID beliefs than under the pre-COVID ones. In other words, the combined effect of the changes in tail risk and debt reduction is actually mildly positive. While the magnitudes are not directly interpretable, our point is simply that rising equity valuations are not evidence against tail risk.

If credit spreads and equity premia are not clear indicators of tail risk, what is? For that, we

Covid Economics 8, 22 April 2020: 1-40

	2019	Scenario 1		Scenario 2	
	Baseline	level	change	level	change
Credit Spreads	0.837%	0.842%	+0.5 bps	0.838%	+0.1 bps
Debt	2.75	2.56	-0.19	2.63	-0.12
Default	2.0%	2.0%	0	2.0%	0
Risk free rate	3.66%	3.46%	-20 bps	3.54%	-11 bps
Equity market value	0.44	0.45	+0.01	0.44	0
SKEW	102.7	111.3	+8.6	104.2	+1.5
Third moment $E\left[\left(R^e - \bar{R}^e\right)^3\right]$	-1.5	-9.8	-8.3	-2.6	-1.1

Table 3: **Changes in financial market variables: Baseline, Scenarios 1 and 2.** *Baseline is the steady state pre-pandemic, under 2019 beliefs. Columns labelled "change" are the raw difference between the long-run average values under 2019 and 2020 beliefs in each scenario. They do not capture any changes that take place along the transition path or during the pandemic. The aggregate market capitalization in our model is the value of the dividend claim times the aggregate capital stock. Third moment is* $E\left[\left(R^e - \bar{R}^e\right)^3\right] \times 10^4$, *where* R^e *is the return on equity. The expectation is taken under the risk-neutral measure. For the no-learning model, all changes are zero.*

need to turn option prices, in particular out-of-the-money put options on the S&P 500, which can be used to isolate changes in perceived tail risk. A natural metric is the third moment of the distribution of equity returns. The last row of Table 3 reports this object (computed under the risk neutral measure). It shows that the perceived distribution after the shock is more negatively skewed.[16]

This might sound inconsistent with the behavior of the SKEW index reported by the CBOE. This showed a short-lived spike at the onset of the pandemic, but recovered quite rapidly. To understand this pattern, note that the SKEW indexes the *standardized* third moment implied by options prices, which is obtained by dividing the third moment from the previous paragraph by the implied standard deviation (or VIX). Tail events typically lead to a spike in market volatility, both realized and implied. This increase in VIX tends to mechanically lower the skewness. More generally, the SKEW index confounds changes in the third moment with the changes in the second moment, which often reflects many other factors. This is the main reason why we focus on the (non-standardized) third moment. As we saw, this measure clearly reveals the persistent change in beliefs and is consistent with evidence from newspapers and surveys in Barrero and Bloom (2020).

[16] It is straightforward to compute this from the SKEW and VIX indices reported by the CBOE. The third central moment under the risk-neutral measure is $E\left(R^e - \bar{R}^e\right)^3 = \frac{100-SKEW_t}{10} \cdot VIX_t^3$. This calculation reveals that, between February and May 2020, the market implied third moment also became significantly more negative (from -0.04 to -0.09).

Covid Economics 8, 22 April 2020: 1-40

Covid Economics 8, 22 April 2020: 1-40

V Liquidity and Interest Rates

In this section, we augment the baseline model to include a liquidity friction. This is motivated by evidence showing liquidity becoming more scarce following the onset of the pandemic – see Boyarchenko et al. (2020). As we will show, a liquidity motive amplifies the effects of tail risk on rates of return for liquid assets, such as Treasuries. This helps bring this dimension of the model's predictions closer to the observed drops in recent months. We also present evidence from bond markets consistent with the rise in liquidity premia.

Recall that, in the baseline model, riskless rates fall in response to higher demand for safe assets. Just as firms react to the increased tail risk by de-leveraging, investors would like to protect themselves against low-return states by holding more riskless assets. They cannot all hold more. Therefore, the price increases (the rate of return falls) to clear the market. Table 3 reports the riskless rate falls by 20 bps (10 bps) in Scenario 1 (Scenario 2). The sign of this change is consistent with what we saw following the onset of the pandemic, but the magnitude is not: interest rates, especially in the Treasury market, fell much more dramatically.

We introduce liquidity considerations using a stylized yet tractable specification, in the spirit of Lagos and Wright (2005).[17] A positive NPV investment opportunity requires liquid funds. Both capital and government bonds provide liquidity (the former only partially). An adverse capital obsolescence shock reduces the value of capital and thus the amount of liquidity it provides. Thus, an increase in the risk of such a shock makes capital liquidity uncertain and raises the value of riskless bonds, which are always retain their full, liquid value. Thus, higher tail risk also raises liquidity risk and makes riskless bonds, which serve as liquidity insurance, even more attractive. This channel amplifies the effect on their return and turns out to be quantitatively very large. The increased tail risk brought on by the pandemic, combined with liquidity risk, will turn out to depress interest rates three and a half times as much as in the model without liquidity risk.

Formally, firms are assumed to have access to a profitable intra-period opportunity, yielding a net return of $H(x_t) - x_t$ where x_t is the amount invested. The net return is maximized at $x_t = x^*$. But, the firm faces a liquidity constraint: x_t cannot exceed the amount of pledgable collateral. Formally,

$$x_t \leq a_t + \bar{d}k_t \tag{22}$$

where the parameter $\bar{d}$ indexes the pledgability of capital and a_t denotes a riskless, fully liquid asset. This can be interpreted narrowly as government bonds[18], but it could also be thought of

[17] See also Kozlowski et al. (2019).

[18] For concreteness, we adopt this assumption in our analysis. The bonds are issued by a government, which

Covid Economics 8, 22 April 2020: 1-40

as the total liquidity available from other sources. Note that the liquidity value of capital is a function of *effective* capital, i.e. net of obsolescence. As a result, shocks to capital obsolescence influence the availability of liquidity.

The supply of the liquid asset is assumed to be fixed at $\bar{a}$. Thus, the amount invested in the opportunity in t is given by $x_t = \min(x^*, \bar{a} + \bar{d}k_t)$. The liquidity premium is the marginal value (in units of consumption) of an additional unit of pledgable collateral:

$$\mu_t = H'(x_t) - 1 \; . \tag{23}$$

The return on government bonds, i.e. the liquid asset, is characterized

$$\frac{1}{R_t^a} = \mathbb{E}_t \left[M_{t+1}(1 + \mu_{t+1}) \right] \; . \tag{24}$$

The final model alternation is that the liquidity premium shows up in the first term of the optimality condition for capital (16), which becomes $\mathbb{E}_t \left[M_{t+1}(R_{t+1}^k + \bar{d}\mu_{t+1}) \right]$.

Parameterization. To set values for the liquidity parameters, we follow the strategy in Kozlowski et al. (2019). We use the following functional form for the benefit to invest on liquid assets: $H(x) = 2\iota\sqrt{x} - \xi$. The parameter that governs how much of capital is a pledgeable, liquid asset, $\bar{d}$, is set to 0.16 to match the ratio of short-term obligations of US nonfinancial corporations to the capital stock in the Flow of Funds. The liquid asset supply $\bar{a} = 0.8$ and the return parameter $\iota = 1.4$ are chosen so that the ratio of liquid assets to capital is 0.08 and their return in the pre-COVID steady state equals 2%. Finally, the parameter $\xi = 1.94$ is set so the net return of the project is close to zero (on average) in the pre-COVID steady state.

Riskless rates with liquidity premia. The purpose of this extension was to explore how liquidity considerations affect scarring-induced changes in riskless rates. To evaluate this, we compute riskless rates in the stochastic steady states associated with the pre- and post-COVID beliefs. These are presented in Table 4. The model predicts that the yield on liquid bonds drops by 67 bps in the new steady state. In contrast, the return on a riskless but completely illiquid asset falls only by 8 bps: in other words, the liquidity premium rises by 59 bps.

The table also shows changes in various market interest rates between January and July 2020. The yield on the 1y and 5y Treasuries were almost 1.4% lower in July 2020 (relative to the beginning of the year). Note that these are not directly comparable to the model numbers. The latter compare steady-states and so are most appropriately thought of as long-run predictions while the current data obviously reflect short-term, more transitory considerations as well. We

balances its budget with lumpsum taxes/transfers.

	Pre-COVID	Post-COVID	Chg
Model			
Riskless rate (liquid) $R_t^a - 1$	2.12%	1.46 %	-67 bps
Riskless rate (illiquid)	4.97%	4.89 %	-8 bps
Data			
1y Treasury yield (nominal)	1.56%	0.14%	-142 bps
5y Treasury yield (nominal)	1.67%	0.28%	-139 bps
5y forward rate (real)	-0.09%	-0.98%	-89 bps
AAA Yield	2.53%	1.48%	-105 bps
AAA Spread (rel. to 5y Treasury)	0.86%	1.20%	34 bps

Table 4: Implications for interest rates with liquidity frictions, model vs data. Data comes from FRED. Pre-COVID (post-COVID) data are for January 1, 2020 (July 16, 2020).

Covid Economics 8, 22 April 2020: 1-40

therefore construct a proxy for the long-run rates using forward rates implied by the Treasury yield and long-term inflation expectations. Specifically, we use the instantaneous rate 5 years forward from the Treasury yield curve and 5y5y inflation expectations[19] to calculate the change in long-term real rates. This shows a decline of about 89 bps, smaller than short-term rates and closer to the model's predictions.

Next, the table also reports the change in the yield on AAA corporate bonds. These securities carry very little default risk, but are not as liquid as Treasuries. As a result, the yield spread on these bonds relative to Treasuries is often viewed as a proxy for liquidity premia – see, e.g., Krishnamurthy and Vissing-Jorgensen (2012) and del Negro et al. (2017). In recent months, this spread rose by 34 bps, consistent with increased liquidity scarcity. The model liquidity premium reported in the table shows a larger rise. This is to be expected since the model object is defined as the spread of a *completely* illiquid security whereas AAA bonds are probably partially liquid.

Figure 9 shows the time path of the natural rate of interest. Notice that the short-run fluctuations are much larger than the long-term effects reported in the table. This is consistent with short-term market disruptions that are now settling down.

Finally, in interpreting recent data, it is worth pointing out that the last few months have seen unprecedented policy interventions in bond markets, which almost certainly have contributed to the drop in interest rates on both liquid and illiquid assets – see Boyarchenko et al. (2020). Our analysis completely abstracts from such interventions[20] so it is perhaps not too surprising that the model under-predicts the fall in interest rates. Overall, these results suggest

[19] 5y5y inflation expectations are the expectations of inflation over the five-year period, starting five years from today. Source: FRED, Federal Reserve Bank of St. Louis. The series tickers are THREEFF5 and T5YIFR respectively.

[20] We do evaluate the effects of a financial policy in the following section.

CEPR PRESS

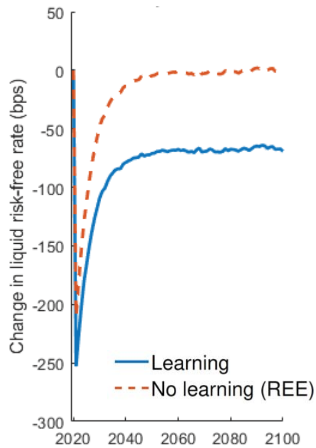

Figure 9: **Belief scarring lowers riskless rate in the long-run.**
Results show the return on a riskless asset, in scenario 1, with scarring of beliefs (solid line) and without (dashed line). Common parameters listed in Table 1. Scenario-specific parameters are– Scenario 1: $\tilde{\phi}_{2020} = 0.90$.

a quantitatively meaningful role for the belief scarring mechanism in the recent behavior of interest rates.

VI The Role of Financial Policy

The COVID-19 pandemic has sparked an unprecedented policy response. These responses fall into three broad categories: social distancing and other mobility restrictions, transfers to households and financial assistance to firms. We explored the consequences of more lax social distance policy in constructing scenarios for our baseline results. Transfers to households has an important role to mitigate the economic fallout, but does not directly affect productive capacity, the key object in our analysis. Financial assistance to firms, on the other hand, can help the economy maintain productive capacity, for example by preventing widespread bankruptcies. In our setting, such a policy would have beneficial long-run effects as well, since they mitigate the consequences of belief scarring. In this section, we use our baseline model to quantify these long-run benefits. We find that the longer-term effects of a policy of debt relief are as much as 10 times larger than the short-run effects.

The need for policy intervention in the model stems from the presence of debt and the associated risk of bankruptcy. Bankruptcy is socially costly because it exacerbates capital obsolescence. Therefore, we model financial policy as designed to prevent/limit bankruptcies by reducing firms' effective debt. This could take the form of the government or other policy-

Covid Economics 8, 22 April 2020: 1-40

maker buying up the debt from private creditors or offering direct assistance to firms. Before examining the effects of such a policy, we perform a simple exercise to quantify the costs of bankruptcy in our baseline model. This is the cost that financial policy might plausibly remedy.

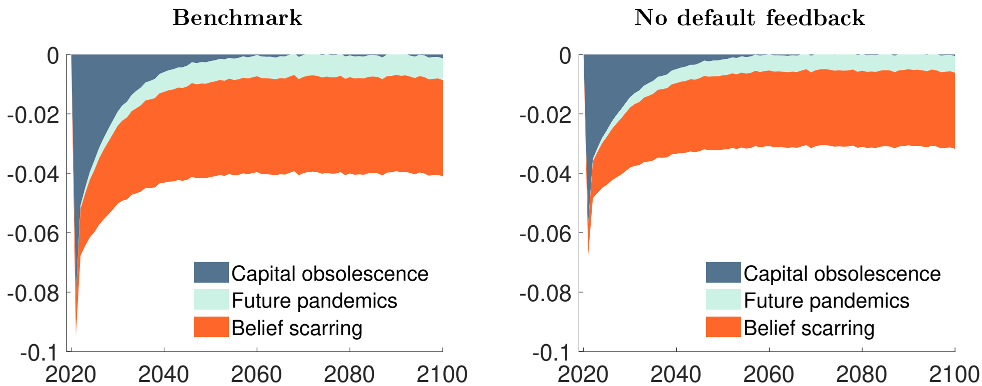

Figure 10: **Default feedback increases long-run effects.**
Results show with scarring of beliefs (solid line) and without (dashed line), often with the two lines on top of each other. Common parameters listed in Table 1. Scenario-specific parameters are: Scenario 1: $\tilde{\phi}_{2020} = 0.90$ *Scenario 2:* $\tilde{\phi}_{2020} = 0.95$

Effect of the default-obsolescence feedback. To understand the role of this feedback rule, suppose that obsolescence is entirely exogenous, i.e. it does not vary with default. This amounts to setting $\mu = 0$ in (11). Figure 10 shows the GDP impact of the COVID-19 shock under our benchmark specification (in the left panel) and without default feedback (in the right panel). The broad patterns are similar with belief revisions accounting for a significant portion of the impact, but the magnitudes are slightly smaller in the right panel (GDP falls by just under 4% in the long-run, relative to a 5% drop in the benchmark). This difference between the two panels is the effect of the the default-obsolescence feedback.

VI.A Financial Assistance Policy.

We consider a simple policy that prevents bankruptcy by reducing firms' debt burden, specifically a reduction in each firm's debt by 10%. This in turn mitigates the effective capital obsolescence and consequently beliefs are slightly less pessimistic going forward. We then simulate the model with these new beliefs and calculate the short- and long-term GDP effects, reported in Table 5.

The table shows that financial assistance of this magnitude only saves 1% of GDP in 2020.

Covid Economics 8, 22 April 2020: 1-40

	No assistance	10% debt reduction	Benefit
GDP drop in 2020	-9%	-8%	1%
NPV of long-term output loss			
from belief scarring	-52%	-45%	7%
from obsolescence	-38%	-34%	4%

Table 5: **Firm Financial Assistance Policy:** No Assistance vs. 10% Debt Reduction
Results are for scenario 1 ($\tilde{\phi}_{2020} = 0.90$). Numbers shown are in percentages of the pre-COVID steady-state GDP.

From that metric alone, one might judge the cost of the policy to be too high.[21] However, preventing bankruptcies in the short-run also helps reduce losses over time. The present discounted value of those savings are worth 11% of 2019 GDP. Of that 11%, 7% comes from ameliorating belief scarring and another 4% comes from the direct effects of limiting capital obsolescence. This exercise shows that considering the long-run consequences can significantly change the cost-benefit analysis for financial policies aimed at assisting firms.

VII What if we had seen a pandemic like this before?

In our benchmark analysis, pre-COVID beliefs were formed using data that did not witness a pandemic (though it did have other tail events like the 2008-09 Great Recession). But, pandemics have occurred before – Jorda et al. (2020) identify 12 pandemics (with greater than 100,000 deaths) going back to the 14th century. This raises the possibility that economic agents in 2019 had some awareness of these past tail events and believed that they could happen again. To understand how this might change our results, we assume that the pre-COVID data sample includes the 1918 episode. Unfortunately, we do not have good data on capital utilization and obsolescence during that period,[22] so we simply use the time series for the capital return shock $\tilde{\phi}_t$ from 1950-2020 as a proxy for the $\tilde{\phi}_t$ series from 1880-1949. In other words, we are asking: What if we had seen all of this unfold exactly the same way before?

The previous data does not change the short-term impact of the shock. But, it does cut the long-term effect of in half. Just before the pandemic of 2020 struck, our data tells us that there has been one pandemic in nearly 140 years. We assess the probability to be about 1-in-140. After 2020, we saw two pandemics in 141 years. Therefore, we revise our perceived probability

[21]Strictly speaking, in our model with a representative agent and lump-sum taxes, there is no real cost to implementing this policy. But, obviously, in a more realistic setting with heterogeneity and distortionary taxation, taking over 10% of corporate debt would entail substantial costs/dead-weight losses.

[22]See Correia et al. (2020) and Velde (2020) for analysis of the economic effects of the 1918 over the short- to medium-term.

Covid Economics 8, 22 April 2020: 1-40

from 1-in-140 to 2-in-141. That is about half the change in probability, relative to the original model where the probability rose from zero to 1-in-70.

But considering data from so long ago does raise the question of whether it is perceived as less relevant. There is a sense that the world has changed, institutions are stronger, science has advanced, in ways that alter the probability of such events. Such gradual change might logically lead one to discount old data.

In a second exercise, we assume that agents discount old data at the rate of 1% per year. In this case, two forces compete. The presence of the 1918 events in the sample reduces the surprise of the new pandemic as before, albeit with a much smaller weight. The countervailing force is that when old data is down-weighted, new data is given a larger weight in beliefs. The larger role of the recent pandemic in beliefs going forward makes belief scarring stronger for the next few decades. These forces more or less cancel each other out leaving the net results indistinguishable, in every respect, from the original results with data only from 1950.

Of course, more recently, we saw SARS, MERS and Ebola arise outside the U.S. Other countries may have learned from these episodes. But the lack of preparation and slow response to events unfolding in China suggests that U.S. residents and policy makers seem to have inferred only that diseases originating abroad stay outside the U.S. borders.

VIII Conclusion

No one knows the true distribution of shocks to the economy. Macroeconomists typically assume that agents in their models know this distribution, as a way to discipline beliefs. For many applications, assuming full knowledge has little effect on outcomes and offers tractability. But for unusually large events, like the current crisis, the difference between knowing these probabilities and estimating them with real-time data can be large. We argue that a more plausible assumption for these phenomena is to assume that agents do the same kind of real-time estimation along the lines of what an econometrician would do. This introduces new, persistent dynamics into a model with otherwise transitory shocks. The essence of the persistence mechanism is this: once observed, a shock (a piece of data) stays in one's data set forever and therefore persistently affects belief formation. The less frequently similar data is observed, the larger and more persistent the belief revision.

When we quantify this mechanism, our model's predictions tell us that the ongoing crisis will have large, persistent adverse effects on the US economy, far greater than the immediate consequences. Preventing bankruptcies or permanent separation of labor and capital, could have enormous consequences for the value generated by the U.S. economy for decades to come.

Covid Economics 8, 22 April 2020: 1-40

CEPR PRESS

Covid Economics 8, 22 April 2020: 1-40

References

Alvarez, F. E., D. Argente, and F. Lippi (2020, April). A Simple Planning Problem for COVID-19 Lockdown. Working Paper 26981, National Bureau of Economic Research.

Angeletos, G.-M. and J. La'O (2013). Sentiments. *Econometrica 81*(2), 739–779.

Atkeson, A. (2020, March). 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.

Baker, S., N. Bloom, S. Davis, and K. Kost (2019, March). Policy news and stock market volatility. Working paper.

Barrero, J. and N. Bloom (2020). Economic uncertainty and the recovery. *Jackson Hole Economic Policy Symposium Proceedings*.

Beaudry, P. and F. Portier (2004). An Exploration into Pigou's Theory of Cycles. *Journal of Monetary Economics 51*, 1183–1216.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018). Really Uncertain Business Cycles. *Econometrica 86*(3), 1031–1065.

Boyarchenko, N., A. Kovner, and O. Shachar (2020). It's What You Say and What You Buy: A Holistic Evaluation of the Corporate Credit Facilities. Technical report, Federal Reserve Bank of New York, Staff Report.

Carvalho, V. M., S. Hansen, Á. Ortiz, J. R. Garcia, T. Rodrigo, S. Rodriguez Mora, and P. Ruiz de Aguirre (2020). Tracking the Covid-19 Crisis with High-Resolution Transaction Data.

Cochrane, J. (2020). What shape recovery? Available at https://johnhcochrane.blogspot.com/2020/04/what-shape-recovery.html (2020/04/06).

Cogley, T. and T. Sargent (2005). The Conquest of US Inflation: Learning and Robustness to Model Uncertainty. *Review of Economic Dynamics 8*, 528–563.

Corbae, D. and P. D'Erasmo (2017, June). Reorganization or liquidation: Bankruptcy choice and firm dynamics. Working Paper 23515, National Bureau of Economic Research.

Correia, S., S. Luck, and E. Verner (2020). Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu. *SSRN Working Paper Series*.

CPPI (2020). Commercial Property Price Index. Technical report, Green Street Adivsors. Accessed: May 15, 2020.

del Negro, M., D. Giannone, M. Giannoni, and A. Tambalotti (2017). Safety, Liquidity, and the Natural Rate of Interest. *Brookings Papers on Economic Activity*, 235–294.

Eaton, J. and M. Gersovitz (1981). Debt with Potential Repudiation: Theoretical and Empirical Analysis. *The Review of Economic Studies*, 289–309.

Eichenbaum, M. S., S. Rebelo, and M. Trabandt (2020, March). The Macroeconomics of Epidemics. Working Paper 26882, National Bureau of Economic Research.

Elenev, V., T. Landvoigt, and S. Van Nieuwerburgh (2020). Can the Covid Bailouts Save the Economy? *COVID Economics 17*.

Fajgelbaum, P. D., E. Schaal, and M. Taschereau-Dumouchel (2017, 05). Uncertainty Traps. *The Quarterly Journal of Economics 132*(4), 1641–1692.

Farboodi, M., G. Jarosch, and R. Shimer (2020, April). Internal and External Effects of Social Distancing in a Pandemic. Working Paper 27059, National Bureau of Economic Research.

Gangestad, S. and D. Buss (1993). Pathogen Prevalence and Human Mate Preferences. *Ethology and Sociobiology 14(2)*, 89–96.

Gourio, F. (2012). Disaster Risk and Business Cycles. *American Economic Review 102(6)*, 2734–66.

Gourio, F. (2013). Credit Risk and Disaster Risk. *American Economic Journal: Macroeconomics 5(3)*, 1–34.

Hall, R. E. (2016). Macroeconomics of Persistent Slumps. In *Handbook of Macroeconomics*, Volume 2, pp. 2131–2181. Elsevier.

Hansen, B. E. (2015). *Econometrics*.

Hansen, L. (2007). Beliefs, Doubts and Learning: Valuing Macroeconomic Risk. *American Economic Review 97(2)*, 1–30.

Huo, Z. and N. Takayama (2015). Higher Order Beliefs, Confidence, and Business Cycles.

Jaimovich, N. and S. Rebelo (2006). Can News About the Future Drive the Business Cycle? *American Economic Review 99(4)*, 1097–1118.

Covid Economics 8, 22 April 2020: 1-40

Johannes, M., L. Lochstoer, and Y. Mou (2016). Learning About Consumption Dynamics. *The Journal of Finance 71*(2), 551–600.

Jorda, O., S. Singh, and A. Taylor (2020). Longer-Run Economic Consequences of Pandemics. *COVID Economics 1*, 1–15.

Kaplan, G., B. Moll, and G. Violante (2020). Pandemics according to HANK.

Kermack, W. O. and A. G. McKendrick (1927). A Contribution to the Mathematical Theory of Epidemics. *Proceedings of the Royal Society of London. Series A 115*(772), 700–721.

Khan, A., T. Senga, and J. K. Thomas (2017). Default Risk and Aggregate Fluctuations in an Economy with Production Heterogeneity. 2017 Meeting Papers 889, Society for Economic Dynamics.

Koren, M. and R. Pető (2020). Business Disruptions from Social Distancing. *COVID Economics 2*, 13–31.

Kozeniauskas, N., A. Orlik, and L. Veldkamp (2014). What Are Uncertainty Shocks? *Journal of Monetary Economics 100*, 1–15.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2019). The Tail that Keeps the Riskless Rate Low. *National Bureau of Economic Research, Macroeconomics Annual 33*(1), 253–283.

Kozlowski, J., L. Veldkamp, and V. Venkateswaran (2020). The Tail that Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation. *Journal of Political Economy 128*(8), 2839–2879.

Krishnamurthy, A. and W. Li (2020, May). Dissecting Mechanisms of Financial Crises: Intermediation and Sentiment. Working Paper 27088, National Bureau of Economic Research.

Krishnamurthy, A. and A. Vissing-Jorgensen (2012). The Aggregate Demand for Treasury Debt. *Journal of Political Economy 120*(2), 233–267.

Krugman, P. (2020). Notes on the Coronacoma. Available at https://www.nytimes.com/2020/04/01/opinion/notes-on-the-coronacoma-wonkish.html (2020/04/01).

Lagos, R. and R. Wright (2005). A Unified Framework for Monetary Theory and Policy Analysis. *Journal of Political Economy 113*(3), 463–484.

Leibovici, F., A. M. Santacreu, and M. Famiglietti (2020). How the Impact of Social Distancing Ripples through the Economy. Available at https://www.stlouisfed.

org/on-the-economy/2020/april/impact-social-distancing-ripples-economy (2020/04/07).

Loehle, C. (1995). Social Barriers to Pathogen Transmission in Wild Animal Populations. *Ecology 76*(2), 326–335.

Lorenzoni, G. (2009). A Theory of Demand Shocks. *American Economic Review 99*(5), 2050–84.

Ludvigson, S. C., S. Ma, and S. Ng (2020, April). Covid19 and the Macroeconomic Effects of Costly Disasters. Working Paper 26987, National Bureau of Economic Research.

Moody's Investors Service (2015). Annual Default Study: Corporate Default and Recovery Rates, 1920-2014. Technical report.

Nimark, K. (2014). Man-Bites-Dog Business Cycles. *American Economic Review 104(8)*, 2320–67.

Reinhart, C. and K. Rogoff (2009). *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press.

Reinhart, C. and K. Rogoff (2020). The Coronavirus Debt Threat. Available at https://www.wsj.com/articles/the-coronavirus-debt-threat-11585262515 (2020/03/26).

Stock, J. H. (2020, March). Data Gaps and the Policy Response to the Novel Coronavirus. Working Paper 26902, National Bureau of Economic Research.

Straub, L. and R. Ulbricht (2013). Credit Crunches, Information Failures, and the Persistence of Pessimism.

Van Nieuwerburgh, S. and L. Veldkamp (2006). Learning Asymmetries in Real Business Cycles. *Journal of Monetary Economics 53(4)*, 753–772.

Velde, F. R. (2020). What Happened to the US Economy During the 1918 Influenza Pandemic? A View Through High-Frequency Data. *SSRN Working Paper Series*.

Veldkamp, L. and J. Wolfers (2007). Aggregate Shocks or Aggregate Information? Costly Information and Business Cycle Comovement. *Journal of Monetary Economics 54*, 37–55.

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

Covid Economics 8, 22 April 2020: 1-40

Covid Economics 8, 22 April 2020: 41-57

A Solution

A.A Equilibrium Characterization

An equilibrium is the solution to the following system of equations:

$$
\begin{aligned}
1 &= \mathbb{E}M_{t+1}\left[R^k_{t+1}\right]J^k(\underline{v}_t) && (25)\\
R^k_{t+1} &= \frac{(1-\alpha)\phi^{\alpha+\nu}_{t+1}\hat{k}^{\alpha}_{t+1}l^{1-\alpha}_{t+1}+(1-\delta)\,\phi_{t+1}\hat{k}_{t+1}}{\hat{k}_{t+1}} && (26)\\
\frac{1-\gamma}{\gamma}\frac{c_t}{1-l_t} &= \frac{(1-\alpha)y_t}{l_t} && (27)\\
(1-\theta)\,\mathbb{E}_t\left[M_{t+1}\underline{v}_t f(\underline{v}_t)\right] &= \left(\frac{\chi-1}{\chi}\right)\mathbb{E}_t\left[M_{t+1}\left(1-F(\underline{v}_t)\right)\right] && (28)\\
c_t &= \phi^{\alpha+\nu}_t\hat{k}^{\alpha}_t l^{1-\alpha}_t+(1-\delta)\,\phi_t\hat{k}_t-\hat{k}_{t+1} && (29)\\
U_t &= \left[(1-\beta)\left(u(c_t,l_t)\right)^{1-\psi}+\beta\mathbb{E}\left(U^{1-\eta}_{t+1}\right)^{\frac{1-\psi}{1-\eta}}\right]^{\frac{1}{1-\psi}} && (30)
\end{aligned}
$$

where

$$
\begin{aligned}
\ln\phi_t &= \ln\tilde{\phi}_t-\mu F(\underline{v}_t)^{1-\varpi}\\
\underline{v}_t &= \frac{lev_{t+1}}{R^k_{t+1}}\\
J^k(\underline{v}_t) &= 1+(\chi-1)\,\underline{v}_t\left(1-F(\underline{v}_t)\right)+(\chi\theta-1)\,h(\underline{v}_t)\\
M_{t+1} &= \left(\frac{dU_t}{dc_t}\right)^{-1}\frac{dU_t}{dc_{t+1}}=\beta\left[\mathbb{E}\left(U^{1-\eta}_{t+1}\right)\right]^{\frac{\eta-\psi}{1-\eta}}U^{\psi-\eta}_{t+1}\left(\frac{u(c_{t+1},l_{t+1})}{u(c_t,l_t)}\right)^{-\psi}
\end{aligned}
$$

A.B Solution Algorithm

To solve the system described above at any given date t (i.e. after any observed history of $\tilde{\phi}_t$), we recast it in recursive form with grids for the aggregate state ($\hat{k}$) and the shocks $\tilde{\phi}$. We then use an iterative procedure:

- Estimate $\hat{g}$ on the available history using the kernel estimator.
- Start with a guess (in polynomial form) for $U(\hat{k},\tilde{\phi}), c(\hat{k},\tilde{\phi}), l(\hat{K},\tilde{\phi})$.
- Solve (25)-(28) for $\hat{k}', lev', l$ using a non-linear solution procedure.
- Verify/update the guess for U, c, l using (29)-(30) and iterate until convergence.

CEPR PRESS

Covid-19: Testing inequality in New York City[1]

Stephanie Schmitt-Grohé,[2] Ken Teoh[3] and Martín Uribe[4]

Date submitted: 16 April 2020; **Date accepted:** 17 April 2020

Covid Economics 8, 22 April 2020: 41-57

Motivated by reports in the media suggesting unequal access to Covid-19 testing across incomes, we analyze zip-code level data on the number of Covid-19 tests, test results, and income per capita in New York City. We find that the number of tests administered is evenly distributed across income levels. In particular, the test distribution across income levels is significantly more egalitarian than the distribution of income itself: The ten percent of the city's population living in the richest zip codes received 11 percent of the Covid-19 tests and 29 percent of the city's income. The ten percent of the city's population living in the poorest zip codes received 10 percent of the tests but only 4 percent of the city's income. At the same time, we find significant disparity in the fraction of tests that come back negative for the Covid-19 disease across income levels: moving from the poorest zip codes to the richest zip codes is associated with an increase in the fraction of negative Covid-19 test results from 38 to 65 percent.

1 We thank for comments Dragan Filipovich, Antti Ilmanen, and Gabriel Picone and seminar participants at Columbia University.

2 Professor of Economics, Columbia University and CEPR Research Fellow.
3 PhD Candidate in Economics, Columbia University.
4 Professor of Economics, Columbia University.

Covid Economics 8, 22 April 2020: 41-57

1 Introduction

The United States is the epicenter of the 2020 coronavirus outbreak. An important aspect of this health crisis is how it impacts different income groups. In particular, a much debated issue is how access to health care is linked to peoples' relative position in the income ladder. This paper addresses a specific aspect of this debate, namely, whether higher income gave privileged access to tests for the Covid-19 disease. Several media outlets and observers have suggested that the distribution of tests inherited the well-known inequalities in the distribution of income. For example, on March 18, 2020, the New York Times ran an article with the headline, "Need a Coronavirus Test? Being Rich and Famous May Help."

To analyze the question of inequality in testing for Covid-19 across income groups, we use zip-code-level data from New York City on the number of Covid-19 tests and test outcomes as of April 2 and April 13, 2020. We combine this information with data on income per capita and population by zip code. New York City is a relevant laboratory for evaluating this question for the following reasons: First, it is the city hardest hit by the coronavirus. Second, the residents of all New York city zip codes are subject to the same health policies and regulations at the local, state, and federal levels. Third, income per capita displays significant variation across zip codes. And fourth, data across all zip codes is produced by the same statistical agencies, guaranteeing cross-sectional comparability. Data on Covid-19 tests and test results for the 177 zip codes in New York come from the New York City Department of Health and Mental Hygiene and data on zip-code level income per capita and population is from the 2014-2018 American Community Survey (ACS).

We find that the number of tests administered is evenly distributed across income levels. In particular, the test distribution is significantly more egalitarian than the distribution of income: The 10 percent of the population living in the richest zip codes received 11 percent of the Covid-19 tests and 29 percent of the city's income. The 10 percent of the population living in the poorest zip codes received 9 percent of the tests and only 4 percent of the city's income.

If disease were equally distributed across income levels, then the reported even distribution of tests could be consistent with equal access to testing. However, if the Covid-19 disease were more prevalent at lower incomes, then this finding could reflect unequal access, for in this case, the poorest income groups should account for a disproportionally larger share of the administered tests. Establishing the spread of the coronavirus would require random testing, which is not available. The available data on test results is influenced by patients' self-selection into testing and health-care providers' selection criteria on whom to test. It is nonetheless of interest to investigate whether the combination of incidence, self-selection, and

Covid Economics 8, 22 April 2020: 41-57

rationing by the health department jointly have an impact on the relation between Covid-19 test results and income. To shed some light on this issue, we analyze the distribution of negative test results across income groups. A negative test result indicates that the patient does not have Covid-19. We find significant disparities across income levels in the fraction of tests for Covid-19 that come back negative: moving from the poorest zip codes to the richest zip codes is associated with an increase in the fraction of negative Covid-19 test results of 27 percentage points, from 38 to 65 percent. Furthermore, controlling for income, the fraction of negative tests is lower in zip codes with larger shares of non-white residents, although the effect is quantitatively small.

Taken together the results on tests and test-outcome distributions by income suggest that an egalitarian distribution of tests may not be tantamount to equal access. The findings are consistent with the non-mutually exclusive hypotheses that Covid-19 is more widespread in poorer zip codes and that the severity of symptoms that triggers a Covid-19 test to be administered is higher for poor New York City residents. The latter hypothesis, in turn, could be driven by both supply and demand factors. For example, a supply side explanation could be that residents of poorer zip codes get a disproportionate share of positive test results because health care providers apply a more stringent symptom threshold to them than to residents of richer zip codes. A demand side explanation could be that poorer residents tend to wait longer before they seek medical assistance.

This paper is related to contemporaneous and independent work by Borjas (2020), who used a similar dataset on Covid-19 tests, test results, population, and income in New York City at the zip code level.[1] The main difference in the empirical strategy is that the present study estimates Lorenz curves and Gini coefficients for testing and negative test results, whereas Borjas focuses on regression analysis. Similar to the present study, Borjas finds that persons who reside in richer neighborhoods were more likely to test negative for Covid-19. However, the methodological differences lead to seemingly different conclusions regarding equity in testing. While the present study finds that the distribution of Covid-19 tests was egalitarian across income, Borjas finds that people residing in poor neighborhoods were less likely to be tested than people residing in rich neighborhoods. These two conclusions are not necessarily inconsistent with each other, for regressions of the odds of Covid-19 tests on household income do not necessarily reflect inequity in the distribution of tests across income levels. The reason is that such regressions do not take into account the weight of the different zip codes in the income distribution, a feature that, by construction, is factored in

[1] There are some differences in the datasets: (1) this paper uses test and test result data for April 2 and 13, whereas Borjas uses data for April 5; (2) this paper uses income and population data from the 2014-2018 ACS, whereas Borjas uses the 2010-2014 ACS; and (3) This paper uses mean per capita income at the zip code level, whereas Borjas uses mean household income at the zip code level.

by Lorenz curves.

The remainder of the paper is in five sections. Section 2 describes the data and their sources. Section 3 documents the distribution of tests for Covid-19 across income levels, and section 4 the distribution of negative test outcomes. Section 5 shows that the findings are robust to extending the sample from April 2 to April 13, 2020, a period in which the number of Covid-19 tests administered in New York City almost tripled. Section 6 concludes with a discussion of the results. An appendix contains the formulas used for the construction of the inequality measures (Lorenz curves and Gini coefficients). The data used in this investigation along with the Matlab code to replicate the tables and figures are available at http://www.columbia.edu/~mu2166/stu_covid19/.

Covid Economics 8, 22 April 2020: 41-57

2 Data

The data used in this investigation include the cumulative number of New York City residents who were ever tested for Covid-19 and the number of residents who tested positive as of April 2, 2020 for each of the city's 177 zip codes. The data source is the New York City Department of Health and Mental Hygiene (DOHMH) Incident Command System for Covid-19 Response (https://github.com/nychealth/coronavirus-data#tests-by-zctacsv). In addition, the dataset includes zip-code level data on per capita income in the past 12 months in dollars of 2018 surveyed over the period from 2014-2018, population, and racial composition as of 2018. The source is the American Community Survey,
series code `B19301_001E` and `B01003_001E`, respectively, (https://www.census.gov/data/developers/data-sets/acs-5year.html).

Table 1 displays summary statistics for Covid-19 testing, per capita income, and population across the 177 zip codes of New York City. The number of tests and of negative test results are cumulative as of April 2, 2020. Each zip code received on average 908 Covid-19 tests per 100,000 residents. On average half of the tests came back negative. The sample displays significant cross sectional variation, with standard deviations of 268 and 9 percent for tests and negatives, respectively. The City also displays large cross sectional variation in income per capita with residents of the poorest zip code receiving $13,394 on average per year and residents of the richest $147,547.

3 Covid-19 Test Distribution By Income

One of the goals of this paper is to document the distribution of tests for Covid-19 across income levels. Figure 1 displays the Lorenz curve of the number of tests for Covid-19 across

Covid Economics 8, 22 April 2020: 41-57

Table 1: Summary Statistics

Statistic	Covid-19 Total Tests per 100,000	Covid-19 Share of Negative Tests (%)	Per Capita Income dollars of 2018	Population
Mean	908	49	44287	47645
Median	860	49	31779	42653
Std.Dev.	268	9	31919	26698
Max	2390	75	147547	112425
Min	450	23	13394	3028

Notes. Summary statistics are computed across the 177 New York City zip codes. Total tests and negative tests are cumulative as of April 2, 2020. Replication code `summary_statistics.m` in `stu_covid19.zip`.

Figure 1: Lorenz Curves of Covid-19 Tests and Mean Income Across New York City Zip Codes

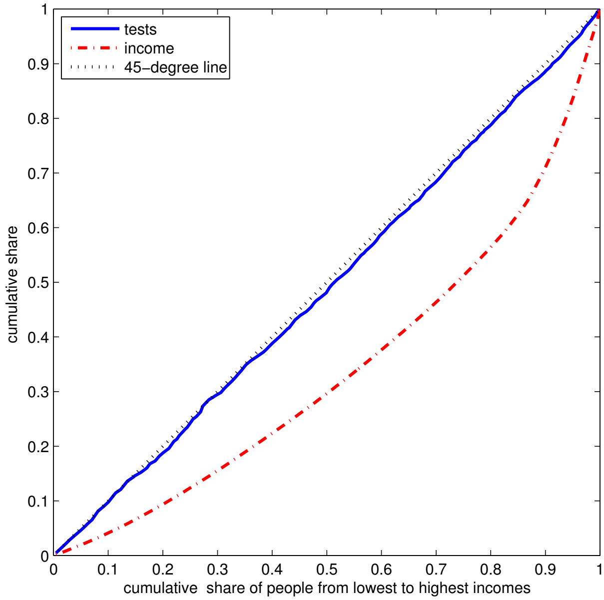

Notes. Own calculations based on data from the New York City Department of Health and Mental Hygiene, as of April 2, 2020, and American Community Survey. Replication code `gini_testing.m` in `stu_covid19.zip`.

Covid Economics 8, 22 April 2020: 41-57

Table 2: Gini Coefficients

Income	0.32
Covid-19 Tests	0.02
Covid-19 Negative Test Results	0.09

Note. Own calculations based on data from the New York City Department of Health and Mental Hygiene as of April 2, 2020, and American Community Survey. Replication code `gini_testing.m` and `gini_negatives_testing.m` in `stu_covid19.zip`.

New York City zip codes as of April 2, 2020. For comparison, the figure also displays the Lorenz curve for the distribution of per capita income. The horizontal axis measures cumulative population shares sorted by income. The vertical axis displays with a solid line the cumulative share of Covid-19 tests and with a broken line the cumulative share of income. The appendix presents the formulas used to construct the Lorenz curves and the Gini coefficients discussed below.

The key message of the figure is that Covid-19 tests are almost perfectly distributed across income groups in New York City. Graphically, this is reflected in the fact that the Lorenz curve is nearly equal to the 45-degree line. The 10 percent of the population living in zip codes with the highest income per capita received 11 percent of all Covid-19 tests administered in the City and the 10 percent of the population living in zip codes with the lowest income per capita received 10 percent of the tests.

A more comprehensive and frequently used measure of inequality is the Gini coefficient, which is given by the ratio of the area between the 45-degree line and the Lorenz curve to the triangular area below the 45-degree line. The Gini coefficient associated with the distribution of Covid-19 testing across income levels is equal to 0.02, with a value of 0 representing a perfectly even distribution.

The evenness of the distribution of Covid-19 tests across income levels contrasts with the inequality in the distribution of income per capita across zip codes. This is reflected in the Lorenz curve for the income distribution being significantly below the 45-degree line. The top decile of the population earns 29 percent of total income, whereas the bottom decile earns only 4 percent. The Gini coefficient of income inequality is 0.32, sixteen times higher than the Gini coefficient of testing inequality. Because of the averaging of income per capita within each zip code, the reported Gini coefficient of New York's income distribution, 0.32, is lower than the one that results from using data at the household level, 0.55 for 2018 according to the American Community Survey. However, the Gini coefficient of 0.32 is the

CEPR PRESS

Covid Economics 8, 22 April 2020: 41-57

Figure 2: Share of Negative Tests and Mean Income Per Capita Across New York City Zip Codes

Notes. Own calculations based on data from the New York City Department of Health and Mental Hygiene, as of April 2, 2020, and American Community Survey. The negative share is defined as the number of negative test results divided by the total number of tests. The solid line is the OLS regression. Replication code `negatives_vs_income.m` in `stu_covid19.zip`.

relevant one for the purpose of the present analysis because the most disaggregated level at which test for Covid-19 statistics are available is the zip code.

4 Test-Result Inequality

The available data make it possible to address the question of how outcomes of Covid-19 tests vary across income at the zip code level. Figure 2 plots with circles the share of tests that came back negative (the patient is not infected with the coronavirus) against income per capita at the zip code level. Income per capita is plotted on a logarithmic scale. The scatterplot displays a clear positive relationship between income per capita and the corresponding share of negative Covid-19 test results. The solid line is the OLS regression, which is given by

$$s_i^n = -0.69 + 0.11 \, \ln y_i^c + \epsilon_i, \quad R^2 = 0.48,$$

where s_i^n, y_i^c, and ϵ_i denote, respectively, the share of negatives, income per capita, and the regression residual in zip code $i = 1, \ldots, 177$. The slope coefficient, 0.11, is significant at the

Covid Economics 8, 22 April 2020: 41-57

Figure 3: Lorenz Curves of Negative Tests and Total Tests for Covid-19 Across New York City Zip Codes

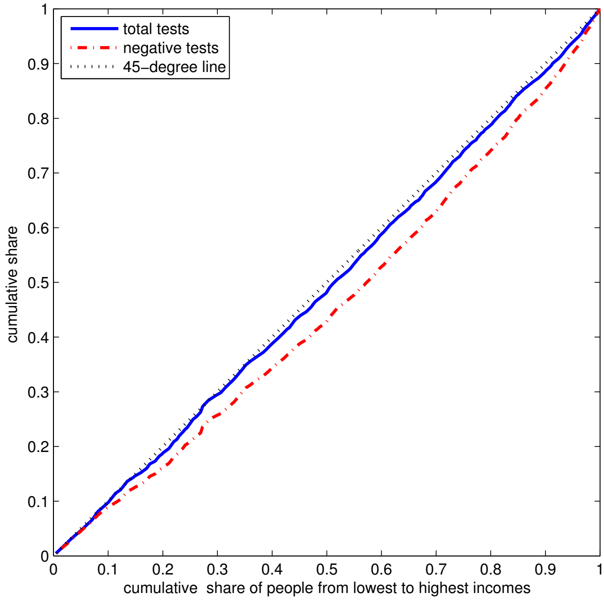

Notes. Own calculations based on data from the New York City Department of Health and Mental Hygiene, as of April 2, 2020, and American Community Survey. Replication code `gini_negatives_testing.m` in `stu_covid19.zip`.

1 percent confidence level. It implies that moving from the poorest zip codes to the richest zip codes is associated with an increase in the share of negative Covid-19 test results of 27 percentage points, from 38 percent to 65 percent. It follows that unlike the distribution of tests for Covid-19, the distribution of test outcomes across income is significantly regressive. This suggests that the observed egalitarian distribution of tests need not reflect equal access to tests.

Expanding the OLS regression to include the share of black residents, s_i^b, and the share of other racial minorities, s_i^o (neither blacks nor whites), yields $s_i^n = -0.44 + 0.09 \ln y_i^c - 0.08\, s_i^b - 0.05\, s_i^o + \epsilon_i$, with $R^2 = 0.51$. According to this expression, race has a negative but relatively minor effect on the share of negative tests. Controlling for income, a one standard deviation increase in the share of black residents (25 percentage points) is associated with a fall in the share of negative test results of 2 percentage points. The association is even weaker for other racial minorities.

To emphasize the finding that the distribution of negative test results is unequal across

CEPR PRESS

income levels, figure 3 displays with a dash-dotted line the Lorenz curve for the distribution of the number of negative test results for Covid-19. For comparison, the figure reproduces with a solid line from figure 1 the Lorenz curve of the distribution of total tests. The Lorenz curve associated with negative test results is farther below the 45-degree line than the one associated with total tests, reflecting more inequality across income levels in test outcomes than in the number of tests administered. The Gini coefficient for the distribution of negatives is 0.09, almost five times larger than the one corresponding to the distribution of total tests.

5 Dynamics

The analysis thus far was conducted on data of cumulated tests and test results as of April 2, 2020. As a robustness check, this section examines data up to April 13. In the intermittent period, the number of administered Covid-19 tests in New York City increased from 73,215 to 182,099 or 150 percent.

Figure 4 displays the Lorenz curve of the Covid-19 test distribution as of April 13, 2020. For comparison, the figure reproduces from Figure 1 the income distribution. The main message conveyed by the figure is that the distribution of tests continues to be egalitarian (i.e., close to the 45-degree line) after the significant increase in the number of New Yorkers that were tested. If at all, it became slightly more in favor of low income groups: Between April 2 and April 13, the fraction of test going to the bottom decile of the income distribution increased from 10 to 11 percent and that going to the top decile fell from 11 to 8 percent.

Figure 5 displays the Lorenz curve of the distribution of negative test results for Covid-19 in New York City as of April 13, 2020. The figure indicates that it continues to be the case that the distribution of negative test results is more unequal than that of administered tests. The bottom decile of the income distribution received 11 percent of the tests but 8 percent of the negative test results, while the top decile received 8 percent of tests administered, but 11 percent of the negative results.

6 Discussion and Conclusion

This paper contributes to the economic analysis of pandemics. It documents that in New York City, the most castigated city by the 2020 coronavirus outbreak, the ex-post likelihood of being tested for Covid-19 was evenly distributed across income levels measured at the zip-code unit. The bottom decile of the income distribution received 10 percent of all tests and the top decile 11 percent. The distribution of test outcomes, by contrast, displays significant

Covid Economics 8, 22 April 2020: 41-57

Figure 4: Lorenz Curves of Covid-19 Tests and Mean Income Across New York City Zip Codes as of April 13, 2020

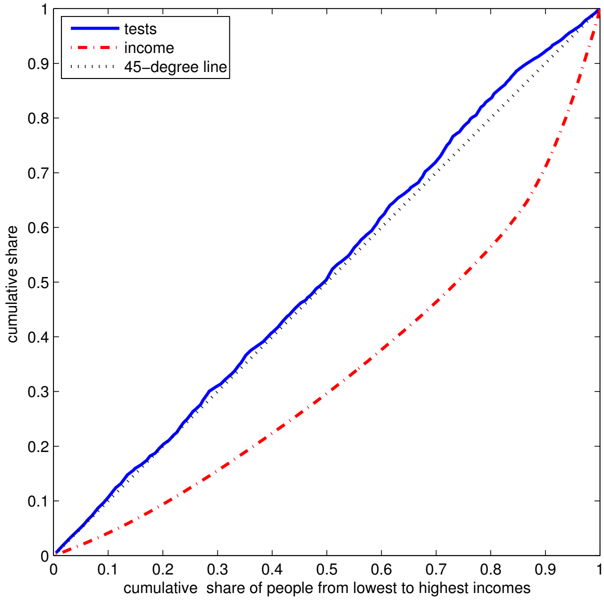

Notes. Own calculations based on data from the New York City Department of Health and Mental Hygiene, as of April 13, 2020, and American Community Survey.

Covid Economics 8, 22 April 2020: 41-57

CEPR PRESS

Covid Economics 8, 22 April 2020: 41-57

Figure 5: Lorenz Curves of Negative Tests and Total Tests for Covid-19 Across New York City Zip Codes, April 13, 2020

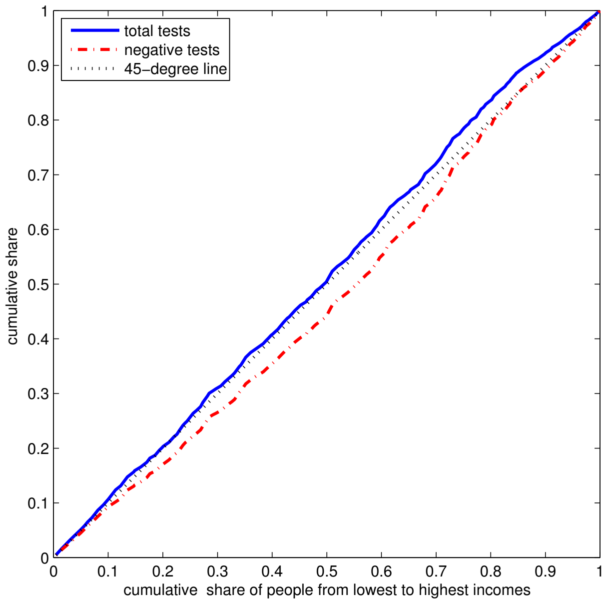

Notes. Own calculations based on data from the New York City Department of Health and Mental Hygiene, as of April 13, 2020, and American Community Survey.

CEPR PRESS

Covid Economics 8, 22 April 2020: 41-57

inequality across income. The ex-post probability of testing negative for Covid-19 in the zip codes with the lowest per capita incomes was 38 percent compared to 65 percent in the zip codes enjoying the highest per capita incomes.

In light of the reported unequal distribution of test outcomes against lower income groups, it is possible that the observed egalitarian distribution of tests was associated with testing not being proportional to incidence. However, the data analyzed in this paper does not provide sufficient information to establish this conclusion. To see this, apply Bayes law, to obtain

$$P(pos|test)P(test) = P(test|pos)P(pos),$$

where P denotes probability, $test$ denotes being tested, and pos denotes a positive test result. Evaluate this expression at the bottom and top deciles of the income distribution, and take ratios, to obtain

$$\frac{P^{poor}(pos|test)}{P^{rich}(pos|test)} \times \qquad \frac{P^{poor}(test)}{P^{rich}(test)} = \qquad \frac{P^{poor}(test|pos)}{P^{rich}(test|pos)} \times \frac{P^{poor}(pos)}{P^{rich}(pos)}$$

The estimates obtained in section 3 suggests that $\frac{P^{poor}(pos|test)}{P^{rich}(pos|test)} = 1.5$ and the estimates of section 4 that $\frac{P^{poor}(test)}{P^{rich}(test)} = 1$. Therefore we can write

$$1.5 \quad \times \qquad 1 \quad = \qquad \frac{P^{poor}(test|pos)}{P^{rich}(test|pos)} \times \frac{P^{poor}(pos)}{P^{rich}(pos)}.$$

If the probability of getting tested for the coronavirus conditional on being infected is the same for the bottom and top deciles, $\frac{P^{poor}(test|pos)}{P^{rich}(test|pos)} = 1$, then it follows that the incidence rate is 50 percent higher in the bottom decile than in the top decile. In this case, the estimated egalitarian distribution of tests would not reflect the relevant incidence of Covid-19 across income groups.

However, testing selection could introduce variation in $P(test|pos)$ across income, which in turn will affect the inference about differences across income levels in incidence of Covid-19, $P(pos)$. In particular, if testing criteria are more stringent in poor neighborhoods than in rich ones, then $P(test|pos)$ could be larger in poor zip codes than in rich ones. To illustrate how testing selection criteria can affect $P(test|pos)$, consider the following example. Suppose there are 100 people in the population, of which 50 have the flu, 30 have corona, and 20 have corona and the flu, so that only 10 have corona but not the flu. People who have the flu have only 1 symptom, say fever. People who have corona have 1 additional symptom, say lack of smell. The health authority has only 5 Covid-19 tests. Suppose initially the testing criterion is just fever. Sixty people therefore meet the criterion, so $P(test|pos) = 1/12$. Now suppose the criterion is fever and lack of smell. In this case only 30 people qualify, namely, the people infested with corona. So $P(test|pos) = 2/12$. This example suggests that if testing selection criteria were more stringent in poor neighborhoods, then, given the results reported in this paper, incidence, $P(pos)$ would be less than fifty percent larger in the poor zip codes than in the rich ones and in principle could even be smaller.

Obtaining reliable measures of incidence would require randomized testing. This would make it possible to design more efficient allocations of tests than relying on what, based on the evidence presented in this paper, appears to be a simple egalitarian rule.

Appendix: Calculation of Gini Coefficients and Lorenz Curves

Let y_i^c denote per capita income in zip code $i = 1, \ldots, 177$. Suppose that y_i^c is sorted in ascending order, so that $y_i^c < y_{i+1}^c$ for any $1 \leq i < 177$ and let p_i be the population of zip code i. Then income in zip code i, denoted y_i is approximated by

$$y_i = y_i^c p_i$$

The share of income of zip code i in total New York City income is defined as

$$s_i^y = \frac{y_i}{\sum_{i=1}^{177} y_i}$$

The cumulative income share up to the ith poorest zip code, denoted S_i^y, is given by

$$S_i^y = \sum_{k=1}^{i} s_k^y$$

Similarly, the population share of the ith poorest zip code, denoted s_i^p is given by

$$s_i^p = \frac{p_i}{\sum_{i=1}^{177} p_i}.$$

And the cumulative population share up to the ith poorest zip code, denoted S_i^p, is given by

$$S_i^p = \sum_{k=1}^{i} s_k^p.$$

Let τ_i denote the number of Covid-19 tests in the ith poorest zip code. Then, the share

Covid Economics 8, 22 April 2020: 41-57

of tests in zip code i, denoted s_i^{τ} is given by

$$s_i^{\tau} = \frac{\tau_i}{\sum_{i=1}^{177} \tau_i}$$

And the corresponding cumulative share up to the ith poorest zip code, denoted S_i^{τ}, is

$$S_i^{\tau} = \sum_{k=1}^{i} s_k^{\tau}.$$

Figure 1 plots the variables S_i^y and S_i^{τ} (vertical axis) against the variable S_i^p (horizontal axis).

The Gini coefficient of the income distribution across zip codes is measured as

$$\text{Gini coefficient of income distribution} = 1 - \frac{\sum_{i=1}^{177} s_i^p S_i^y}{\sum_{i=1}^{177} s_i^p S_i^p},$$

and the Gini coefficient of the Covid-19 testing distribution across income levels by

$$\text{Gini coefficient of Covid-19 testing across income levels} = 1 - \frac{\sum_{i=1}^{177} s_i^p S_i^{\tau}}{\sum_{i=1}^{177} s_i^p S_i^p}$$

References

Borjas, George J., "Demographics Determinants of Testing Incidence and Covid-19 Infections in New York City Neighborhoods," NBER working paper 26952, April 2020.

Covid Economics 8, 22 April 2020: 41-57

CEPR PRESS

Covid Economics 8, 22 April 2020: 41-57

it follows that the incidence rate is 50 percent higher in the bottom decile than in the top decile. In this case, the estimated egalitarian distribution of tests would not reflect the relevant incidence of Covid-19 across income groups.

However, testing selection could introduce variation in $P(test|pos)$ across income, which in turn will affect the inference about differences across income levels in incidence of Covid-19, $P(pos)$. In particular, if testing criteria are more stringent in poor neighborhoods than in rich ones, then $P(test|pos)$ could be larger in poor zip codes than in rich ones. To illustrate how testing selection criteria can affect $P(test|pos)$, consider the following example. Suppose there are 100 people in the population, of which 50 have the flu, 30 have corona, and 20 have corona and the flu, so that only 10 have corona but not the flu. People who have the flu have only 1 symptom, say fever. People who have corona have 1 additional symptom, say lack of smell. The health authority has only 5 Covid-19 tests. Suppose initially the testing criterion is just fever. Sixty people therefore meet the criterion, so $P(test|pos) = 1/12$. Now suppose the criterion is fever and lack of smell. In this case only 30 people qualify, namely, the people infested with corona. So $P(test|pos) = 2/12$. This example suggests that if testing selection criteria were more stringent in poor neighborhoods, then, given the results reported in this paper, incidence, $P(pos)$ would be less than fifty percent larger in the poor zip codes than in the rich ones and in principle could even be smaller.

Obtaining reliable measures of incidence would require randomized testing. This would make it possible to design more efficient allocations of tests than relying on what, based on the evidence presented in this paper, appears to be a simple egalitarian rule.

Appendix: Calculation of Gini Coefficients and Lorenz Curves

Let y_i^c denote per capita income in zip code $i = 1, \ldots, 177$. Suppose that y_i^c is sorted in ascending order, so that $y_i^c < y_{i+1}^c$ for any $1 \leq i < 177$ and let

CEPR PRESS

p_i be the population of zip code i. Then income in zip code i, denoted y_i is approximated by

$$y_i = y_i^c p_i$$

The share of income of zip code i in total New York City income is defined as

$$s_i^y = \frac{y_i}{\sum_{i=1}^{177} y_i}$$

The cumulative income share up to the ith poorest zip code, denoted S_i^y, is given by

$$S_i^y = \sum_{k=1}^{i} s_k^y$$

Similarly, the population share of the ith poorest zip code, denoted s_i^p is given by

$$s_i^p = \frac{p_i}{\sum_{i=1}^{177} p_i}.$$

And the cumulative population share up to the ith poorest zip code, denoted S_i^p, is given by

$$S_i^p = \sum_{k=1}^{i} s_k^p.$$

Let τ_i denote the number of Covid-19 tests in the ith poorest zip code. Then, the share of tests in zip code i, denoted s_i^τ is given by

$$s_i^\tau = \frac{\tau_i}{\sum_{i=1}^{177} \tau_i}$$

And the corresponding cumulative share up to the ith poorest zip code, denoted S_i^τ, is

$$S_i^\tau = \sum_{k=1}^{i} s_k^\tau.$$

Figure 1 plots the variables S_i^y and S_i^τ (vertical axis) against the variable S_i^p (horizontal axis).

The Gini coefficient of the income distribution across zip codes is mea-

Covid Economics 8, 22 April 2020: 41-57

sured as

$$\text{Gini coefficient of income distribution} = 1 - \frac{\sum_{i=1}^{177} s_i^p S_i^y}{\sum_{i=1}^{177} s_i^p S_i^p},$$

and the Gini coefficient of the Covid-19 testing distribution across income levels by

$$\text{Gini coefficient of Covid-19 testing across income levels} = 1 - \frac{\sum_{i=1}^{177} s_i^p S_i^\tau}{\sum_{i=1}^{177} s_i^p S_i^p}$$

References

Borjas, George J., "Demographics Determinants of Testing Incidence and Covid-19 Infections in New York City Neighborhoods," NBER working paper 26952, April 2020.

Covid Economics 8, 22 April 2020: 41-57

Covid Economics Issue 8, 22 April 2020

A workable strategy for Covid-19 testing: Stratified periodic testing rather than universal random testing[1]

Matthew Cleevely,[2] Daniel Susskind,[3] David Vines,[4] Louis Vines,[5] and Samuel Wills[6]

Date submitted: 15 April 2020; **Date accepted:** 17 April 2020

This paper argues for the regular testing of members of at-risk groups more likely to be exposed to SARS-CoV-2 as a strategy for reducing the spread of Covid-19 and enabling the resumption of economic activity. We call this 'stratified periodic testing'. It is 'stratified' as it is based on at-risk groups, and 'periodic' as everyone in the group is tested at regular intervals. We argue that this is a better use of scarce testing resources than 'universal random testing', as recently proposed by Paul Romer. We find that universal testing would require checking over 21 percent of the population every day to reduce the effective reproduction number of the epidemic, R', down to 0.75 (as opposed to 7 percent as argued by Romer). We obtain this rate of testing using a corrected method for calculating the impact of an infectious person on others, where testing and isolation takes place, and where there is self-isolation of symptomatic cases. We also find that any delay between testing and the result being known significantly increases the effective reproduction number and that one day's delay is equivalent to having a test that is 30 percent less accurate.

1 We are especially grateful to David Cleevely, CBE, FREng, Chair of the COVID Positive Response Committee, Royal Academy of Engineering, and Frank Kelly, CBE, FRS, Emeritus Professor of the Mathematics of Systems in the Statistical Laboratory and Fellow of Christ's College, University of Cambridge, for their editorial contributions to the paper. In addition, we have received many helpful comments and suggestions from Eric Beinhocker, Adam Bennett, Wendy Carlin, Daniela Massiceti, Warwick McKibbin, Bob Rowthorn, Stephen Wright, and Peyton Young.

2 Founder, 10to8 ltd.
3 Fellow in Economics, Balliol College, University of Oxford.
4 Emeritus Professor of Economics and Emeritus Fellow, Balliol College, University of Oxford; CEPR Research Fellow.
5 Data scientist.
6 External Research Associate, School of Economics, University of Sydney.

Covid Economics 8, 22 April 2020: 58-84

Covid Economics 8, 22 April 2020: 58-84

1 Introduction and Summary

Governments around the world are looking for a testing strategy for Covid-19. And they are keen to see how this testing strategy might help ensure that the escape from lockdown is as speedy as a possible.

In a talk given on 3 April 2020, Paul Romer proposed population-wide testing and isolation – what we call “universal random testing”.[1] He made the important points that the economic benefit of a speedier recovery would be measured in $trillions and this would easily justify spending $billions on testing, and that there is no necessity for such testing to be highly accurate. Unfortunately, as we will show, the model used in that talk is flawed. Our corrected model suggests that Romer’s proposal is unlikely to work in practice. Instead, we show that targeted testing of particular groups, what we call “stratified periodic testing”, and the subsequent isolation of those who test positive, would be a more effective tool in reducing the spread of Covid-19.

Romer uses a simple model to show how random testing of 7 percent of the population per day for evidence of infection (using antigen tests) would be sufficient to halt the pandemic. While already a big ask, we argue that his proposal ignores asymptomatic cases and rests on a mistaken calculation. A corrected model implies that the required proportion of daily testing would be more than 21 percent of the population – an impossible task – and indicates that attempting to test the entire population at random would be a waste of resources.

Instead we recommend testing smaller stratified specific groups who are at particular risk of spreading the infection. These would include healthcare staff, key workers, and others who are at high risk of creating cross-infection (See Section 6 below). Like Romer, we believe that high frequency testing of such groups would be manageable. But, unlike Romer, we also believe that only targeted testing would be feasible, and that, on its own, such testing would be a means of stopping the virus from spreading in the whole population.

We also recommend that testing of at-risk groups should be done periodically, rather than randomly. This ensures that each member is certain to be tested at regular intervals. This improves the efficiency of testing: for example, it can prevent some individuals being tested on two successive days, and others having to wait a disproportionate length of time to be re-tested. We show that this increases the efficiency of any given number of tests by approximately 35 percent, at no extra cost.

So long as testing remains a scarce and relatively expensive resource, we argue that testing of the general population should be reserved for immunity tests (antibody tests), which would allow those that have been infected to get back to work.

[1] See https://bcf.princeton.edu/event-directory/covid19_04/.

Covid Economics 8, 22 April 2020: 58-84

The paper is organised as follows. In Sections 2 and 3 we use our corrected model to calculate the frequency of random testing that would be required to halt the spread of the virus. We use our model to suggest that Romer's strategy would not halt the spread of the virus and that it is best thought of as a risky "throttle" strategy. In Section 4 we examine the robustness of these conclusions.

Section 5 sets out what we think is a workable testing strategy: testing people at regular intervals in carefully selected groups which have a higher rate of spread of the infection relative to others. This would enable greater rates of isolation in these groups, preventing the epidemic spreading where it matters most. We outline a number of practical suggestions that explain how these groups can be identified. And then in Section 6 we show how to formally calculate the required testing rate to sufficiently reduce the spread of the infection in these groups, when testing is done periodically rather than randomly.

Section 7 presents our conclusions.

In Appendix 1 we set out how and why we think that Romer made his error. Appendix 2 provides a simple formula which can be used to support the conclusions in the body of the paper that Romer's strategy would not control the spread of the virus, and which also provides some simple intuition as to what would be required to achieve such an outcome with universal random testing.

We should make it clear immediately that our paper is not intended as a criticism of Romer. His focus on the critical importance of testing, the ideas presented in his lecture of 3 April 2020, including his emphasis on the need to test key groups very frequently, and his simulations of the spread of the epidemic (which we discuss below) are all immensely valuable. We just want to "check the math" – as Romer has asked us all to do[2] -- and, by doing so, we demonstrate that doing the calculations properly shows that universal random testing should *not* be part of a Covid-19 testing strategy.

2 The basic ideas

The reproduction number R in an epidemic is the expected number of cases directly generated by any one infectious case. The *basic* reproduction number, R_0, is the initial value of R.[3] For Covid-19, Romer takes this value for the whole population to be $R_0 = 2.5$ and so will we.

[2] See Paul Romer's tweet, https://twitter.com/paulmromer/status/1248719827642068992.

[3] One can think of an *original* value for R_0 which is the reproduction number of the infection where all individuals are susceptible to infection, and no policy interventions have been adopted. But of course we also want to allow for the effects of policy interventions which might reduce that rate, such as wearing face masks, undertaking social distancing, or allowing the infection to spread so as to reduce the susceptible population. Such features can be included on our model. One way of doing this is to model these effects explicitly, which is

Covid Economics 8, 22 April 2020: 58-84

We know that an epidemic can only be controlled if the value of *R* is brought below 1. When that happens, each infected person infects less than one new person, and the epidemic will die out.

Romer wants to get the *effective* reproduction number R' (*R* prime) down to $R' = 0.75$ from the basic reproduction number of $R_0 = 2.5$, by randomly testing a fraction of the entire population each day and then isolating those who are found to be positive. In Romer's analysis, which we follow, the effective reproduction number R' is the product of the basic reproductive number, R_0, and the fraction of the infectious population that is not isolated. R' is below R_0 because a fraction of the population is tested each day and those found to be infectious are isolated.

Let φ be the proportion of the infectious population which is isolated. Then we can write Equation (1):

(1) $$R' = (1 - \varphi)R_0$$

For Romer's objective of $R' = 0.75$, Equation (1) implies that that $\varphi = 0.7$, *i.e.* that 70 percent of the infectious population is isolated, and that sufficient tests and isolation are carried out to make this possible.

Drawing on his calculations, Romer suggests that this can be done by randomly testing 7 percent of the population each day, *i.e* testing everybody randomly, at a frequency of about once every two weeks. He believes that this would achieve $R' = 0.75$. With a population of 300 million in the US testing on this scale would require about 20 million tests a day. In the UK with a population of 60 million, this would require about 4 million tests a day. Romer proposes the immediate allocation of $100bn in the US to make such an outcome possible.

This calculation, though, is mistaken. In Appendix 1 we identify the error in Romer's analysis of the relationship between the level of testing that is required to achieve a particular level of φ, and explain how we think that Romer came to make his mistake.

We now set out our alternative calculation which shows that, if random testing followed by isolation were adopted, more than 21 percent of the population would need to be tested each day. This, we show, is what would be necessary to get to a position in which 70 percent of the infectious population were isolated (*i.e.* $\varphi = 0.7$), so that $R' = 0.75$. To do this would, we show, require everyone in the population to be tested about every five days.

how we will allow for the effects of self-isolation. Alternatively one can be implicit, by using an *adjusted* value for R_0 as an input to the calculations, in order to reflect the characteristics of a population, or group of individuals, who have particular attributes – like a high proportion of susceptible people, or indeed recovered people - or are subject to already existing interventions - like lockdown.

3 Calculating the required test rate for universal random testing

3.1 Finding the testing rate which would get R down to *R'* = 0.75

We assume, like Romer, that the whole population is randomly tested. We let t be the proportion of the population tested each day, *i.e.* the probability, for each person, of being tested each day and then isolated. Romer allows for false negatives in his tests. He lets n be the proportion of false negatives and assumes that this is 0.3 (i.e. 30 percent of those who are infected test negative). We follow him in assuming such a high number.[4] In Section 4 we look at the effect of different proportions of false negatives.

We assume that the number of days an infected person is infectious is $d = 14$. Fourteen days is familiar in the analysis of Covid-19 as the period after which the person is either dead or – much more likely – recovered, but no longer infectious. Of course, this number may not be the best one to use. We use $d = 14$ here mainly because it is the number used by Romer as the number of days that each person who tests positive is placed in isolation. We discuss this issue further in Section 4 below and discuss Romer's procedure in Appendix 1.[5]

As a preliminary, let us consider the impact of an infectious person on others. We assume that if $R_0 = 2.5$, and if there were no testing which led to the isolation of infected people, then an infectious person would infect 2.5 people in total, or $2.5/d$ persons per day for d days. That is what we take $R_0 = 2.5$ to actually mean. Note that we discount any idea of a person being more or less infectious during the period of d days. See a brief discussion of this point in Section 4 below.

We now consider the impact of an infectious person on others when he or she has a probability t of being tested each day, and so of being placed in isolation immediately if the test is positive.[6] For clarity, it is helpful to think about the 'discovery rate' x, where $x = t(1-n)$ and n is the number of tests which show false negatives, which we take to be $n = 0.3$ as specified by Romer. The variable x shows the probability that, on any day on which this person is infectious, he or she is isolated. This means that $(1 - x)$ is the probability that this person will not be isolated, and so be able to infect people on the next day.

[4] There is a good reason for this. Massive testing – even of the amount which we contemplate – may well make it impossible to ensure that tests are accurate. Romer rightly argues that, whilst accurate tests are absolutely necessary for clinical reasons when treating an individual person, much more rough-and-ready testing is satisfactory if the purpose of this testing is epidemiological control.

[5] 7 days has been the standard advice for isolation or 14 days if more than one in a household, however recent data shows that the infectious period may last much longer. A recent detailed study of repeatedly tested individuals in Taiwan found a long tail for infectiousness. For further discussion of this point See Section 4 below. See also https://focustaiwan.tw/society/202003260015.

[6] This testing regime assumes that on the day you are tested, if tested positive and isolated, you cannot infect someone else. This simplifying assumption can be thought of in practical terms as a rapid 'early morning' test. In the section below on robustness we investigate the effect of a much more cautious assumption supposing that the infected individual remains infectious for the entire day that they are tested, with corresponding increase in required testing rate t for any given R'.

Covid Economics 8, 22 April 2020: 58-84

In almost all countries, an important part of the current policy response to Covid-19 is that those who show symptoms are required to self-isolate. Because of this, only asymptomatic people and those who have chosen not to self-isolate, for example with mild or mistaken symptoms, will be spreading the disease, and will be included in those who are tested. Romer discussed this issue in his lecture but did not actually allow for it in his calculations.

The proportion of asymptomatic patients isn't really known, and estimates vary wildly. The WHO suggests that as many as 80 percent of cases are asymptomatic or mild.[7] Very inexact data from Iceland suggest that all infectious patients are asymptomatic for the first five days and that, after this time, only about half become symptomatic.[8] We will use these Icelandic figures when constructing the "base case" in our calculations.

We construct our analysis as follows. We want to find the value of *x* that would give Romer's desired value for *R'* of 0.75. We will then work out the required probability of testing *t*, given that we take as given Romer's assumption that the proportion of false negative tests, *n*, is equal to 0.3.

If there were *no* self-isolation of those who became symptomatic, then the probability of an infectious person remaining undetected on day *j* is $(1-x)^j$ and the expected number of infections caused by an infected person on day j is r_j. Therefore, in expectation, an infected person would infect $r_1(1-x)$ persons on day 1 of their infection $r_2(1-x)^2$ persons on day 2, and so on, up to $r_{14}(1-x)^{14}$ persons on the final day *d*. We assume that the infection rate is constant over the infected period and so $r_j = R_0/d$.

We allow for self-isolation of those who are symptomatic as follows. We suppose that all of those infected are asymptomatic for five days (and are subject to random testing during that time). But that, after 5 days, a proportion of the population display symptoms and self-isolate.

Let α be the proportion of population who display symptoms and self-isolate. This means that only a proportion $(1-\alpha)$ go on being tested from day 6 onwards. We can thus write our key equation as follows:

(3a) $$R' = [(1-x) + (1-x)^2 + ... (1-x)^5 + (1-\alpha)\{(1-x)^6 + (1-x)^7 ... + (1-x)^d\}]/d$$

or, in more mathematical notation:

(3b) $$R' = \frac{R_0}{d}\sum_{j=1}^{d}(1-x)^j - \frac{\alpha R_0}{d}\sum_{j=6}^{d}(1-x)^j$$

[7] See https://www.medrxiv.org/content/10.1101/2020.02.20.20025866v2.
[8] See https://edition.cnn.com/2020/04/01/europe/iceland-testing-coronavirus-intl/index.html.

Covid Economics 8, 22 April 2020: 58-84

Our ambition, like Romer's, is to get to a position in which each infectious person on average only infects 0.75 other people, so that the virus will die out. So, we seek to find the value of *x* which would the total number of infections caused by such a person to 0.75. Recalling that $R_0 = 2.5$, and letting $d = 14$ we then can solve for *x* from Equation (4):

(4a) $R_0\,[(1-x) + (1-x)^2 + \ldots + (1-x)^5 + (1-\alpha)\{(1-x)^6 + (1-x)^7 \ldots. + (1-x)^{14\}}\}]/14 = 0.75$

or, in more mathematical notation,

(4b) $\frac{R_0}{d}\sum_{j=1}^{d}(1-x)^j - \frac{\alpha R_0}{d}\sum_{j=6}^{d}(1-x)^j = 0.75$

The solution to this equation can be obtained numerically for various values of α. When α = 0.5, as roughly observed in Iceland, $x \approx 0.146$.

But $t = x/(1-n)$. And *n*, the proportion of false negatives, is equal to 0.3. So the required probability of testing on each day, *t*, is given by $t = 0.146/0.7 \approx 0.209$. That is to say, the proportion of people tested on any day must be as high as 21 percent in order to achieve the required 15 percent discovery rate *x*.

Recall from Equation (1) that $R' = (1-\varphi)R_0$, where φ is the proportion of the infectious population which is isolated. With $R_0 = 2.5$ and $R' = 0.75$, this means that $\varphi = 0.7$. Thus, over the 14 days in which a person is infectious, this person will, on average, be in isolation for 70 percent of the time. We have shown that to achieve such a very striking outcome, the probability of testing an infected person who is not yet isolated on any day must be at least 21 percent. With random, population-wide testing, this means, in effect, that everybody in the population has to be tested about every five days.

3.2 Identifying the testing "threshold" at which *R' = 1*

The dotted line in Figure 1 plots *R'* as a function of the proportion of the population tested. As we have already seen, when half of those who are infected self-isolate from day 6 onwards, *i.e.* $\alpha = 0.5$, *t* needs to be equal to about 21 percent to get *R'* down to 0.75. More than this, the Figure also displays the different testing rates which are required to obtain a range of different values for *R'*. And it does this for different values of α as well.

Covid Economics 8, 22 April 2020: 58-84

Figure 1
Over 20 percent of the population needs to be tested to stop the infection if no-one self isolates when they show symptoms

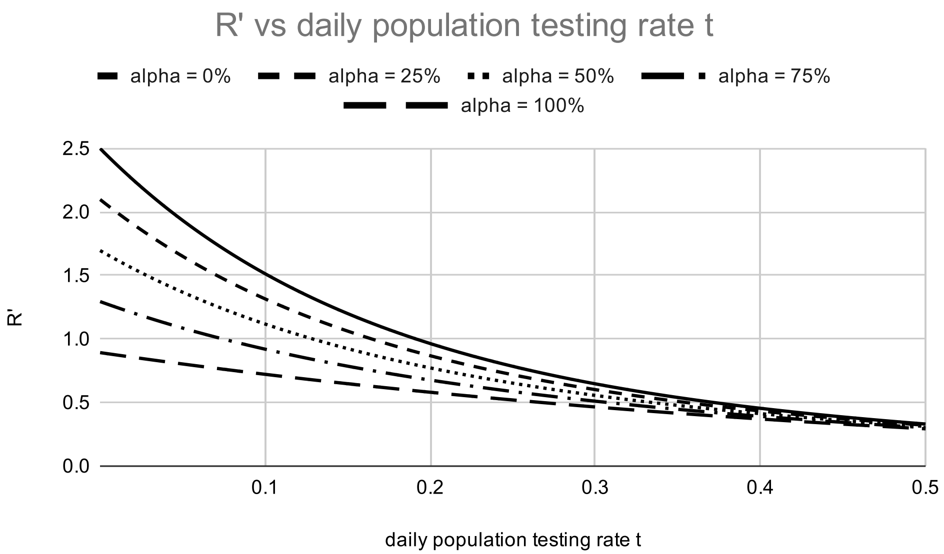

α is the proportion of infected people who display symptoms and self-isolate. Higher levels of α mean that the rate on infection (R') is lower for a given level of testing.

Figure 1 enables us to identify the threshold testing rates, t^* that reduce R' to exactly 1, and so just stop the epidemic from exploding.[9] For α = 0.5, this threshold testing rate is t^* = 13 percent. To achieve this, everyone would need to be tested, on average, every eight days. This is still way above Romer's proposed testing rate of 7 percent. Figure 1 also displays the

[9] It would be good to find a simple way of calculating the threshold testing rates, t^*, for any population, based on the value of R_0 for that population, and given any assumed value for α, without having to solve the complex non-linear model being discussed in this Section.

It turns out that we can do this by using a simple approximation that ignores the dynamics of the infection process, thereby producing an equation which is easy to solve. In Appendix 2, we set out this simple approximate method for calculating the threshold testing rates t^* for a population, depending on the values of R_0.and of α that are appropriate for that population. We show that the approximation is relatively accurate, even although the calculation rests on two simplifying assumptions.

It might be useful to carry out the kind of calculations described in Appendix 2, even despite the fact that, in this paper, we are recommending using periodic testing rather than the random testing being discussed here. This is because, as we show in Section 5 of the paper, periodic testing is more efficient than random testing. As a result, calculating the amount of testing needed to bring R' down to 1, if testing were to be random, might well provide a *lower bound* for the threshold testing rate when using periodic testing. It might be useful to have a simple way of calculating this lower bound, even although the method of calculation is only an approximate one.

Covid Economics 8, 22 April 2020: 58-84

Covid Economics 8, 22 April 2020: 58-84

sensitivity of the results to various values of α, the proportion of infected people who display symptoms and self-isolate. It shows that, at one extreme, when all cases are symptomatic after 5 days and then self-isolate (*i.e.* when α = 1.00), the situation seems just about manageable: the epidemic would stop exploding even without any testing. Nevertheless, the required value of t which would ensure that R' = 0.75, is about 8.2 percent. This percentage is actually a little *above* that proposed by Romer because infected asymptomatic people do a lot of damage in the first five days! At the other extreme, with α = 0.00, the situation is much worse: the probability of testing per day, t^*, which is required to stop the epidemic exploding is now nearly 20 percent (19.1 percent) and the value of t required to ensure that R' = 0.75 is about 26 percent (26.1 percent).

Data for the extent to which infectious people become symptomatic is extremely unreliable and furthermore the range of possibilities seems very wide. Even if over half of infected people self isolate, around 20 percent of the population would need to be tested every day to get R' to 0.75. The outcomes depicted in Figure 1 suggest that, on the balance of probabilities, Romer's strategy of universal random testing would be unworkable.

3.3 Romer's strategy is really a risky "throttle" strategy

Returning to the share of asymptomatic cases observed in Iceland, (*i.e.* the case when α = 0.5), Figure 1 shows that by setting a testing rate of 7 percent of the population, the epidemic remains explosive with R' equal to about 1.3. In this situation, infection would spread rapidly in the earlier stages, since the testing rate is not high enough to slow the spread in a controlled manner; all testing would do is slow down the inevitable spread of the disease. However, ultimately, once the contagion reaches a certain size, the effect of testing, together with the fact that more and more people have had the disease and so are immune, will begin to slow the spread. The proportion of those infected will tend towards a constant level at which there is what has come to be called "herd immunity". It can be shown, for any R, that this proportion is given by $(1 – 1/R)$. Without testing, with a value of R_0 of 2.5, the herd immunity proportion is 60 percent.

Romer's strategy, with a massive amount of testing, would reduce R' to 1.3, and so, using the above formula, it would reduce the herd immunity proportion, to which the population is tending in the long run, to about 23 percent. But in the earlier stages of infection, when there is no immunity, the disease could still spread rapidly as the testing rate would not be high enough to slow the spread in a controlled manner. Thus, in sum, we can say that Romer's testing strategy would slow the spread, and would reduce the level of herd immunity, but would not control the initial explosive phase of the epidemic.

In fact, Romer is really proposing a random testing strategy which could be used as a throttle to control an inevitable spread, along with a view that, at 7 percent testing, this throttle strategy might be 'good enough'. But such a control mechanism does not stop very large numbers of people being infected, it merely "flattens the curve". Of course, the hope is that something else intervenes, like a vaccine.

Romer has posted a very detailed and helpful model of the spread of the epidemic in this manner, one which avoids the problems of the model put forward in his April 3 talk. See https://paulromer.net/covid-sim-part1/ and https://paulromer.net/covid-sim-part2/.

In Romer's simulations of this model, the spread is very rapid: there is a peak of infections of between 3 percent and 18 percent of the population. This would overwhelm any national health system since even 3 percent of the population is an enormous number. There is also a chance that 20 percent of the population would be infected at the same time. That would be a national calamity.

Romer shows in his simulation model that using such a strategy over the course of 500 days might result in about 30 percent of the population contracting Covid-19. This is a risky strategy. Peak levels of infection might rise out of control. Even if testing rates were then increased, lags in responses would mean that the spread of infection would only be gradually reduced. Meanwhile, the virus would go on spreading towards the herd immunity level.

4 The robustness of our conclusion that universal random testing is unworkable

Of course, there are many changes to our assumptions which could modify our calculations.

In particular, there would be significant *reductions* in required testing rates if whole households were to be self-isolated if anyone in the household tested positive. If, for example, only one person in a household were tested at any time, and households consisted on average of two people, then each positive test would remove two people into isolation. That is – testing would become more effective.

On the other hand, our calculations have deliberately assumed a very speedy testing strategy: we have supposed that a test done on any day which finds the person to be infectious causes that person to immediately self-isolate even on that day; an extreme assumption. We have redone the calculations using the more cautious assumption that a positive test result for a test performed on any day does not lead to the person isolating until the next day. The relevant equation now becomes:

(5a) $R' = R_0[1 + (1 - x) + (1 - x)^2 + ... + (1 - x)^4 + (1 - \alpha)\{(1 - x)^5 + (1 - x)^6 .... + (1 - x)^{13}\}]/14$

or, in more mathematical notation.

(5b) $R' = \frac{R_0}{d}\sum_{j=1}^{d}(1-x)^{j-1} - \frac{\alpha R_0}{d}\sum_{j=6}^{d}(1-x)^{j-1}$

These results are *much* worse than those described in Section 3. With α = 0.5, the critical testing rate, t^*, is now 16 percent, and the testing rate required to get R' down to 0.75 is now as high as 27 percent.

Furthermore, we have been assuming, like Romer, that there is uniform contagion throughout the 14-day period. But the medical data shows an asymptomatic infectious period followed by a hump of maximum infectivity as symptoms develop and a tail as symptoms resolve.

We conclude that for the random universal testing proposed by Romer to be workable (involving testing, say, less than 10 percent of the population per day) policymakers would require confidence that:

- i) nearly all infected patients are symptomatic and self-isolate, reducing the burden on testing after the incubation period,
- ii) the tests are sufficiently effective, and complied with, that they capture more than 70 percent of infected cases (and ideally close to 100 percent),
- iii) testing is conducted quickly, early in the morning, and people are isolated on the day of the test, and
- iv) whole households are isolated when any member is infected.

Unfortunately, we do not feel that all these conditions can be met given our current state of knowledge about the virus so we do not believe that whole-population random testing would be a good use of resources.

5 A workable strategy of stratified periodic testing[10]

5.1 Stratified testing

We argue that testing should be carried out at different frequencies for different stratified groups, based on their likelihood of infecting others. This likelihood can be deduced from their occupation, geography, and other factors. Testing at rates above 20 percent per day could be done for carefully selected groups which have a high basic reproduction number (R_0) relative to others. This would enable greater rates of isolation in these groups, lowering their effective reproduction number (R') and helping to prevent the epidemic spreading where it matters most.[11] This appears to be a much lower-risk strategy to contain the spread of infection, and could be done with cheap tests, even if they are somewhat inaccurate.

[10] Some of what follows comes from suggestions made to us by Eric Beinhocker, for which we are very grateful.
[11] It might even enable the general lockdown to be eased, so that other lower-risk groups could keep working and not need to be isolated.

Covid Economics 8, 22 April 2020: 58-84

Broadly, there are two types of people that are likely to have a particularly high basic reproduction number relative to others. The first are those who have a high basic reproduction number to begin with. These are individuals who would have been more likely to infect others, before the infection had begun to spread and any policy interventions had been adopted. Doctors are a one example. They have very frequent and unavoidable close contact with others. Their basic reproduction number will, as a result, be very high and, very frequent testing will be necessary to ensure that their effective reproduction number is low enough. As a result, there could be very frequent testing for doctors in hospitals to ensure that the effective reproduction rate for them is brought well below 1. It appears likely that the relevant calculations will show that doctors actually need to be tested every day. This is something which Romer already suggests in his talk.

But there are also other people that will have a high basic reproduction number because of the uneven application of the lockdown and other factors that might affect variation in infectiousness across groups.[12] For many people, for instance, lockdown means that they are confined to their homes (including many workers who are able to work from home), reducing their basic reproduction number well below 1. But key workers, who are encouraged to keep working in spite of the lockdown, will have a higher basic reproduction number as a result (*e.g.* those involved in food production and distribution). The same will be true for all those who are unable to work from home and are given permission to avoid the lockdown (*e.g.* those involved in construction and manufacturing). Another group that is likely to have a higher basic reproduction number involves those who are more exposed to people who are particularly susceptible to the infection (*e.g.* prison warders and care workers). As testing is rolled out it will become increasingly appropriate to frequently test all such people.

One challenge in all this is that the basic reproduction number may be high in particular groups for idiosyncratic reasons that are hard to anticipate. The kinds of calculation described in the next section could easily be carried out for structured samples in different locations, in order to identify these pockets of infectiousness.

The general principle, then, is that testing should be concentrated in groups that have high basic reproduction numbers relative to others. But this principle should not be interpreted too strictly. In certain cases, other criteria may also be important: for instance, we may want to regularly test groups that interact with those who are more likely to die from the infection (this is another reason to test care workers more frequently) or groups whose absence would have a greater economic impact than others (this is another reason to test key workers more frequently).

[12] See footnote 3.

CEPR PRESS

Covid Economics 8, 22 April 2020: 58-84

5.2 Clarifying how our strategy differs from that of Romer

We hope that Romer would agree with what we have just written above. Indeed, in the version of his plan that he set out on Twitter, Romer has himself provided useful suggestions about who might have priority as tests are rolled out.[13] But from here on we part company.

Romer goes on to suggest that, once tests have been rolled out for these most important groups, there be a further vast expansion of testing, enabling mass random testing *for the whole population,* in order to get the effective reproduction number down *for the whole population.* The version of his plan on Twitter makes this very clear. It concludes as follows:

> "When you strip away all the noise and nonsense, note that once we cover essential workers, it's easy to test everyone in the US once every two weeks. Just do it. Isolate anyone who tests positive. Check your math. Surprise, $R_0 < 1$. Pandemic is on glide path to 0. No new outbreaks. No need for any more shutdowns."[14]

Instead, we argue that testing must be focused on particular groups. This is because our findings, discussed in Section 3 above, show that a mass testing plan would still leave the effective reproduction number significantly above 1 unless it was carried out infeasibly frequently.

Nevertheless, there will still need to be random testing of groups in the population, and some random testing of the whole population. But this testing would be for *informational* purposes only and would only involve testing very small samples of those involved.

Such informational testing will be needed for two reasons. First, random testing of small samples from particular groups will be necessary be to track groups in which the basic reproduction number is already known to be high, and where there is greater potential for a high rate of spread. Once identified, these groups will then need very frequent testing of everyone in the group, for the reasons which we have been discussing in this paper. But testing of small samples of wider groups in the whole population will also be needed to identify new groups where the basic reproduction number is high. As before, this may be for idiosyncratic reasons that are hard to anticipate. Once identified, such groups will then also need very frequent testing of everyone in the group.

But the accuracy of this testing for informational reasons will be determined by the sample size, rather than population size. The samples required for these informational purposes will be *very* small relative to the size of the whole population.

[13] See: https://threadreaderapp.com/thread/1248712889705410560.html.
[14] Romer uses R_0 here to stand for what we call *R'*. See again: https://threadreaderapp.com/thread/1248712889705410560.html

Covid Economics 8, 22 April 2020: 58-84

5.3 Periodic tests rather than random tests

Once the frequency of testing has been decided and testing kits are available, testing can begin for everyone in the identified groups. But it is important that this testing be done periodically for each person, rather there being a random choice of those who are to be tested in each time period.

The rationale underlying periodic testing can be explained by the "waiting-time paradox".[15] Random testing, say of 20 percent of a group each day, wastes many resources. This is because every day some of those tested will have actually been tested the day before, whilst others of those who are infectious will, nevertheless not be tested and so will possibly continue infecting people. By contrast, periodic testing of 20 percent of a group means that, on days 1 to 5, a different fifth of the group will be tested each day, and that on day 6 the first fifth of the group will tested again, and so on. It is clear that this means each person tested will have been tested exactly five days previously, removing the problem that some tests are being wasted and that other tests are being postponed for too long. Because of this, you need to test far fewer individuals in a group to get the same reduction in *R.*

In the next Section, we provide a simple account of this issue, and show how important it is likely to be. We show that with high testing rates, periodic testing beats random testing by a very significant factor. For example, in the model which we examined in Section 3, in the special case in which there is no self-isolation of symptomatic people (*i.e.* with $\alpha = 0$) and perfect testing ($n = 0$), our testing rate required for R' to be 0.75 was 18.3 percent with random testing. With periodic testing this rate falls to 13.5 percent, a 26 percent reduction. This is a big improvement at *no* extra cost.

5.4 The testing system: running two kinds of tests in parallel

In this paper, we have been discussing antigen testing (i.e. testing for active infections) as opposed to antibody testing (i.e. testing for those who have had the disease and are both immune and non-infectious). A combination of the two might be effective and realistic if the testing capacity for active infection remains constrained, but that one-time antibody tests become widely available. One might then proceed as follows:

- Immediately and frequently perform antigen tests on groups and areas with a high R' and immediately isolate those found to be positive. Trace[16] and test the contacts of

[15] The waiting time for a Poisson bus service is twice the waiting time for a periodic bus service with the same rate for a randomly arriving traveller.

[16] Contact tracing can be as simple as testing those in the household and workplace of those who are infected. Technological solutions exist to perform more detailed contact tracing, *e.g.* using mobile phone movements. However, these involve privacy concerns which, crucially, may take time to debate and resolve. The authors believe that simple, immediate testing of infected households and workplaces is preferable to detailed tracking of mobile phone movements at some months delay. As contact tracing apps become more

those who test positive, as these now have a higher probability of also testing positive.

- Self-isolate anyone developing symptoms for a minimum of 7 days. These people would not be tested unless medically necessary. Trace and test their contacts.
- If there were enough tests then one could test people at the *end* of their isolation period to show that they were clear of virus before they were allowed to come out of isolation.
- Widespread home kit antibody testing for anyone to see if they had had the virus - these would be one-off tests that would not need to be repeated.
- A system to track people with immunity who could then circulate freely if they had either a) had a positive antibody test, b) had a positive active infection test more than some specified number of days ago, or c) had a negative active infection test after their symptoms resolved.

All of this could be done using cheap antigen tests, even if they were somewhat inaccurate. It is not the case that "no test is better than an unreliable test". Our calculations show that, whilst accurate antigen tests are absolutely necessary for clinical reasons when treating an individual person, much more rough-and-ready testing is satisfactory if the purpose of this testing is epidemiological control through isolation. (In our baseline calculations discussed below, for instance, we assume that 30 percent of infected people wrongly test negative, $n = 0.3$).

Doing all of this will help governments to track spread and to determine where hotspots are flaring up. Such information will help them to work out how to selectively tighten, or loosen, containment measures when needed.

Such a testing procedure would involve doing two things at once: the stratified periodic antigen testing which we have been discussing would be designed to damp the spread of the disease in key groups, by catching those in these groups who were infectious but asymptomatic, or pre-symptomatic, or post-symptomatic, and so not self-isolating. At the same time, antibody testing for the entire population would separate out the immune population; passing an antibody test would enable such people to return to work.

6 Calculating the required test rate with random testing[17]

6.1 Finding testing rate which would get R down to R' = 0.75 using periodic testing

In this Section we show how to solve for the required testing rate when there is periodic testing. As in our discussions of random testing in Section 3 we aim to find the testing rate which would get R down to R' = 0.75, and also to Identify the testing "threshold" at which R'

widespread the degree of contact (separation distance, length of interaction, *etc.*) will be available to help target testing resources.

[17] We are grateful to Frank Kelly for his assistance in preparing this Section.

Covid Economics 8, 22 April 2020: 58-84

= 1. One of our aims is to show how much more effective periodic testing might make the testing process, when compared with the random testing process discussed in Section 3.

As in that discussion of random testing, our first objective was to get the value of R down from R_0 = 2.5 to R' = 0.75. In our discussion here, we will include the effects of self-isolation, *i.e.* the results in Section 3 with which we will compare our findings here are those in which α = 0.5. Our results suggest that periodic testing might be about 37 percent more effective than random testing, at no extra cost.

As in Section 3, we let R' be the expected number of people that a randomly chosen infected person infects before that person is positively tested (or stops being infective, if sooner). Let r_j be the expected number of individuals infected by an individual on day of his/her infection, for j =1,2,...,d where the length of infectivity is d. Thus $R' = \sum_{j=1}^{d} r_j$ Now suppose an individual is tested every N days, and that for high risk groups we test very frequently, so that $N < d$. If the time of the infection is random and the individual is infective for the day of the test (as we considered in more likely in our robustness analysis in section 4), then

(6) $$R' = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{i} r_j$$

From this we can deduce that

(7) $$R' = \frac{1}{N}\sum_{j=1}^{N}(N - j + 1)r_j$$

Here we introduce our self-isolation factor α into Equation (7) by noting that r_j during the infectious and asymptomatic period (first 5 days), which we denote as $d_0 = 5$, is still $r_0 = \mathrm{R}_0/d$ and thereafter $r_\alpha = (1-\alpha)\,\mathrm{R}_0/d$ for the infectious period in which a person may self-isolate. Substituting these values into r_j we have three alternative situations to consider: (i) the case when $N < d_0$– the testing rate is more frequent than the move to becoming symptomatic that occurs after day d_0 and therefore Equation (7) represents this outcome entirely, (ii) the case when $d_0 < N < d$ for which the testing period N is between 6 and 13 days inclusive, and finally (iii) the case when $N \geq d$.

The following is for the case when $d_0 < N < d$ which is the testing range that tends to deliver an $R' < 1$ for this model of testing

(8) $$R' = \frac{1}{N}\sum_{j=1}^{N}(N - j + 1)r_j = \frac{1}{N}\sum_{j=1}^{d_0}(N - j + 1)r_0 + \frac{1}{N}\sum_{j=d_0+1}^{N}(N - j + 1)r_\alpha$$

Covid Economics 8, 22 April 2020: 58-84

We take $r_\alpha = (1-\alpha)R_0/d$ to the daily likelihood an infected individual will infect another person during the period in which they may become symptomatic and self-isolate with probability α. This gives[18]

(9) $$R' = \frac{R_0}{2Nd}[(N+1)N - \alpha(N - d_0 + 1)(N - d_0)]$$

We can now compare our findings with those in Section 3, for the case in which α = 50 percent. We found there that that to bring *R* down from $R_0 = 2.5$ to $R' = 0.75$ would – with random testing and next day test results - require a testing rate of 27 percent of the population a day, (or about one test every four days) . But those results assumed that 30 percent of tests failed. If tests were perfect those results imply an equivalent correct identification rate of 0.19 per day (or a 100 percent accurate test about every five days).

We now use Equation (9) to solve for the value of *N,* the number of days between each test, that is required to bring *R* down from $R_0 = 2.5$ to $R' = 0.75$ when there is periodic testing. This equation shows that the period between testing for each individual would need a test rate of 17 percent or a test every 6 days[19]. This assumes a test that has the same 30 percent false negative rate to estimate the periodicity of the test. If these tests were perfect, then the test rate could be every 8.2[20] days or an identification rate of 0.12 - that is a test with a correct identification rate carried out on 12 percent of the population per day.

This is a thirty-seven percent reduction in the rate of testing required, as compared with the case of random testing. We can see that doing random testing would provide a big improvement at *no* extra cost.

6.2 Identifying the testing "threshold" at which R' = 1

The dotted line in Figure 2 plots R' as a function of the proportion of the population tested.[21] As we have already seen, when half of those who are infected self-isolate from day 6 onwards, i.e. α = 0.5, t needs to be equal to about 17 percent to get R' down to 0.75.

[18] For the case N < d_0 the solution remains unchanged as the frequency of testing is above that which would allow the self-isolation process to occur: $R' = R_0\frac{N+1}{2d}$
and for the case $N \geq d$:
$$R' = R_0\left[1 - \frac{1}{2N}(d-1)\right] - \frac{\alpha R_0}{2dN}(d-d_0)(2N - d0 - d + 1)$$

[19] We make here a simplifying assumption that in order to reproduce an effective testing rate x given the false negative rate of a test n, that the required *t* is *t = x/(1-n).*
[20] Of course, in reality such a number would need to be rounded up or down to a full number of days.
[21] It is interesting to note that when N<d_0 we are testing such that the virus spreads only for N days precisely (the time between tests) which means that R' changes linearly when we reach this high frequency level of testing.

More than this, the Figure 2 also displays the different testing rates which are required to obtain a range of different values for R'. And it does this for different values of α as well. Figure 2 enables us to identify the threshold testing rates, t* that reduce R' to at which the disease exactly 1, that is to the value which divides outcomes in which the epidemic dies out from outcomes in which it explodes

For α = 0.5, this threshold testing rate implied to bring R' to 1 is t* = 11 percent – testing 11 percent of the population each day. To achieve this, everyone would need to be tested every nine days. So with periodic testing of an entire population with good isolation of infected individuals, the population of infected individuals may maintain a stable size (R'=1).

Figure 2 also shows us how sensitive this testing rate is to the effectiveness of a population's self-isolation when infectious. For testing rates that are below 20 percent per day (*i.e.* less frequent than once every 5 days), the period in which a person is assumed to be asymptomatic) we can clearly see the risk of those who are infectious not successfully self-isolating.

This is an important figure as α may change from population group to population group. For example, contract workers who are paid by the hour have a direct incentive to ignore symptoms. This would translate to a lower α for this group and consequently a much higher R' for any rate of testing. Workers in this category who also have high contact rates as part of their job would therefore be expected to require the highest rate of testing.

In the case in which all individuals become symptomatic and so self-isolate after 5 days (when α = 1.00) R' ≈ 1. This means that the pandemic doesn't necessarily die out and therefore to bring this down to zero, by bringing R' to 0.75 we need a 'perfect' (no false negatives) test every 12 days. As we show in the following section, if we have same day testing and a perfect population and test, we can reduce the testing rate to every 19 days.

Figure 2
With periodic testing the proportion of the population checked each day can vary depending on how likely it is that infected people self-isolate.

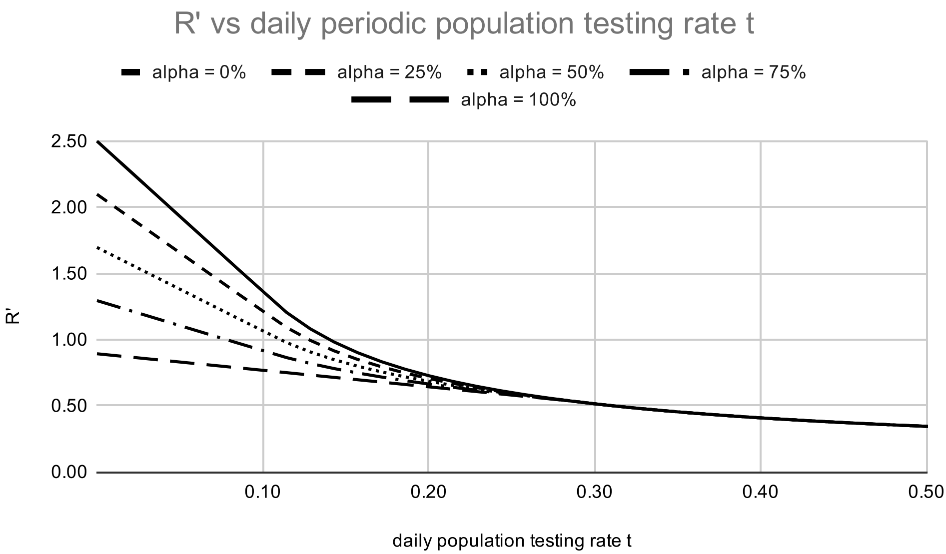

α is the proportion of infected people who display symptoms and self-isolate. Higher levels of α mean that the rate on infection (R') is lower for a given level of testing.

6.3 The impact of an instant test

In Section 3 we started by looking at a theoretically perfect test for which a positive identification of an infected individual would result in them not being able to infect anyone that day. Whilst this instantaneous test is unrealistic, it is useful to compare the results of our periodic model against those of the model in section 3.

We can solve the model introducing a delay with similar results[22]; Again, the impact of differing self-isolation likelihoods is prevalent when the testing rate falls much below 20

[22] To model the instant test we adjust our model to remote the anticipated extra day of infection that would have otherwise occurred, deprecating the sum used in earlier sections from j=1 to i, to j=– 1 to i-1. For completeness we now explicitly set out the equations for these calculations explicitly.

(a) For low-risk individuals we can test much less frequently, so that $N \geq d$. Then we obtain:

$$R' = \frac{1}{N}\left(\sum_{i=1}^{M}\sum_{j=1}^{i-1} r_j + (N-d)R_0\right)$$

and hence

$$R' = R_0 - \frac{1}{N}\sum_{j=1}^{d} j r_j$$

So, if $r_j = R_0/d$ we can solve for the required rate of testing. It is

$$R' = R_0\left[1 - \frac{1}{2N}(d+1)\right] - \frac{\alpha R_0}{2dN}(d-d_0)(2N - d0 - d - 1)$$

percent. However, as the tests are now instant, the amount of testing in order to reduce R has fallen. For α = 0.5 the required testing rate to achieve a value of R = 0.75 is now every 7.7 days (t = 13.4 percent) as opposed to every 6 days for a test that gave the results a day later. It is interesting to note that the effect of an unreliable test (30 percent false negatives versus no false negatives) is similar in scale to the impact of having to wait for a day for the results of a test (so that there is an additional day on which the person can spread the infection). The figures show an effect of 13 percent for an unreliable test versus 12 percent for a perfect test which would deliver results a day later.

Figure 3
If the results of the test are known instantly then less of the population needs to be checked every day to reduce R' to a given level

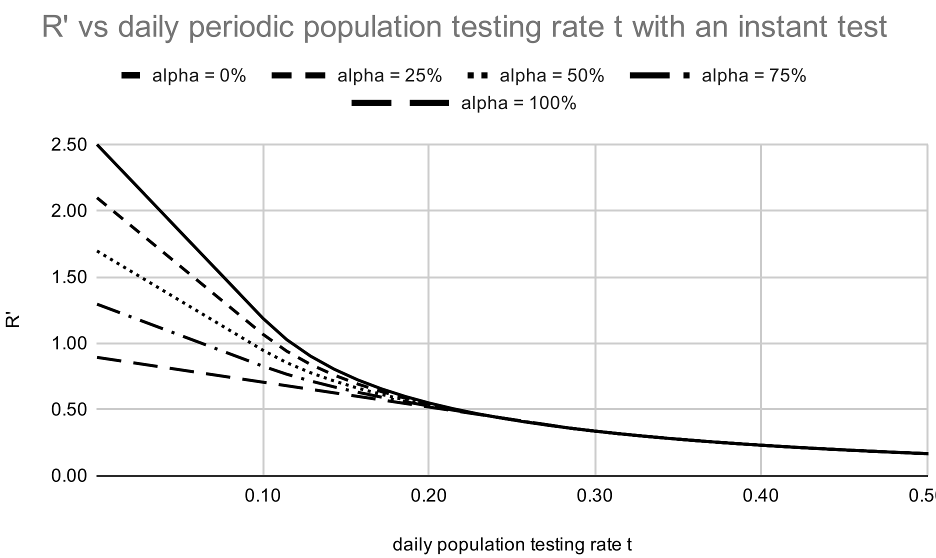

α is the proportion of infected people who display symptoms and self-isolate. Higher levels of α means that the rate on infection (R') is lower for a given level of testing.

(b) Likewise for testing when $d_0 < N < d$ we now sum to i-1 as opposed to i:

$$R' = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{i-1} r_j$$

Then;

$$R' = \frac{1}{N}\sum_{j=1}^{N}(N-j)r_j$$

Giving:

$$R' = \frac{R_0}{2Nd}[(N-1)N - \alpha(N - d_0 - 1)(N - d_0)]$$

(c) For high frequency testing for when $d_0 > N$ and therefore α is not relevant to the dynamics:

$$R' = \frac{1}{N}\sum_{j=1}^{N}(N-j)r_j$$

$$R' = \frac{R_0}{2d}(N-1)$$

Covid Economics 8, 22 April 2020: 58-84

6.4 Allocating testing over populations with different R_0

If scarce testing is to be allocated over individuals with different prior probabilities of infection, these formulas can be used to optimize the allocation.

If we take the simplest form of our periodic testing model, when α = 0, for high frequency testing where an individual has a test after a smaller number of days than the length of an infection (N<d), then equation (7) simplifies to:

(10) $$R' = \frac{R_0}{2d}(N+1)$$

Suppose that individual *k* has prior probability p_k of infection and that we wish to choose N_k so as to allocate a given amount of testing over a set of individuals so as to maximally reduce *R'*. Then, using frequent testing where an individual has a test after a smaller number of days than the length of an infection (N<d), the optimal allocation of a given amount of testing should choose $N_k \propto 1/\sqrt{p_k}$.

This shows very clearly that there are diminishing returns to making N very small – that is conducting more frequent tests results in increasing but diminishing returns to R'. The amount of testing allocated, therefore, to higher risk individuals is naturally higher, but not proportionately so.

6.5 Summary

Periodic testing periods of around 6 days may be sufficient to control the propagation of infection on their own. This is testing approximately 17 percent of the population each day. This figure is highly sensitive to changes in self-isolation behaviour and the characteristics of the test.

This model moves us towards a useful framework for determining the frequency with which a population need be tested in order to treat that group's R_0. For any particular population, the frequency of testing which is required to reduce transmission sufficiently is determined by the both the population's properties and the properties of the test. A population's key factors are the capability of its individuals to self-isolate if symptomatic, and the initial propagation rate R_0. The important figures for the test being administered is the speed of results and the accuracy with which it reports positive results. All of these can be traded off against one another using our model.

7 Conclusions

This paper proposes 'stratified periodic testing' as a strategy for ending the lockdown and returning economies to work, while preventing an explosive re-emergence of Covid-19 (keeping R'< 1). The testing would be 'stratified', in the sense that it would focus on specific

Covid Economics 8, 22 April 2020: 58-84

Covid Economics 8, 22 April 2020: 58-84

subsets of the population who currently have the highest basic reproduction number R_0. The criteria could be amended to also take into account the vulnerability of subgroups, or the loss of economic activity if they are forced to self-isolate at home. The testing would also be 'periodic', in the sense that each member of the subset would be tested at regular, defined intervals, rather than testing within the group being done at random. This ensures that infected people can be identified and isolated quickly. Those who test positive would be quickly isolated at home, as would anyone with symptoms. The tests need not be perfect: if they are cheap but deployed widely within particular groups, then false negatives can be offset by the scope of testing. The effectiveness of the program can be improved by simple tracing of the contacts of those infected: for example testing those in their workplaces and households, rather than fully tracking mobile-phone movements with the associated privacy concerns.

We argue that this is better than 'universal random testing' which is currently being discussed globally. Romer suggests that by testing 7 percent of the population every day we can get the effective reproduction number of Covid-19 to around 0.75 and curb the epidemic. Unfortunately, these calculations contain errors. By correcting this method, and using reasonable assumptions about asymptomatic carriers, we believe that at least 21 percent of the population would need to be tested each day to get the effective reproduction number well below 1 (*i.e.* to the value of 0.75). For obvious reasons, we do not see this as a feasible population-wide strategy.

Any testing strategy should be thought of as a complement to other measures that can reduce the spread of Covid-19 at little economic cost. For example, those that can work from home with little loss of productivity should continue to do so, retirees should continue to self-isolate, and people in public places should wear masks and regularly wash their hands. Stratified periodic testing can then help those sectors in the economy that cannot operate from home get back to work quickly and safely. This should continue until widespread vaccines or treatment for the virus are available

Appendix 1: **Romer's analysis of testing**

We now explain why we think there is a mistake in the way in which Romer calculates φ, the proportion of the infectious population which is isolated. We then present our attempt to understand how and why he made his error.

A.1.1 Romer's analysis

Romer assumes random testing of the whole population. In his calculations, he lets t be the proportion of the population tested each day, *i.e.* the probability that each person is tested each day. He supposed that t = 0. 07.

Romer allows for false negatives in tests. He lets n be the proportion of false negatives. Romer assumes that this proportion is 0.3.

Romer lets l be the number of days that each person who tests positive is placed in isolation. He assumes $l = 14$.

Romer then computes φ, the proportion of the infectious population which is isolated, as follows. He writes something similar to, but not the same as what we have called Equation 1 in our paper.

(1) $\varphi = t(1 - n)l$

Just to be clear, this equation here comes directly from the slides which accompanied Romer's talk.[23] We have no background on why he wrote down this equation, and we think that it is incorrect. We say this because in our paper above Equation (1) reads $\varphi = t(1 - n)d$. This has the variable d on the right-hand side, showing the number of days for which a person remains infections. By contrast Equation (1) above has l on the right-hand side the variable l which is the number of days that each person who tests positive is placed in isolation. The nature of what we think is Romer's error is discussed immediately below.

Since $t = 0.07$, $(1 - n) = 0.7$ and $l = 14$ Romer claims that $\varphi = 0.69$. This value of $\varphi = 0.69$ would, he says, produce his desired value for R', since:

$R' = (1 - \varphi) R_0$, or $R' = (1 - 0.69) \times 2.5 \approx (1 - 0.7) \times 2.5 = 0.75$.

Drawing on these calculations, Romer suggests that there should be testing of 7 percent of the population each day.

Notice that, although Romer mentioned self-isolation of those who have symptoms in his lecture, there is no allowance for such an action in any of the calculations in his slides.

A.1.2 Our Criticism

It is helpful to try to understand how Romer made what we think is an error.

To see most clearly, and simply, why his calculation cannot be right, imagine what would happen if there were to be double the amount of testing proposed by Romer, *i.e.* suppose that $t = 0.14$. Then, using his formula for φ we would get $\varphi = t(1 - n)l = 0.14$ times 0.7 times 14 = 1.38; the person would be in isolation for more than all of the period of 14 days! So the equation must be wrong.

How can we understand the inclusion of 'l', the number of days that an infected person is placed in isolation, on the right-hand side of this equation? One possibility is that the inclusion of 'l' is simply a mistake. Romer states that $\varphi=t(1\text{-}n)l$, but φ and $t(1\text{-}n)l$ appear to be very different things. Because $t(1\text{-}n)$ is equal to the probability that an infectious person is put into isolation on any day, it follows that $t(1\text{-}n)l$ is equal to the expected length of isolation any infected person is likely to face, after one round of testing. But φ is the fraction of the infected population that are isolated -- which is clearly not the same thing as the

[23] See minutes 16 to 20 of the Romer talk, and the accompanying slides.

Covid Economics 8, 22 April 2020: 58-84

expected length of isolation any infected person is likely to face, $t(1\text{-}n)l$. So it appears that stating $\varphi=t(1\text{-}n)l$ is a mistake.

Another possibility is to make a set of assumptions about Romer's set-up that bring the meaning of φ and $t(1\text{-}n)l$ closer together. For instance, consider the following approach. First, interpret '*l*' as the 'number of days that an infected person is infectious', rather than 'is placed in isolation'. Secondly, define '*Z*' as the number of infected people not in isolation. Thirdly, imagine there are '*l*' periods where, in each period, a fraction $t(1\text{-}n)$ of those infected people not in isolation, *Z*, are removed and put into isolation. And finally, assume that in each period the number of infected people not in isolation, *Z*, remains the same (i.e. the infected who are put into isolation are replaced with newly infected people). Then it follows that, after '*l*', periods, $t(1\text{-}n)l *Z$ people will be in isolation. Now, $t(1\text{-}n)l$ is indeed equal to φ, the proportion of the infected population not in isolation who are put into isolation – but with two very significant caveats. First, it assumes that everyone who will be isolated over the '*l*' days is isolated on the first day. And secondly, because *Z* is constant over time, it follows that φ may also be greater than one if *t* or *l* is large enough, or *n* is small enough – which, as shown before, is exactly the problem with Romer's analysis.[24]

Appendix 2 A simple method for calculating the testing threshold when testing is random

In this Appendix we set out a simple method for calculating the threshold testing rate, for random testing, which would reduce the effective reproduction number, *R'*, to the value at which the disease does not die out. The calculation employs a simple approximation which ignores the dynamics of the infection process. If we ignore the dynamics, we do not have to solve a complex equation like Equation (3) which sums a number of effects in a non-linear way, over many time periods.

Our method of calculation builds on the following insight: for *R'* to be less than 1 when there is universal random testing then, on any given day, a person with Covid-19 is more likely to go into isolation than to spread it to someone else. Relying on this insight, we can ignore the dynamics of the process and simply solve for the value of *t* for which this condition will hold. We proceed in two steps.

(a) For simplicity, we first examine the extreme case in which none of those who are infected become symptomatic and self-isolate; this corresponds to the case considered in Section 3 in which $\alpha = 0$.

[24] Intuitively, the problem here is that you are taking a fraction $t(1\text{-}n)$ of the infected population not in isolation, *Z*, and putting them into isolation in each period – but because *Z* replenishes over time, if you isolate a large enough proportion of Z, $t(1\text{-}n)$, enough times, *l*, then you will end up with more infected people in isolation, $t(1\text{-}n)l*Z$, than there are infected people not in isolation, *Z*.

Covid Economics 8, 22 April 2020: 58-84

Consider any group of z, as yet unidentified, infectious people. Assuming that this group is a small fraction of the overall population, the number of people who will be infected by this group on any given day is (R_0/d) times z, where d is the number of days that an infectious person remains infectious.

The number of these z people who, on this same day, will go into isolation because they have tested positive will be $t(1-n)$ times z. But there will be additional infectious people who cease to infect others because, although they did not test positive on that day, the period during which they had the disease and were infectious will have come to an end. This happens with probability $1/d$; so there will be $\{[1-t(1-n)]/d\}$ times z such people[25]. Thus, for R' to be less than 1, we require that:

(1) $\{t(1-n) + [1-t(1-n)]/d\} > R_0 / d$.

This means that, for this extreme case, the threshold testing rate is given by

(2) $t^* = (R_0 - 1)/[(d-1)(1-n)]$

If $R_0 = 2.5$, $d = 14$, and $n = 0.3$, we get $t^* = 16.5$ percent . That is, this method says that, to get R less than 1 by randomly testing the whole population, one needs to test at least 17 percent of the population. That is, the threshold testing rate, t^*, is 17 percent.

This is a lower value than what we found for t^* using the full dynamic model in Section 3 when $\alpha = 0$. The result there was that $t^* = 19.1$ percent. The discrepancy between these two results arises precisely because of the dynamic process of the epidemic: the simple calculation carried out here ignores the fact that, as time passes, testing will remove some of the infected people, so that they are no longer available to be tested on later days. That is what made Equation (3) so complex.[26] For this reason the result produced using this method will always underestimate the required testing rate. This simple method thus provides a (quick and dirty) lower bound for the true value of t^*. Nevertheless, the fact that this calculation is so simple, and the intuition provided by thinking about the problem in this way, may make it useful to carry out this calculation.

(b) This calculation can be readily extended to include the more general cases considered in Section 3 in which a proportion of those who are infected become

[25] This $1/d$ probability is the chance that an infected person becomes non-infectious independently of testing. An intuitive way to think about this is that, in choosing someone at random, there is a $1/d$ chance that that person is on their last day of infection and so will become non-infectious the following day. This is only an approximation since it requires that the value of R' resulting from testing is equal to 1. That is because if $R'>1$ then the virus would be spreading and hence an individual would be less likely to be on their last day of infection; conversely if $R'<1$ then a higher proportion of the infected population would more likely to be about to end their infectious period.

[26] It is possible that this is what Romer was effectively assuming in his analysis. See the final paragraph of Appendix 3.

symptomatic after a certain number of days and so self-isolate. Consider here a proportion α who self-isolate after a number of days d_0 out of the total number of days of infection d. We can approximate an adjusted value of α, which we call α'. This is the probability that someone who is infected self isolates on any particular day (independently of any test or of reaching the end of their infectious period), such that at the end of the infectious period the chance the individual has self-isolated is α, namely;

(3) $\alpha' = \alpha/d$

We then introduce α' in Equation (1) to create a new condition that the testing rate, t, must satisfy for a population following this self-isolation rule:

(4) $\{t(1\text{-}n) + [1\text{-}t(1\text{-}n)]/d + (1\text{-}t(1\text{-}n))\ (1\text{-}\alpha')/d\} > R_0/\ d.$

This means that the threshold testing rate is now given by

(5) $t^* = (R_0 - d\alpha' - 1 - \alpha') / [(d - d\alpha + \alpha' - 1)(1\text{-}n)]$

Suppose, as in the previous case, that $R_0 = 2.5$, $d = 14$, and $n = 0.3$. Then for a population for whom the first 5 days are asymptomatic, who then become symptomatic and self-isolate with probability $\alpha = 0.5$, Equation (4) produces vale for α'. This leads, using Equation (5) to a correspondingly reduced threshold testing rate of $t^* = 11.0$ percent. In other words, this method says that, to get R' less than 1 by randomly testing the whole population, one would only need to test 11 percent of the population because some of the population will self-isolate after 5 days.

This is a smaller value than what we found for t^* in Section 3, in this case with $\alpha = 0.5$, using the full dynamic model. The result there was that $t^* = 13$ percent. A discrepancy between these two results arises partly for the same reason that it did in the case in which $\alpha = 0$: the simple calculation carried out in both cases ignores the fact that, as time passes, testing will remove some of the infected people, so that they are no longer available to be tested on later days. But in addition, in this case here with $\alpha = 0.5$ we are making a second simplifying assumption, that there is a ‘random’ self-isolation process α' each day, such that at d = 14 days the chance that someone has self-isolated is exactly equal to α. This is as opposed to the detailed model in Section 3 in which, after day 5, there is a step of size α in the chance of someone self-isolating, so that nobody self-isolates for the first 5 days and then a proportion α self-isolate for a 9 day period with certainty.

Nevertheless, despite these two simplifying assumptions, it might still be useful to carry out the simple calculation described here, even despite the fact that, in this paper, we are recommending using periodic testing rather than the random testing being discussed here. This is because, as we have shown, periodic testing is more efficient than random testing. As

Covid Economics 8, 22 April 2020: 58-84

a result, calculating the amount of testing needed to bring R' down to 1, if testing were to be random, might well provide a *lower bound* for the threshold testing rate when using periodic testing. It might be useful to have a simple way of calculating this lower bound, even although the method of calculation is only an approximate one. We say this in the light of the current uncertainty about the true value of α in populations and the strong impact that this will have on the infection rate. In such circumstances, it seems helpful to have a calculation for *t** which one can carry out quickly for different values of α, without having to solve for *t** over and over again, using the complex non-linear approach presented in Section 3.

Covid Economics 8, 22 April 2020: 58-84

Covid Economics **Issue 8, 22 April 2020**

Working from home across countries[1]

Charles Gottlieb,[2] Jan Grobovšek[3] and Markus Poschke[4]

Date submitted: 14 April 2020; **Date accepted:** 15 April 2020

We study how the share of employment that can work from home changes with country income levels. We document that in urban areas, this share is only about 20% in poor countries, compared to close to 40% in rich ones. This result is driven by the self-employed workers: in poor countries their share of employment is large and their occupational composition not conducive to work from home. At the level of the entire country, the share of employment that can work from home in poor countries compared to rich countries depends on farmers' ability to work from home. This finding is due to the high agricultural employment share in poor countries.

1 We thank Joel Frischknecht and Joern Onken for outstanding research assistance. Research funding from Leverhulme, GFF Fund of the University of St. Gallen and ESRC-DFID (ES/L012499/1) is gratefully acknowledged. All errors are ours.

2 Assistant Professor of Economics, University of St Gallen.

3 Assistant Professor of Economics, University of Edinburgh.

4 Associate Professor and William Dawson Scholar, Department of Economics, McGill University.

Covid Economics 8, 22 April 2020: 85-105

Covid Economics 8, 22 April 2020: 85-105

1 Introduction

Many countries are implementing drastic measures of social distancing to tame the spread of COVID-19. These measures often involve closure of workplaces to limit interpersonal contact. While they are in place, work can only continue if it can be conducted from workers' homes.[1] The extent to which work can be conducted from home therefore is a key factor determining the economic consequences of social distancing policies.

The ability to work from home (WFH) has been measured for the United States (Dingel and Neiman, 2020*a*, henceforth DN) and for a set of European countries (Barrot et al., 2020; Boeri et al., 2020). These authors have found that around 40% of jobs could potentially be carried out from home.[2] Evidence on the ability to work from home in poorer countries is more scant, with the exception of two papers, which we discuss below. Such evidence is particularly timely and valuable as some low-income countries have started to implement social distancing policies. We put a particular emphasis on how differences in the economic structure across countries contribute to differences in the ability to work from home.

The starting point of our analysis are the occupation-level measures of ability to work from home computed by Dingel and Neiman (2020*a*). We combine these measures with information on the distribution of employment over occupations across countries to obtain measures of the aggregate ability to work from home by country and by country income group. We obtain this information from a micro level dataset we built, which consolidates information from labor force and household surveys for 612 country years for 57 countries.[3] The key advantage of this data is that it allows for the analysis of detailed subgroups. This is important, since lockdown policies affect such groups very differently.

Our main analysis focusses on urban areas. We find that the ability to work from home is significantly lower in poor countries. Only about 22% of workers can work from home, in contrast to 37% in rich countries.

We then investigate the extent to which this conclusion is driven by two particularities of the employment structure in poor countries. First, we show that the lower ability to WFH in poor countries is particularly pronounced for the self-employed. For wage and salary workers, WFH

[1]Exceptions consist in sectors considered to be essential.

[2]Hensvik et al. (2020) find that in the US, the share of workers who actually worked from home in 2011 to 2018 is around 15%, with substantial variation across occupations.

[3]Table 4 provides an overview of all the data sources.

ability in poor countries is not far below that in rich countries. This implies that the large share of self-employment in poor countries contributes to the low WFH ability in these countries. We verify that this is also the case when we use a new measure of the ability to run a household enterprise from home, which we compute using data from the Indonesia Family Life Survey.

Second, we go beyond urban areas and compute measures of WFH ability at the level of the entire country. Due to the predominance of agricultural employment in rural areas of poor countries, the WFH ability of farmers crucially affects our findings here. If farmers are assumed to have a negligible ability to work from home, as indicated by DN's measure, the gap in WFH ability between poor and rich countries is even larger, 15 and 35%, respectively. If, in contrast, we assume that all farmers can work from home, the aggregate WFH ability in poor countries in fact exceeds that of rich countries.

In summary, the share of workers in urban areas who can work from home is clearly lower in poor countries. This result is principally driven by urban self-employed workers. At the level of the aggregate economy, poor countries may or may not have lower ability to work from home, depending critically on the WFH ability of farmers. A lower ability to work from home implies a greater potential cost of social distancing policies.[4] The trade-off between the costs and benefits of such policies might thus be different in low-income countries. The existing literature has pointed out several other reasons why the trade-off may differ across countries (Mobarak and Barnett-Howell, 2020; Loayza and Pennings, 2020). Our findings constitute an additional factor. They also point to a particularly important role of self-employment and agriculture.

Related literature. We are aware of two other efforts to build WFH ability measures for poor countries. Dingel and Neiman (2020b) combine their WFH measures with ILO data on the distribution of occupations across countries. They find that the share of employment that can be done from home is significantly lower in poor countries. Saltiel (2020) analyzes data for urban areas in ten developing economies. Using a country-specific measure of WFH ability, he finds a similar cross-country pattern. He also investigates how the WFH ability is related to individual characteristics. While we find similar results to this work at the aggregate level, our analysis also allows us to point out the main sources of differences in WFH ability across countries.

[4] Our analysis does not address additional factors that might reduce the ability to work from home in poor countries even further, in particular the digital infrastructure. See e.g. Chiou and Tucker (2020).

Covid Economics 8, 22 April 2020: 85-105

2 The distribution of occupations across countries and the ability to work from home

In this section, we measure the share of employment that can be done remotely, across countries of different levels of income per capita. To do so, we use the classification by Dingel and Neiman (2020*a*) to measure the share of jobs that can be done from home for each ISCO-1 level occupation.[5] As in DN, the share of WFH jobs refers to the fraction of detailed occupations within a broad occupation group that can be done from home. The measure is computed based on characteristics of each occupation. It does not depend on the distribution of employment in the United States.

Table 1 shows that the ability to work from home differs very strongly across broad occupation groups. In managerial and professional occupations, the majority of jobs could be carried out from home, at 76.8 and 70.6%, respectively. In contrast, very few elementary occupations or occupations involving plant or machine operation (common in manufacturing) can be done remotely. In particular, 96.1% of craft or trade occupations are tied to the location of the activity. The ability to work from home in services and sales occupations is also relatively low.

Table 1: Percent of detailed occupations that can be done from home by main occupation category

	Occupation, ISCO 1 digit	WFH (in %)
1	Managers	76.8
2	Professionals	70.6
3	Technicians and Associate Professionals	39.6
4	Clerical Support Workers	49.6
5	Services and Sales Workers	20.7
6	Skilled Agricultural, Forestry and Fishery Workers	8.3
7	Craft and Related Trades Workers	3.9
8	Plant and Machine Operators and Assemblers	7.4
9	Elementary Occupations	9.6

Note: We take the classification based on ONET data provided by Dingel and Neiman (2020*a*) and use a cross-walk to the ISCO-1 classification.

The distribution of employment across occupations varies significantly with economic development. We show this using a dataset we built combining household surveys and labor force surveys from 57 countries, covering 612 country years. The total sample size approaches 18 million

[5] This is the level of aggregation at which occupation data can be harmonized across countries. DN's measure is reported using the SOC occupation classification. We use a crosswalk to map this into the ISCO classification. We report WFH shares from DN at the ISCO-2 level in the Appendix, Table 5.

Covid Economics 8, 22 April 2020: 85-105

Figure 1: Distribution of occupations by country income level, urban areas

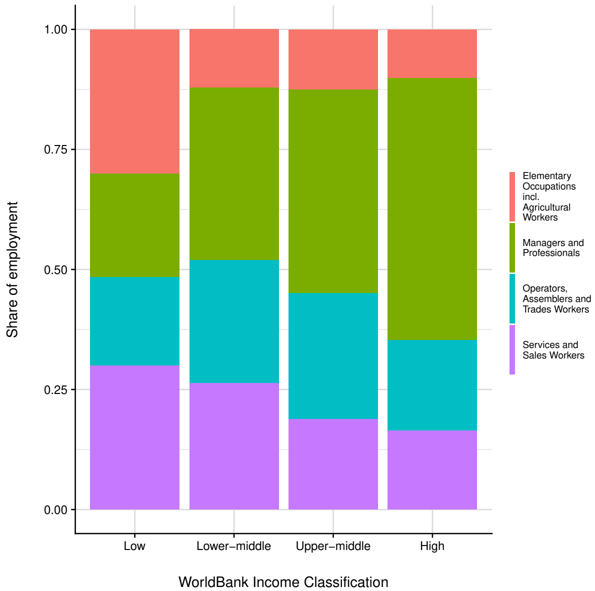

Note: This figure reports the share of occupations in employment of all countries that belong to a certain country income category as defined by the World Bank. The occupation categories are defined as follows, whereby the number refers to the rows (ISCO codes) of Table 1: Managers and Professionals = 1-4, Operators, Assemblers and Trade Workers = 7-8, Elementary Occupations (incl. Ag Workers) = 6+9, Services and Sales Workers = 5. Data sources: The occupation data are computed from the data sets listed in Table 4, GDP per capita is taken from Penn World Tables (Feenstra et al., 2015; Zeileis, 2019).

observations. Country coverage ranges from countries that are among the poorest, like Ethiopia and Uganda, via middle-income countries to high-income countries including the United States and many European countries. Table 4 in the Appendix contains the full list. The advantage of this dataset is that it allows cross-country comparisons over the entire income spectrum, and allows us to measure occupational composition for many subgroups, in particular by geographic area (urban or rural) and employment status (employee or self-employed).[6]

In this section, we measure the occupation distribution in urban areas. We begin here, since these are more comparable across country income groups. Measures for urban areas are also less sensitive to the treatment of farmers, which we explore in Section 4.

Figure 1 shows employment shares in four broad occupation groups

[6] In the Appendix, we present alternative calculations based on ILO data, which have a somewhat more comprehensive coverage and include more recent observations for some countries. Results are similar to the ones in the main text (see D for details). They are also similar to the results computed by Dingel and Neiman (2020*b*) using ILO data. The disadvantage of the ILO data is that they do not permit a disaggregation of the occupational composition along several dimensions at once, and therefore do not allow analyzing urban wage and self-employment separately.

Covid Economics 8, 22 April 2020: 85-105

Table 2: Percent of workers who can work from home by country income level

	Low	Lower-middle	Upper-middle	High
Urban	22.1	29.6	31.2	37.1
Urban, wage employed	28.0	32.9	31.7	36.7
Urban, self-employed	15.5	23.8	28.8	40.4
Urban, WFH for self-empl. from IFLS	19.5	24.6	27.6	33.1
Urban and rural	14.7	24.8	28.8	34.7
Urban and rural, WFH for farmers =1	64.3	42.9	34.2	37.5

Note: The numbers represent averages across country-years' WFH employment shares within each income group as defined by the World Bank classification in 2018.

for four country income groups. It is evident that in high income countries, a very large share of employment is in managerial and professional occupations.[7] This share decreases monotonically as one goes from the highest to the lowest country income group, from 55 to 22%. In contrast, employment in low income countries is concentrated in elementary occupations and agricultural activities (30%). The share of employment in such activities is minor in rich countries (10%). The share of employment in services and sales occupations is also much larger in low income countries (30%) than in high income ones (17%). The share of employment as operators, assemblers and trades workers is hump-shaped in country income per capita.

The large differences in the occupation composition of employment with income per capita, combined with large differences in the ability to work from home across occupations, imply that the ability to work from home varies strongly with income per capita. Figure 2 shows that the share of workers with occupations that can be done from home is increasing with income levels. The first line of Table 2 proposes a summary, grouping countries by income levels defined by the World Bank. While in the least developed countries the share of occupations that can be executed from home accounts for just over 20% of workers, this share rises to close to 40% in the most developed countries.

This analysis applied the WFH measures by (Dingel and Neiman, 2020*a*) to all countries, so that cross-country differences only reflect differences in the composition of employment across occupations. The next section addresses another potentially important difference between rich and poor countries that affects the ability to work from home, namely the prevalence of self-employment in poor countries. The section after that investigates the importance of the agricultural sector.

[7] We include technicians and clerical support workers in this broad group. See Table 6 in the Appendix for exact figures for all groups.

Covid Economics 8, 22 April 2020: 85-105

Covid Economics 8, 22 April 2020: 85-105

Figure 2: Percent of urban workers who can work from home by income per capita

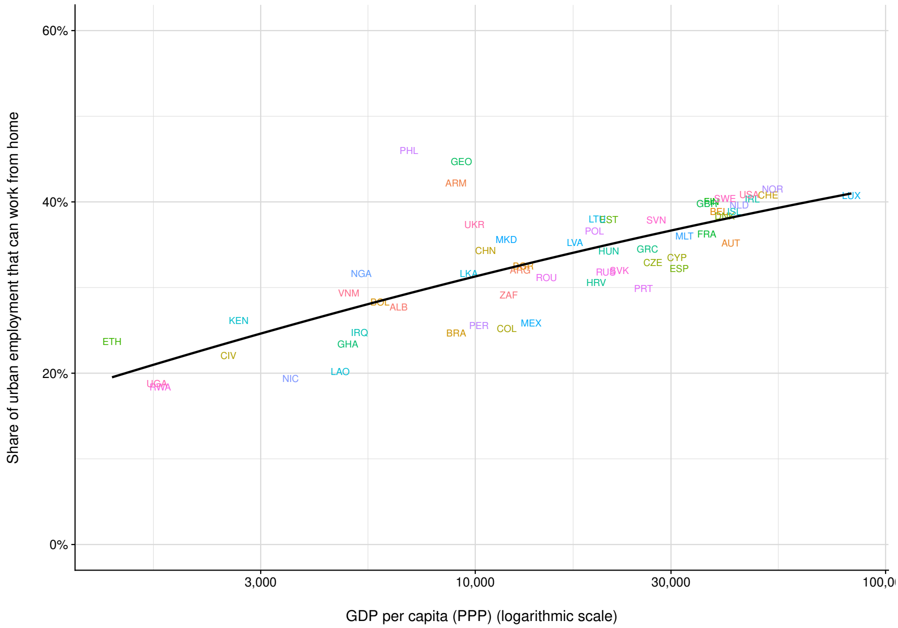

Note: Figure 2 shows the share of the urban employed population with an occupation that can be executed remotely by country year. The data sources for the occupation employment shares are displayed in Table 4. The GDP data is taken from Feenstra et al. (2015); Zeileis (2019), and the share of WFH jobs by occupation is from Table 1.

3 Ability to work from home by employment status

The organization of work differs significantly with country income per capita. In particular, the importance of self-employment varies very strongly with income per capita (Gollin, 2008). While in low-income countries, more than half of the working population is self-employed, only 10% of the working population is self-employed in rich countries. To assess the importance of this pattern, we analyze the WFH ability of the self-employed and wage employees separately. We also compute an alternative measure of WFH ability for household enterprises.

3.1 Ability to work from home for self-employed workers

3.1.1 Baseline results

The third line of Table 2 summarizes the WFH employment shares of the urban self-employed by country income group. Notice that the gap between low and high-income countries is substantially larger than in line 1. In other words, the self-employed in low income countries are

Covid Economics 8, 22 April 2020: 85-105

Figure 3: Distribution of occupations by country income level, urban areas, by type of employment

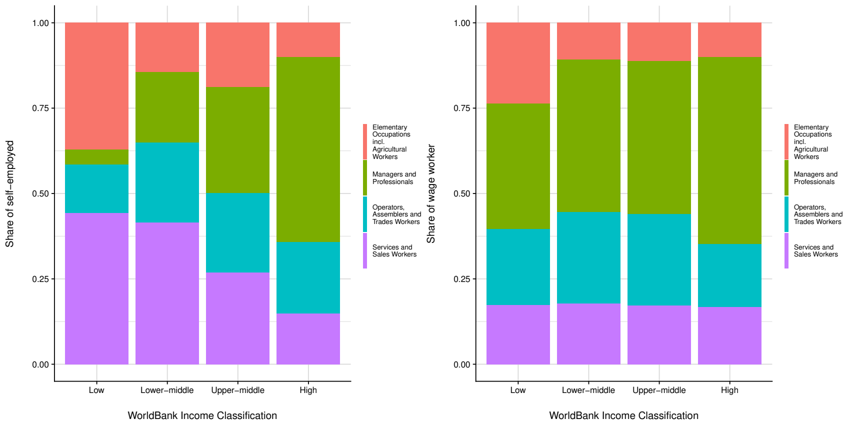

(a) Occupations of urban self-employed (b) Occupations of urban wage workers

Note: See Figure 1.

particularly limited in their options to carry out work from home.

This is due to the cross-country variation in the occupation composition of the self-employed. Figure 3 documents the occupation distribution of urban employment for wage employees (panel (a)) and the self-employed (panel (b)) separately, again by country income level (see Table 6 for the corresponding numbers). What stands out is that in rich countries, the occupational composition of the self-employed is similar to that of employees, and therefore to the aggregate occupation composition. In poor countries, in contrast, household enterprises are concentrated in occupations characterized by low WFH scores (notably elementary occupations and services and sales occupations), with only a negligible share of employment in the high-WFH score managerial and technical professions.

3.1.2 Alternative WFH measure

How easy is it to operate a household business from home? It is conceivable that the WFH measures computed by DN do not fully capture the ability to run a small household business from home in a poor country, given that they are based on a survey of work arrangements from a country where employment is concentrated in relatively large firms. For example, it may be possible to operate small production businesses, e.g. for food or garments, from the household. To assess the ability of running a household enterprise from home, we therefore compute an alternative

CEPR PRESS

Covid Economics 8, 22 April 2020: 85-105

WFH measure, directly using data on household businesses.

A WFH measure for household enterprises. For this, we use the 2014 Indonesia Family Life Survey (IFLS Wave 5).[8] The survey is useful for our purposes as it records information on the location of business activity and on job characteristics. It collects detailed information on household businesses, including sector, ownership, and many others. We use information on urban non-farm businesses without paid employees. The survey also records, for each business, the identity of the household member most involved in the business. We use this to match their occupation to the business. We restrict our analysis to those who are self-employed as their main activity, to ensure that the recorded occupation actually refers to the household business.

We build a WFH measure based on two criteria, paralleling Dingel and Neiman (2020*a*). First, the survey records whether a business is operated entirely or partially outside the household's home, or not. Our first "loose" measure for the ability to run a household business from home is one for businesses not operating outside the home, and zero otherwise.[9] Second, the survey records information on job characteristics at the individual level. The one that most closely matches our objective is "My job requires skill in dealing with people." Our strictest measure for the ability to operate a household businesses from home takes the value one if the loose measure is one and the reply to this question is "None/Almost none of the time." We also define an intermediate measure, which is one if the loose measure is one and the reply to this question is "None/Almost none of the time" or "Some of the time". These two stricter criteria capture the fact that even when the location of a business is in the household's home, it may still require interaction with people from outside the household. This can be close, as in the case of a hairdresser, or more distant, as in the case of a business selling prepared food (a very common type of business).[10]

Table 3 shows the proportion of household businesses that can be operated from home, for the three measures, by ISCO 1 occupation. While a significant fraction of businesses are operated from home (loose criterion) in several broad occupation groups, our measures for the ability

[8]Indonesia is a lower-middle income economy. The IFLS has been used very widely in research.

[9]Since this question asks whether a business is currently operated at home, and not whether it could in principle be operated from home, this aspect of our criterion is stricter than DN.

[10]Ideally, the question would ask about the frequency or importance of customer interaction, not the required skill. Yet, we presume that if no skill in dealing with people is required, this probably indicates no or very few interactions with people.

Table 3: Percent of household businesses that can operate from home by ISCO1 occupation

	WFH criterion		
Occupation, ISCO 1 digit	loose	inter.	strict
Managers	0.0	0.0	0.0
Professionals	30.0	0.0	0.0
Technicians and Associate Professionals	27.6	6.3	0.0
Clerical Support Workers	6.9	0.0	0.0
Services and Sales Workers	21.7	10.8	4.2
Skilled Agricultural, Forestry and Fishery Workers	18.2	7.4	6.6
Craft and Related Trades Workers	23.6	22.4	12.0
Plant and Machine Operators and Assemblers	6.3	0.0	0.0
Elementary Occupations	12.8	3.5	0.0

Note: Data sources: Indonesia Family Life Survey (IFLS) 2014. Loose criterion: The business is not operated outside the home (question NT05b). Intermediate/strict criterion: Loose, and the job of the main person responsible for the business requires skill in dealing with people some or none of the time/never.

Covid Economics 8, 22 April 2020: 85-105

to WFH decline to very low levels once customer interaction is taken into account.

In the following, we will use the intermediate measure. Depending on occupation, the share of businesses that can be operated from home ranges from zero to 22% according to this measure. Compared to the figures for employees shown in Table 1, a significantly larger share of craft and related trades can be operated from home, if they are conducted by the self-employed. In contrast, household enterprises in managerial or professional occupations, technicians, and clerical support work can barely be conducted from home (note though that there are very few household enterprises of these types). The ability of service work to be conducted from home is also lower for household enterprises. Overall, this measure thus reports a lower ability to WFH.

Ability to work from home. We next compute the share of urban employment that can WFH using the measure of WFH ability for household enterprises shown in Table 3 (intermediate criterion). We continue to use the measure by DN for wage employees. Results are shown in Figure 4 and summarized in line 4 of Table 2.

In line with the lower ability to WFH of the measure for household enterprises, this Figure shows a generally somewhat smaller share of employment that can be done from home. It drops from around 37% to 33% for the richest countries. The drop is similar for the poorest countries,

Covid Economics 8, 22 April 2020: 85-105

from 22% to 19.5%.

To summarize, the high levels of self-employment in poor countries, combined with its concentration in occupations where it is difficult to work from home, contributes significantly to the lower ability to WFH in poorer countries.

Figure 4: Percent of urban workers who can work from home by income per capita, with employment-type specific WFH score

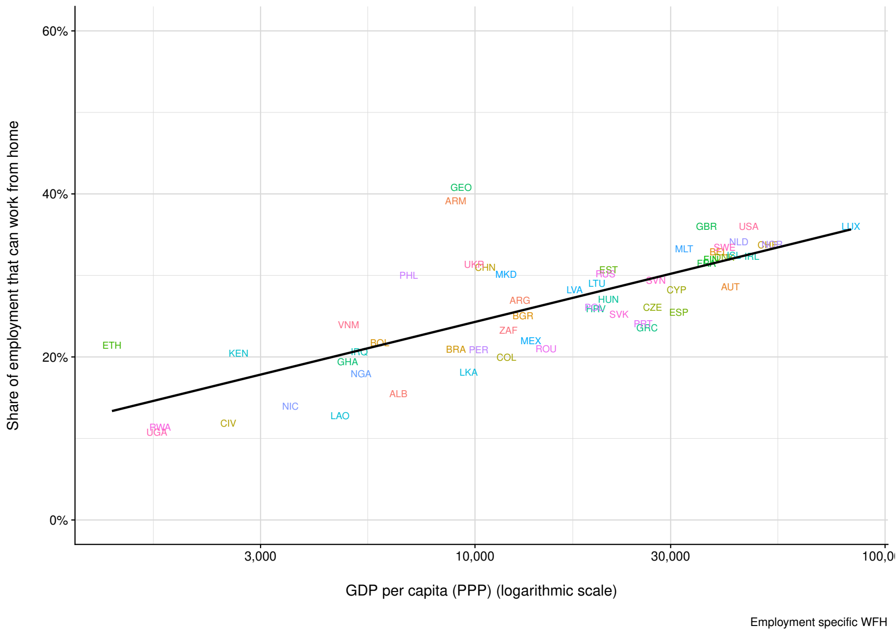

Note: Figure 4 shows the share of WFH employed population when WFH wage employment and self employment specific by income per capita. Data sources as in Figure 2. the share of WFH jobs for wage workers is based on table 1, and the share of WFH jobs for self-employed workers is taken from WFH in table 3 (intermediate WFH criterion).

3.2 Wage employees

Figure 5 shows the ability to work from home across countries for wage employees only, again using data for urban areas. Panel (a) shows the share of wage employees in each country that can work from home. For rich countries, the differences between this figure and Figure 2 are small, reflecting the dominance of wage employment in aggregate employment in these countries. Yet for poor countries, differences are notable: the share of wage employees who can work from home in the poorest countries reaches almost 30%, significantly exceeding the aggregate share of urban employment that can be done from home. This can also be seen by comparing the first two lines of Table 2. The reason for this is that the occupation distribution of wage employment differs much less across

Covid Economics 8, 22 April 2020: 85-105

countries than that of all employment. In particular, employees in poor countries are not as concentrated in elementary and services and sales occupations as the self-employed are. (See Figure 3 and Table 6.)

Panel (b) of Figure 5 depicts the share of the wage bill accounted for by urban employees able to work from home. It varies less systematically by income per capita. Compared to panel (a), there is an additional composition effect at work: occupations with high WFH scores, which already are high-wage occupations in the US (DN), tend to pay even higher wages in poor countries. As such occupations are skill-intensive (managers, professional), this is likely a reflection of the scarcity of skill supply in these occupations in developing countries. To the extent that wages are informative of efficiency units of labor, we conclude that the fraction of efficiency units of wage employment that can be provided from home is weakly correlated with the level of development.

Figure 5: Ability to work from home for wage employees

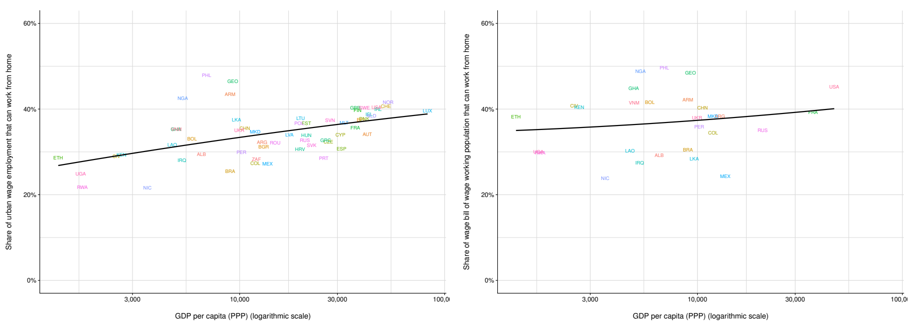

(a) Percent of urban wage employees who can work from home by income per capita

(b) Wage bill share of employees who can WFH by income per capita.

Note: Figure 5a shows the share of the urban wage-working population with an occupation that can be executed remotely by country year. Figure 5b displays the share of the wage bill that is spent on wage jobs. The country year coverage is smaller since wage information is only available for a subset of the surveys. Data sources as in Figure 2.

4 The role of farmers' ability to work from home

A second specificity of poor countries is the much larger share of agricultural employment. This did not affect results in the main analysis, since that focussed on urban employment. However, results for rural areas or at the national level will crucially depend on the ability of farmers to work from home.

The Dingel and Neiman (2020*a*) classification finds that farmers can barely work from home. It is not clear to what extent this is applicable in farming in poor countries, which occurs in very different technological

Figure 6: Percent of a country's workers who can work from home by income per capita

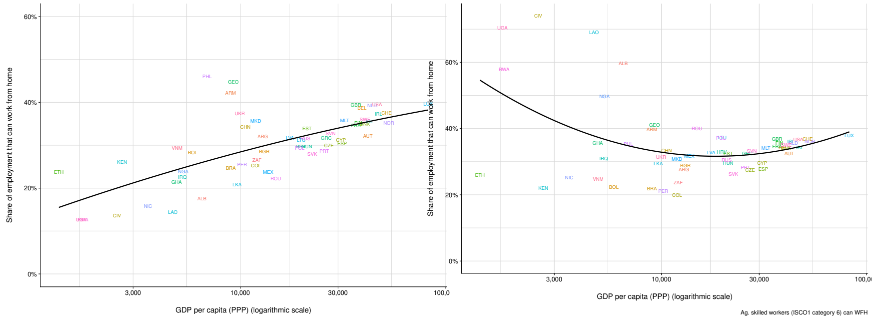

(a) Baseline WFH scores

(b) WFH score of 1 for agricul. workers

Note: Data sources as in Figure 2. Panel (a) is analogous to that figure, using data for the entire country. Panel (b) is similar, except for the assumption that the ability to WFH is 1 for the occupation "Skilled Agricultural, Forestry and Fishery Workers".

Covid Economics 8, 22 April 2020: 85-105

and geographical settings. In rural areas, a very large fraction of households engage in some farming. If plots are close to home, or adjacent to home, farming may be possible from home, at least for some time. Similarly, in such a setting, a large fraction of output is consumed within the household, and not sold to market. This could also be sustained while working from home.

We therefore next show the ability to WFH for the country as a whole, under two alternative assumptions on the share of farmers able to do WFH: 0.083, as in Table 1, or one, almost the polar opposite.[11] Our findings will give an indication of how much restrictions on farmers' ability to work affect overall labor supply.

Results are shown in Figure 6. These figures show that for the aggregate ability to WFH in poor countries, farmers' ability to WFH is crucial. If farmers cannot work from home, the share of workers who can work from home in the poorest countries is extremely low, at less than 20%. If, in contrast, farmers are assumed to be able to WFH, this rises to 30 to 70%, somewhat higher than the average for rich countries.[12] The bottom half of Table 2 summarizes these results clearly. In the baseline scenario, the WFH gap between low (14.7%) and high-income countries (34.7%) is particularly pronounced. In the second scenario, on the other hand,

[11] A score of 1 probably exceeds the true ability to work from home even for subsistence farmers. Yet, it illustrates the importance of this number very powerfully. Note that while occupation 9 also contains agricultural workers, they are mostly wage workers, and therefore could typically not work from home.

[12] Note that the low productivity of agriculture in poor countries implies that, in such a scenario, the negative effect of only being able to WFH on aggregate output may still be larger in poor than in rich countries.

CEPR PRESS

the gap reverses, with WFH averaging 64.3% in low and 37.5% in high-income countries. The rigidity of social distancing rules for farmers will thus affect the ability to WFH for a significant share of the population.

5 Concluding remarks

The ability to work from home is an important instrument to soften the economic fallout resulting from social-distancing measures to stem the COVID-19 pandemic. We document that the occupational composition in urban areas provides less scope for WFH in developing than in developed countries. This result is particularly driven by self-employed workers: they represent the bulk of employment in developing countries, working in occupations that can hardly be accomplished away from the production site or the customer base. The country-level ability to work from home depends crucially on the WFH ability of farmers.

Covid Economics 8, 22 April 2020: 85-105

Appendix

A Data sources

Figure 2 uses our individual level dataset that consolidates labor force surveys and the labor force section of household surveys from many countries. This dataset harmonizes information on individual characteristics and labor supply. It contains information on employment status, job type, occupation and sector of activity. Table 4 lists all data sources used to construct the dataset.

Covid Economics 8, 22 April 2020: 85-105

Covid Economics 8, 22 April 2020: 85-105

Table 4: Individual level dataset. Information on data sources, sample size and country years covered.

Name	Years	Sample size (in thds)	GDP per capita (PPP)	Source
Albania	2002–2012	23	4'845–9'918	LSMS
Argentina	2004–2006	127	12'074–13'770	LFS
Armenia	2013–2013	1	8'979–8'979	STEP
Austria	1999–2017	1'034	34'938–51'524	LFS
Belgium	1999–2017	474	32'357–46'522	LFS
Bolivia	2012–2012	2	5'860–5'860	STEP
Brazil	2002–2006	723	8'358–9'515	LFS
Bulgaria	1995–2017	177	6'390–20'027	LSMS, LFS
China	2012–2012	1	10'596–10'596	STEP
Colombia	2012–2012	2	11'934–11'934	STEP
Cote d'Ivoire	1985–1988	13	2'429–2'734	LSMS
Croatia	2002–2017	155	13'750–24'368	LFS
Cyprus	1999–2017	207	25'255–36'137	LFS
Czech Republic	2002–2017	663	21'374–36'061	LFS
Denmark	1999–2017	511	33'525–49'607	LFS
Estonia	1999–2017	118	10'772–31'013	LFS
Ethiopia	2013–2014	46	1'248–1'357	LFS, UES
Finland	1999–2017	207	31'433–42'902	LFS
France	2003–2017	812	31'567–40'975	LFS
Georgia	2013–2013	1	9'254–9'254	STEP
Ghana	2013–2015	6	4'875–4'910	STEP, LFS
Greece	1999–2017	1'143	22'683–31'340	LFS
Hungary	2001–2017	1'179	16'448–27'531	LFS
Iceland	1999–2017	54	37'732–51'316	LFS
Iraq	2006–2006	27	5'223–5'223	LSMS
Ireland	1999–2017	1'071	33'680–73'297	LFS
Kenya	2013–2013	2	2'652–2'652	STEP
Laos	2012–2012	2	4'693–4'693	STEP
Latvia	2001–2017	154	10'921–26'643	LFS
Lithuania	1999–2017	277	10'373–30'936	LFS
Luxembourg	1999–2017	168	64'436–99'477	LFS
Macedonia	2013–2013	2	11'910–11'910	STEP
Malta	2009–2017	76	26'792–41'847	LFS
Mexico	2005–2005	163	13'691–13'691	LFS
Netherlands	1999–2017	834	37'786–50'024	LFS
Nicaragua	2005–2005	12	3'548–3'548	LSMS
Nigeria	2010–2018	18	4'971–5'641	LSMS
Norway	2005–2017	111	49'908–63'768	LFS
Peru	2009–2014	115	8'515–11'086	LFS
Philippines	2015–2015	1	6'896–6'896	STEP
Poland	2006–2017	1'155	16'416–28'420	LFS
Portugal	1999–2017	771	22'413–28'567	LFS
Romania	2009–2017	694	16'752–25'262	LFS
Russian Federation	2004–2015	77	12'554–25'777	RLMS-HSE
Rwanda	2013–2016	49	1'551–1'872	LFS
Slovakia	2007–2017	354	22'724–30'433	LFS
Slovenia	2005–2017	297	26'506–33'947	LFS
South Africa	2012–2019	243	11'965–12'201	QLFS
Spain	1999–2017	920	25'102–37'233	LFS
Sri Lanka	2012–2012	1	9'653–9'653	STEP
Sweden	1999–2017	1'441	34'468–47'892	LFS
Switzerland	2010–2017	232	54'028–62'927	LFS
Uganda	2009–2013	21	1'571–1'759	LSMS
Ukraine	2012–2012	1	9'956–9'956	STEP
United Kingdom	1999–2017	702	31'110–42'138	LFS
United States	1998–2004	220	43'625–49'138	CEPR
Viet Nam	2012–2012	2	4'917–4'917	STEP
		17'892	1'248–99'477	

B Working from home by more detailed ISCO occupations

Table 5: Working from home by occupation category ISCO-2.

Occupation, ISCO 2 digit	Share of WFH occupations
Chief Executives, Senior Officials and Legislators	87.7
Administrative and Commercial Managers	89.9
Production and Specialized Services Managers	69.1
Hospitality, Retail and Other Services Managers	46.3
Science and Engineering Professionals	66.0
Health Professionals	11.0
Teaching Professionals	96.6
Business and Administration Professionals	95.1
Information and Communications Technology Professionals	100.0
Legal, Social and Cultural Professionals	68.5
Science and Engineering Associate Professionals	19.7
Health Associate Professionals	6.0
Business and Administration Associate Professionals	70.8
Legal, Social, Cultural and Related Associate Professionals	58.0
Information and Communications Technicians	81.8
General and Keyboard Clerks	100.0
Customer Services Clerks	28.3
Numerical and Material Recording Clerks	51.9
Other Clerical Support Workers	63.3
Personal Services Workers	23.8
Sales Workers	21.1
Personal Care Workers	21.9
Protective Services Workers	11.8
Market-oriented Skilled Agricultural Workers	10.0
Market-oriented Skilled Forestry, Fishery and Hunting Workers	9.6
Subsistence Farmers, Fishers, Hunters and Gatherers	0.0
Building and Related Trades Workers (excluding electricians)	1.5
Metal, Machinery and Related Trades Workers	0.0
Handicraft and Printing Workers	15.9
Electrical and Electronics Trades Workers	0.0
Food Processing, Woodworking, Garment and Other	7.9
Stationary Plant and Machine Operators	0.0
Assemblers	0.0
Drivers and Mobile Plant Operators	23.7
Cleaners and Helpers	0.0
Agricultural, Forestry and Fishery Labourers	0.0
Labourers in Mining, Construction, Manufacturing and Transport	8.3
Food Preparation Assistants	0.0
Street and Related Sales and Services Workers	0.0
Refuse Workers and Other Elementary Workers	25.0

Note: We follow the classification provided by Dingel and Neiman (2020*a*) who use two ONET surveys with information on work context and generalized work activities for many jobs. They consider a job to not be teleworkeable requires amongst others the handling of equipment or contact with the public or if the job has a work context that requires the handling of objects and tools (that are not computers). We use a cross-walk to map DN's measures to the ISCO-2 classification.

Covid Economics 8, 22 April 2020: 85-105

CEPR PRESS

C Additional tables

Table 6: Average occupation share by country income level

Urban, all	Low	Lower-middle	Upper-middle	High
Managers and Professionals	0.216	0.359	0.424	0.546
Services and Sales Workers	0.301	0.264	0.189	0.165
Elementary Occupations incl. Agr. Workers	0.299	0.121	0.125	0.101
Operators, Assemblers & Trades Workers	0.184	0.256	0.261	0.189
Urban self-employed				
Managers and Professionals	0.044	0.206	0.310	0.541
Services and Sales Workers	0.443	0.416	0.269	0.148
Elementary Occupations incl. Agr. Workers	0.371	0.144	0.188	0.100
Operators, Assemblers & Trades Workers	0.143	0.235	0.234	0.210
Urban employee				
Managers and Professionals	0.368	0.446	0.448	0.547
Services and Sales Workers	0.174	0.178	0.173	0.167
Elementary Occupations incl. Agr. Workers	0.236	0.108	0.111	0.101
Operators, Assemblers & Trades Workers	0.222	0.268	0.267	0.186

Note: The numbers represent averages across countries' occupation shares, conditional on employment status, within each income group as defined by the World Bank classification in 2018. The first occupation group (ISCO 1-4) consists of *Managers*, *Professionals*, *Technicians and associate professionals*, and *Clerical support workers*; the second (ISCO 5) of *Services and sales workers*; the third (ISCO 6 & 9) of *Skilled agricultural, forestry and fishery workers* and *Elementary occupations*; and the fourth (ISCO 7 & 8) of *Craft and related trades workers* and *Plant and machine operators*.

D Robustness: ILO data on occupation employment shares

The main results are based on the occupational composition in our assembled dataset. Here, we re-compute some of the results using occupational employment provided by ILO data. Figure 7 focuses on urban employment and is analogous to Figure 2. We confirm a positive correlation between the share of WFH employment and GDP per capita. Figures 8a depicts the share of WFH for the aggregate economy (both urban and rural) using the baseline WFH scores, while Figure 8b does the same while attaching a WFH score of 1 to agricultural workers. They confirm the trends portrayed in Figure 6. Table 7 summarizes the findings by country income groups. Most importantly, the main specification in the first line suggests that 22.1% of workers can execute their work from home in low-income countries, as opposed to 37.4% in high-income countries.

Covid Economics 8, 22 April 2020: 85-105

Covid Economics 8, 22 April 2020: 85-105

Figure 7: Percent of urban worker that can work from home, ILO data

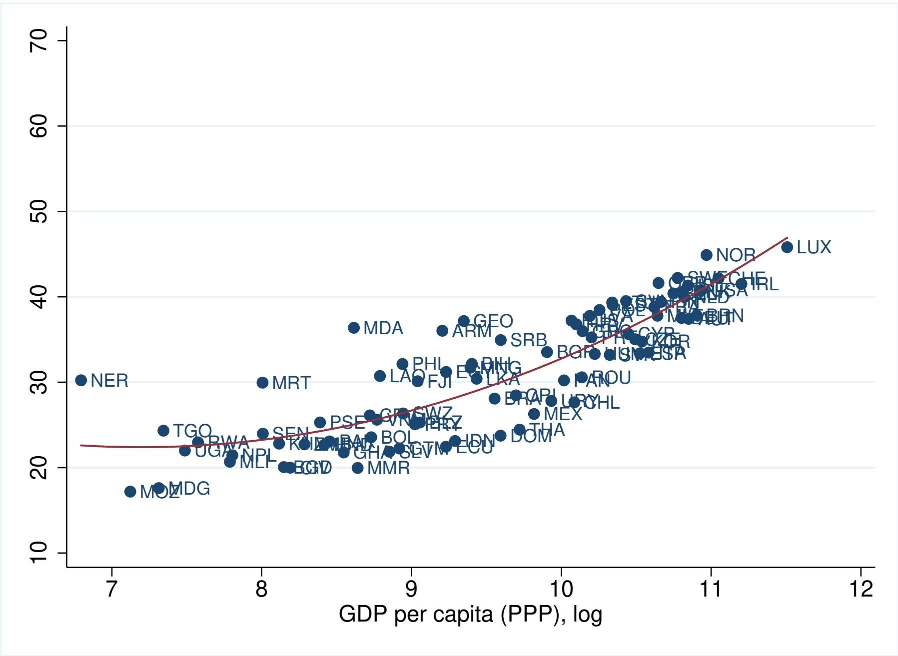

Note: The employment share of WFH combines WFH scores from Table 1 and ISCO-1 employment by occupation in urban areas from ILO. GDP data is from Feenstra et al. (2015); Zeileis (2019). Each country is the most recent annual observation over the period 2015-2019 for which occupational and GDP data are available. The regression line is a quadratic fit.

Figure 8: Percent of all workers (both urban and rural) that can work from home, ILO data

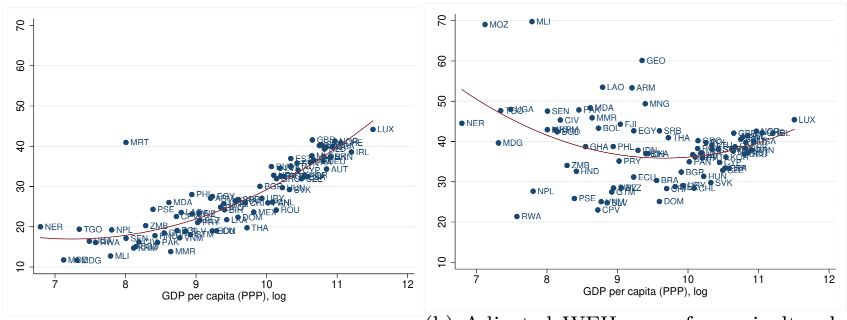

(a) Baseline WFH scores

(b) Adjusted WFH score for agricultural workers

Note: In panel (a), the employment share of WFH combines WFH scores from Table 1 and ISCO-1 employment by occupation in both urban and rural areas from ILO. In panel (b), the WFH score of skilled agricultural workers (ISCO code: 6) is set to 1. GDP data is from Feenstra et al. (2015); Zeileis (2019). Each country is the most recent annual observation over the period 2015-2019 for which occupational and GDP data are available. The regression line is a quadratic fit.

Covid Economics 8, 22 April 2020: 85-105

Table 7: Percent of workers who can work from home by country income level, ILO occupation data

	Low	Lower-middle	Upper-middle	High
Urban	22.1	24.5	29.2	37.4
Urban, wage employed	–	–	–	–
Urban, self-employed	–	–	–	–
Urban, WFH for self-empl. from ILFS	–	–	–	–
Urban & rural	15.4	20.6	24.0	35.0
Urban & rural, WFH for farmers =1	48.6	38.8	36.7	37.5

Note: The numbers represent averages across countries' WFH employment shares within each income group as defined by the World Bank classification in 2018.

References

Barrot, J.-N., Grassi, B. and Sauvagnat, J. (2020), 'Sectoral effects of social distancing', *Covid Economics* **3**.

Boeri, T., Caiumi, A. and Paccagnella, M. (2020), 'Mitigating the work-security trade-off while rebooting the economy', *Covid Economics* **2**.

Chiou, L. and Tucker, C. E. (2020), 'Social distancing, internet access and inequality', *NBER Working Paper* **26982**.

Dingel, J. and Neiman, B. (2020*a*), 'How many jobs can be done at home?', *Covid Economics* **1**.

Dingel, J. and Neiman, B. (2020*b*), 'How many jobs can be done at home?', *Becker Friedman Institute White Paper* .

Feenstra, R. C., Inklaar, R. and Timmer, M. P. (2015), 'The next generation of the penn world table', *American Economic Review* **105**(10), 3150–3182.
URL: *http://www.ggdc.net/pwt/*

Gollin, D. (2008), 'Nobody's business but my own: Self-employment and small enterprise in economic development', *Journal of Monetary Economics* **55**(2), 219–233.
URL: *http://linkinghub.elsevier.com/retrieve/pii/S0304393207001493*

Hensvik, L., Le Barbanchon, T. and Rathelot, R. (2020), 'Which jobs are done from home? evidence from the american time use survey', *IZA Discussion Paper* **13138**.

Loayza, N. and Pennings, S. (2020), 'Macroeconomic policy in the time of covid-19: A primer for developing countries', *World Bank Research & Policy Briefs* **28**.

Mobarak, A. M. and Barnett-Howell, Z. (2020), 'Poor countries need to think twice about social distancing', *Foreign Policy* .

Saltiel, F. (2020), Who can work from home in developing countries?

Zeileis, A. (2019), *pwt9: Penn World Table (Version 9.x)*. R package version 9.1-0.
URL: *https://CRAN.R-project.org/package=pwt9*

CEPR PRESS

Covid Economics Issue 8, 22 April 2020

Welfare resilience in the immediate aftermath of the Covid-19 outbreak in Italy[1]

Francesco Figari[2] and Carlo V. Fiorio[3]

Covid Economics 8, 22 April 2020: 106-133

Date submitted: 15 April 2020; **Date accepted:** 16 April 2020

This paper analyses the extent to which the Italian welfare system provides monetary compensation for those who lost their earnings due to the lockdown imposed by the government in order to contain the Covid-19 pandemic in March 2020. In assessing first-order effects of the businesses temporarily shut down and the government's policy measures on household income, counterfactual scenarios are simulated with EUROMOD, the EU-wide microsimulation model, integrated with information on the workers who the lockdown is more likely to affect. This paper provides timely evidence on the differing degrees of relative and absolute resilience of the household incomes of the individuals affected by the lockdown. These arise from the variations in the protection offered by the tax-benefit system, coupled with personal and household circumstances of the individuals at risk of income loss.

1 We thank Manos Matsaganis, Holly Sutherland and Alberto Tumino for comments on a preliminary version. The methodology draws on Fernandez Salgado et al. (2014). We use EUROMOD (version I2.0+) which is developed and managed by the Institute for Social and Economic Research (ISER) at the University of Essex, in collaboration with the European Commission - JRC Seville and national teams from the EU member states. We are indebted to Holly Sutherland and the many people who have contributed to the development of EUROMOD. The process of extending and updating EUROMOD is financially supported by the European Union Programme for Employment and Social Innovation 'Easi' (2014-2020). We make use of microdata from the Italian versions of the SILC data made available by ISTAT. Data providers bear no responsibility for the analysis or interpretation of the data reported here. Any mistakes are the authors' only.

2 Associate Professor of Public Finance, University of Insubria.

3 Jean Monnet Chair in EU Public Finance and Professor of Public Economics, University of Milan.

1. Introduction

The COVID-19 pandemic can lead to a worldwide economic downturn worse than the one that characterised the 2008 Great Recession. The potential impact on GDP, although mostly unpredictable today without a clear knowledge of the boundaries of the health emergency, can lead to a massive slump in economic development (Dorn et al. 2020) depending on the scenarios.

Italy has been the European front runner in terms of infection rates and deaths in the population, as it experienced a sudden outbreak at the end of February 2020. As a consequence, the Italian government issued various decree laws which limited and shut down economic activity, in order to prevent contagion through social contacts and to limit the virus spread. Dorn et al. (2020) estimates that a two-month shutdown would lead to a reduction of annual GDP growth by 8–13 percentage points. Qualitative indicators already show the effect of unprecedented demand and supply shocks due to the COVID-19 pandemic. The business confidence climate index crashed from 97.8 to 81.7. The confidence index in manufacturing reduced sharply from 98.8 to 89.5 (Istat, 2020)

OECD estimates of the initial direct impact of shutdowns reveal that the output decline would be of roughly 20%-25%, with consumer expenditure dropping by 33%. Such a decline in the level of output would correspond to a decline in annual GDP growth of around 2 percentage points for each month of shutdown (OECD 2020a).
Focusing on the situations faced by workers, the International Labour Organization estimates a rise in global unemployment of between 3% and 13%, with underemployment expected to increase on a large scale and the decline in economic activity and travel limits impacting both manufacturing and services (ILO, 2020)

The adverse impact of the necessary containment measures to the COVID-19 pandemic has determined unprecedented demand and supply shocks to international growth prospects. Financial markets reacted with a sharp increase of volatility and fall in asset prices. The outlook for world trade, which was already declining in January, worsened dramatically in March (Istat, 2020). Despite the negative outlook, the cost of a government inaction would have been much higher in terms of human lives and long-term recovery.

The picture described above, as well as the lessons of previous recessions, suggest that the downturn due to the COVID-19 pandemic will overshadow European economies for years to come, through a legacy of unemployment, public debt and long-lasting impacts on household incomes as already experienced during the Great Recession (Jenkins et al., 2013). Furthermore, Saez and Zucman (2020) argue that governments “can prevent a very sharp

Covid Economics 8, 22 April 2020: 106-133

but short recession from becoming a long-lasting depression" by acting as payer of last resort: providing insurance to the affected workers and making sure that cash flows to idle workers and businesses immediately. To this end, governments have introduced discretionary policy measures to support the most vulnerable (OECD, 2020b).

However, a word of caution should be cast in that Dolls et al. (2012) show that automatic stabilizers differ greatly across countries, particularly in the case of asymmetric shocks. The observation is particularly relevant in the case of the Italian tax-benefit system, whose income stabilisation mechanisms may be limited by design in times of emergency.

The primary aim of this paper is to offer a scenario, rather than a forecast, in order to understand in a timely fashion the extent to which the Italian tax-benefit system provides income stabilisation in the first month of the health emergency for those who lost their earnings at the very beginning of the COVID-19 pandemic. In particular, we aim to measure the amount of income insurance that individuals and their households receive from the Welfare State against the hazard of the economic shutdown. The consequences of the shutdown on the most vulnerable individuals depend on their individual characteristics and the interaction between their labour market participation, their living arrangements and the capacity of the tax and benefit systems. We do not consider other aspects such as the reduced likelihood to get a job for those who are looking for one and the wider consequences of macroeconomic feedbacks.

Lack of longitudinal up-to-date information on household income and labour market circumstances, usually available a few years after the economic shock and in a limited number of countries only, constrains the possibilities for empirical analysis. To address this limitation, we assess the impact of the economic lockdown on household income by means of simulating counterfactual scenarios with a fiscal microsimulation approach (Figari, Paulus, Sutherland, 2015). First, we attempt to identify the workers affected by the lockdown by using aggregate data on employment shares by activity sectors. Second, we estimate the household incomes for individuals who lose their earnings, considering the direct cushioning effect of the tax-benefit system in relation to how they depend on the remaining household market income as well as personal and household characteristics. The use of tax-benefit microsimulation models to consider how the welfare systems protect people from an extreme shock is known as a "stress test" of the tax-benefit system (Atkinson, 2009) and has become increasingly popular in analysing consequences of the Great Recession (Figari et al., 2014, Jenkins et al., 2013).

We highlight the main motivations to exploit such an approach in Section 2. Therewithin we introduce EUROMOD and we describe the indicators we

Covid Economics 8, 22 April 2020: 106-133

apply to capture the resilience of the welfare system in both relative and absolute terms. Section 3 provides a snapshot of the characteristics of those affected by earning loss.

The current analysis focuses on Italy but it is about to be extended to other EU countries in order to highlight the interaction between the country-specific effects of the pandemic and the policy responses implemented by national governments, and also to generalise the impact of the COVID-19 pandemic in a cross-country perspective. The most relevant features of the policy measures included in the analysis are described in Section 4.

Empirical evidence on the different income stabilisation aspects of the Italian tax-benefit system is presented in Section 5, which shows differing degrees of how individual loss of earnings can reduce household incomes, as well as to what extent those incomes are resilient upon intervention . Section 6 concludes, summarising the main findings and suggesting future work and improvements in light of ongoing developments as data is made available.

2. Empirical methodology

2.1. Stress testing the tax-benefit systems

In the presence of a sudden economic shock with direct consequences for the labour market participation of individuals, coupled with fiscal policies implemented to react to unexpected earning losses, understanding how contemporary tax-benefit systems react to changes in individual circumstances is essential. More importantly, it is fundamental to assess the extent to which household incomes are protected by the tax-benefit systems.

The stress test approach is common in financial institutions to test the sensitivity of a portfolio to a set of extreme but plausible shocks and to assess the significance of the system's vulnerabilities (Jones et al., 2004). We follow Tony Atkinson's suggestion of extending the same approach to tax-benefit systems in order to predict the cushioning effects of the social protection schemes in the event of a loss of market incomes and to assess overall income stabilisation after a macroeconomic shock (Atkinson 2009, Fernandez Salgado et al., 2014).

A stress test exercise can provide evidence of the effects of either a hypothetical macroeconomic shock or a contemporary shock for which survey data covering the period of interest are not yet available. The latter option is the approach we follow to assess the variation in the social impact of the earning loss due to the economic shutdown at the very beginning of the COVID-19 pandemic in Italy. In due course, survey data collected over the period of the pandemic will provide evidence of the evolution of income

Covid Economics 8, 22 April 2020: 106-133

distribution, while analysis of longitudinal data will show how incomes changed for those directly affected by the lockdown.

Moreover, it is important to assess the economic impact of specific aspects of the pandemic and to inform the policy debate in a timely fashion. By using a fiscal microsimulation model which combines detailed survey data, representative of the national population, on market incomes and household characteristics with tax-benefit rules (Figari, Paulus, Sutherland, 2015), we can determine the different components of household disposable income under different counterfactual scenarios in which we identify the individuals more likely to lose their earnings as a result of an economic shock.

The simulated household disposable income as related to the individuals losing from the lockdown depends on the cushioning effect of automatic stabilizers existing in the country in the form of (a) income taxes and social contributions, (b) contributory benefits for those who lose their earnings (if entitled), (c) other means-tested benefits and tax credits designed to protect families on low income, and (d) other household incomes, in the form of earnings of those still in work as well as pensions and benefits, received by other household members. In addition, it is crucial to capture the effects of the discretionary policies that the government might decide to implement in order to prevent a sudden fall in household income.

The stress test approach allows us to focus on a specific aspect of the economic shock, highlighting the direct compensation provided by tax-benefit systems rather than that arising from other adaptive changes in individual behaviours. In this paper we focus exclusively on the loss of earnings as one of the channels through which the COVID-19 pandemic directly affects individual well-being. The overall effect of the pandemic on income distribution is likely to be affected by general equilibrium consequences and other behavioural responses. However, individuals and households directly affected by earning loss suffer to a large extent and it is important to assess the extent to which the welfare system helps to stabilise their income, and whether there are specific weaknesses in the policy instruments in operation.

2.2. Counterfactual scenario derived using EUROMOD

We exploit the potential of the microsimulation techniques to define the counterfactual scenario (Figari et al., 2015), based on survey data representative of the national population before the onset of the pandemic, in which we impute the earning loss as observed in March 2020 and we simulate the discretionary policy measured implemented in the same month.

We make use of EUROMOD, the EU-wide tax-benefit microsimulation model. EUROMOD simulates tax liabilities (direct tax and social insurance

Covid Economics 8, 22 April 2020: 106-133

contributions) and benefit entitlements for the household populations of EU Member States in a comparable way across countries on the basis of the tax-benefit rules in place and information available in the underlying datasets. The components of the tax-benefit systems which are not simulated (e.g. old age pensions) are extracted from the data, along with information on original incomes. The simulation of the Wage Supplementation Scheme (*Cassa Integrazione Guadagni*) is based on reported earnings, where relevant, and under assumptions about past contributions derived from the limited information available in the data. See Sutherland and Figari (2013) for further information.

The underlying micro data come from the 2017 national version of the EU-SILC provided by Istat. The analysis in this paper is based on the tax-benefit rules in place in 2019 (as of June 30th), which are essentially identical to those in place in March 2020. Monetary values of non-simulated income components referring to 2016 were updated to 2019 according to actual changes in prices and incomes over the relevant period, as documented in the Italian EUROMOD Country Report (Ceriani et al. 2019). No adjustment is made for changes in population composition between 2016 and 2019.

In the analysis we focus on what happens in a single month, i.e. March 2020. We compute household disposable income, taking account of the discretionary measures included in the Decree Law 18/2020 ("*Cura Italia*") and detailed in the next section.

Given the extraordinary and sudden decision of the government to impose a generalised economic lockdown, the traditional automatic stabilizers embedded in the tax-benefit systems are not allowed to operate, with the exception of income tax and social contributions which are lower due to the lower level of earnings. The existing income-tested benefits (I.e. bonus IRPEF, Family allowances (ANF), Citizenship income (RdC)) based on the income and means-test of the previous fiscal year do not react to the loss of earnings experienced in March 2020. The opportunity to modify the design of the existing income support mechanism to deal with the economic effects of the pandemic is part of the policy debate in Italy (Forum Diseguaglianze Diversità and ASviS, 2020) and we refer to this in the conclusion.

We aim to highlight the amount of insurance coverage guaranteed directly by government, independently of any potential change in the behaviour of family members which could occur in the short or long term. Furthermore, considering the incidence of the shadow economy in Italy, gross self-employed income has been calibrated so as to obtain an aggregate amount corresponding to that reported in fiscal data (Fiorio and D'Amuri, 2006) and we assume there are no changes in the tax evasion behaviour as a consequence of the shock.

2.3. Income stabilisation indicators

Our analysis focuses on both relative and absolute resilience provided by the welfare state, taking into account the interactions of the tax-benefit policies with other existing household income and household composition.

First, in order to assess the level of stabilisation of incomes with respect to the pre-shock baseline, we employ the Net Replacement Rate approach (Immervoll and O'Donoghue, 2004). This gives an indication of the extent of the remaining disposable income for those affected by the economic lockdown and is computed as follows:

$$\text{Net Replacement Rate} = \frac{Y_{post}}{Y_{pre}}$$

where Y is Household Disposable Income made up of Original Income plus Benefits, minus Taxes; Y_{post} and Y_{pre} refer to the income after and before the earning shock, respectively.

In addition to any form of market income, Original Income includes also other sources of personal income, such as private inter-household transfers and alimonies. Even in the lockdown scenario where we simulate the earning shock, household original income may be positive due to income from savings, private pensions, inter-household transfers or the earnings of other household members. Income from savings could be seen as another channel of self-insurance but, given the poor quality of the underlying data, we treat it as one of the components of Original Income, without highlighting its specific role.

To analyse the transmission channels of relative resilience, we decompose the Net Replacement Rate by income source:

$$\text{Net Replacement Rate} = \frac{O_{post} + B_{post} - T_{post}}{Y_{pre}}$$

where O is the Original Income, B is the sum of Benefits and T includes Income Taxes and Social Insurance Contributions paid by employees and the self-employed.

Benefits comprise *(1)* Wage-integration Benefits (*Cassa Integrazione Guadagni*), *(2)* COVID Benefit,i.e. newly discretionary policies such as lump sum transfers to self-employed and employees, (3) Housing Benefits,i.e. amount equivalent to the mortgage instalment for the main residence, (4) Other Benefits,i.e. pension and invalidity benefits, minimum income schemes, family benefits.

Covid Economics 8, 22 April 2020: 106-133

Moreover, to measure the extent of protection offered by public support, we use an indicator developed in Figari et al. 2014, Compensation Rate, which measures the proportion of net earnings lost due to the economic lockdown, compensated by public transfers net of taxes:

$$\text{Compensation Rate}=\frac{\left(B_{post}\text{-}B_{pre}\right)\text{-}\left(T_{(B_{post})}\text{-}T_{(B_{pre})}\right)}{\left(E_{pre}\text{-}E_{post}\right)}$$

where the difference in net earnings before and after the shock represents the income lost due to the lockdown, which in turn is compensated by more generous net benefits. To derive net measures, taxes are allocated proportionally to each income source.

This new indicator allows us to isolate the net public support from the effect of other earnings present in the household of a worker affected by the lockdown, which usually play an important role in determining the income after an individual employment shock. The compensation rate gives us a direct indication of the net public contribution as a proportion of the net market income lost due to the lockdown. Furthermore, we decompose the compensation rate in the same way as the Net Replacement Rate to highlight the contribution of each group of benefits.

In order to test whether the income stabilisation offered by the tax-benefit systems prevents those affected by the lockdown from falling below an absolute income threshold, we compare the equivalised disposable income before and after the lockdown to the poverty threshold at 60% of the median in the pre-shock baseline, without and with the discretionary policy measures implemented by the government.

Our approach is equivalent to calculating absolute poverty rates with a fixed poverty line and resembles the suggested practice in the measurement of poverty during an economic crisis using a threshold fixed in real terms (Jenkins et al., 2013). Such an indicator can be considered as an appropriate proxy for the experience of impoverishment that an individual faces, comparing their current condition with their own status before the income shock (Matsaganis and Leventi, 2011). A normative judgment of the proper level of protection provided by the welfare systems is beyond the scope of this paper and should be evaluated considering the minimum levels of living standards guaranteed by the welfare system as a whole (Boadway and Keen, 2000). However, given the policy goal of limiting the numbers of individuals at risk of poverty, it is implicit that household income of those affected by the lockdown should not fall below the poverty threshold.

Before moving to the results of the empirical analysis, it is important to reiterate that we consider the hypothetical situation of one month in isolation

Covid Economics 8, 22 April 2020: 106-133

only (i.e. March 2020, when the government imposed the lockdown and the first compensation measures have been implemented). Our considerations abstract from the smoothing possibilities of the income shock that an individual can exploit over a longer period of time.

Furthermore, our main indicators – Net Replacement Rate, Compensation Rate and poverty status of individuals affected by the shutdown – refer to the set of individuals identified as those affected by the earning loss and depend on their characteristics and the assumptions we have made on 100% benefit take-up. These indicators are not affected by the absolute numbers of individuals identified as those affected by a loss of earnings. As opposed, estimates of budgetary costs and those of poverty and inequality in the overall population are affected instead by the absolute numbers of individuals considered and this should be borne in mind when interpreting the results.

Covid Economics 8, 22 April 2020: 106-133

3. The characteristics of those affected by earning loss

The analysis focuses on employed and self-employed individuals who lost their earnings in the immediate aftermath of the COVID-19 outbreak.

We consider economic sectors at 6-digit level, as classified by ATECO, that were listed in the Decree Law imposing the shutdown of economic activities.[1] Although SILC microdata lack information on business activities at 6-digit level, we draw on other detailed available statistics released by Istat (namely, the operating firms archive (ASIA), the national labour force survey (RCFL) and National Accounts) in order to compute the occupation shares in each sector subject to shut down.

The left enclave in Table 1, based on Istat detailed statistics, shows that[2] 39% of Italian active workers are subject to the shutdown, on average. The shares of workers affected are different across economic sectors: while more than 60% of the active workers in the manufacturing and construction sectors are affected, the shares corresponding to affected workers in the wholesale and retail trade sectors, as well as accommodation and food service activities are of more than 80%. All workers in real estates, arts, entertainment and recreation activities are affected by the shutdown of economic activity.

We then randomly select the individuals, with a positive income source from either employment or self-employment. We perform this selection by sector

[1] (Decree Law of the Minister of Economic Development which updates the DPCM 22/3/2020 available here https://www.gazzettaufficiale.it/eli/id/2020/03/26/20A01877/sg

[2] For the sake of simplicity, the table reports economic sectors at 1-digit level as per the ATECO classification.

of employment at 2-digit ATECO level, which we relate to data in EUROMOD in order to get the same occupation shares subject to shut down. On the other hand, in EUROMOD we identify 27 million individuals with a positive income source from employment, temporary jobs or self-employment reported in the income reference year (i.e. 20016). As expected, this figure is higher than the 23 million individuals reported by Istat which refers to those with regular employment contracts.

Moreover, shares shown in the right enclave of Table 1, do not always correspond to those on the left. They can be lower as the salaries of individuals working in the public sector are not affected by the shutdown (and hence not selected in EUROMOD) but can also be higher as the number of individuals observed in some sectors is too limited to select the right amount from the left enclave.

Covid Economics 8, 22 April 2020: 106-133

Table 1. Workers subject to shutdown by sectors of economic activities

	Economic activity	ISTAT Workers	ISTAT Workers subjects to shut down		EUROMOD Workers	EUROMOD Workers subjects to shut down	
		thousands	thousands	%	thousands	thousands	%
A	AGRICULTURE, FORESTRY AND FISHING	909	55	6.02	1,167	15	1.32
B	MINING AND QUARRYING	25	15	60.65	81	58	71.79
C	MANUFACTURING	4,321	2,825	65.38	5,087	3,627	71.30
D	ELECTRICITY, GAS, STEAM AND AIR CONDITIONING SUPPLY	114	0	0.00	135	0	0.00
E	WATER SUPPLY; SEWERAGE, WASTE ACTIVITIES	243	0	0.00	181	0	0.00
F	CONSTRUCTION	1,339	806	60.17	2,022	1,230	60.80
G	WHOLESALE AND RETAIL TRADE; REPAIR OF VEHICLES	3,287	2,711	82.48	3,804	3,220	84.66
H	TRANSPORTATION AND STORAGE	1,143	0	0.00	1,322	0	0.00
I	ACCOMMODATION AND FOOD SERVICE ACTIVITIES	1,480	1,271	85.86	1,522	1,323	86.93
J	INFORMATION AND COMMUNICATION	618	0	0.00	562	0	0.00
K	FINANCIAL AND INSURANCE ACTIVITIES	636	0	0.00	839	0	0.00
L	REAL ESTATE ACTIVITIES	164	164	100.00	114	113	99.52
M	PROFESSIONAL, SCIENTIFIC AND TECHNICAL ACTIVITIES	1,516	78	5.15	1,909	69	3.60
N	ADMINISTRATIVE AND SUPPORT SERVICE ACTIVITIES	1,028	362	35.22	902	282	31.21
O	PUBLIC ADMINISTRATION, DEFENCE; SOCIAL SECURITY	1,243	0	0.00	1,680	0	0.00
P	EDUCATION	1,589	0	0.00	2,107	0	0.00
Q	HUMAN HEALTH AND SOCIAL WORK ACTIVITIES	1,922	0	0.00	2,125	0	0.00
R	ARTS, ENTERTAINMENT AND RECREATION	318	318	100.00	268	221	82.54
S	OTHER SERVICE ACTIVITIES	712	523	73.50	895	740	82.62
T	ACTIVITIES OF HOUSEHOLDS AS EMPLOYERS	739	6	0.75	421	17	4.12
U	ACTIVITIES OF EXTRATERRITORIAL ORGANISATIONS	14	14	100.00	8	1	14.06
		23,360	9,148	39	27,151	10,916	40.21

Notes. Our elaboration using ASIA, RCFL and National Accounts and SILC data.

Overall, we identify 11 million workers potentially at risk of losing their earnings as they are active, with a private employer, in one of the economic sectors subject to the shutdown.

We plan to extend this analysis as soon as administrative data (COB) or Labour Force Survey data are made available where one can identify those who actually suffered the income loss.

Table 2 reports some characteristics of those affected by the economic shutdown: 37% of them lives in households with some children; 41% of them come from one-earner households and for them the temporary shutdown of their activities imply the loss of the main income source.

The distribution of those affected by the lockdown by household income quintile groups (assessed before the earning loss) shows an increasing pattern with quintile shares ranging from 15% at the bottom of the distribution to 24% at the top.

Table 2. Characteristics of those affected by earnings losses

Presence of children %	36.60
Number of earners %	
1	40.74
2	42.24
3+	13.35
Household income quintile %	
Bottom	14.76
2nd	16.62
3rd	20.99
4th	23.50
Top	24.13

Notes: Summary statistics for those affected by income losses as identified in EUROMOD data. Quintile groups based on household equivalised disposable income in the baseline. Source: EUROMOD version I2.0+.

4. Income protection policies

The existence in all European countries of a developed welfare state (Schubert et al., 2009), that is intended, among other things, to protect people and their families against economic shocks, is one of the main differences between the

Covid Economics 8, 22 April 2020: 106-133

crisis faced today and that of the 1930s. However, the sudden and unexpected shock due to the COVID-19 pandemic forced European governments to adapt existing measures and to define new discretionary and bold measures in order to support those who are bearing a disproportionate share of the economic burden (OECD, 2020)

Table 3 provides a summary of the most important measures implemented by the Italian government, including the Decree Law 18/2020 ("Cura Italia") to support individuals and their families.. The same Decree Law imposes that firms cannot fire employees after February 23, 2020: this implies that existing Unemployment Insurance Schemes do not apply to workers affected by an earning loss due to the COVID-19 pandemic.

In order to compensate the earning loss suffered by the employees, the government extended the existing Wage Supplementation Scheme (i.e. *Cassa Integrazione Guadagni, CIG*) relaxing the eligibility conditions and allowing most of employees to be entitled to the scheme. Only domestic workers and consultants (i.e. *parasubordinati*) are not eligible. The Wage Supplementation Scheme provides a replacement of 80% of earnings subject to a maximum cap, which is fully covered by the National Institute of Social Security (INPS). As INPS payments usually take 2 or 3 months, in an attempt to limit delays, the government reached an agreement with commercial banks that anticipate the transfers on behalf of the government and disburse the owed amounts to entitled workers. If monthly earnings are below 2,160 euro, *Cassa Integrazione Guadagni* cannot exceed 940 euro, while if earnings are above the threshold the *Cassa Integrazione Guadagni* is capped at 1,130 euro. This implies that the replacement can be substantially below 80% for most workers. The government expects to transfer up to 3.4 billion euros on this scheme,in addition to 1.7 billion euros for figurative contributions. This amount represents the maximum expense allowed by the government and transfer payments are subject to income taxes.

In order to compensate the earning loss incurred by the self-employed, the government defined a new lump-sum transfer of 600 euro to be paid for the month of March to *all* self-employed, irrespective of whether they incurred a loss or not. The self-employed in specific professional bodies (e.g. lawyers, accountants, notaries, etc.) are eligible for the lump-sum transfer only if their 2019 income was below 35,000 euro. Rules are such that self-employed must apply for this transfer, and there has been a delay in the processing times due to the high volume of applications with the tax authority, INPS, so that the first transfers reached beneficiaries in mid April. The estimated maximum

Covid Economics 8, 22 April 2020: 106-133

CEPR PRESS

binding expenditure for the first month is roughly 3.1 billion euros. The transfer is not subject to income tax and does not enter in any means-test of other benefits.

Employees bound to continue work on company premises and those who cannot typically work from home are entitled to a lump-sum transfer of 100 euro to be paid for the month of March. We arbitrarily assume that 50% of employees working in the economic sectors that are not subject to the shutdown still work on company premises. The estimated maximum binding expenditure is about 0.8 billion euros. The transfer is not subject to income tax and does not enter in any means-test of other benefits.

Self-employed can ask to suspend the mortgage on their main residence.[3]

In addition to the policies listed in Table 3, the government allowed employees in the private sector with children up to 12 years old to take parental leave for 15 days at 50% of the earnings' level or, alternatively, to have a babysitting bonus of 600€ (incremented to 1000€ for those working in the health system). We do not simulate these measures due to data unavailability but we focus instead on simulations involving the realistic take-up of these schemes.

[3] This is a reduction in current expenditures, which in our simulations is considered as a transfer. Arguably, other naturally reduced costs (e.g. commuting or childcare costs) should have received the same treatment but we decided to consider this expenditure solely because it is the only one clearly defined by the Decree Law and properly guaranteed for by a Fund that covers such expenditures.

CEPR PRESS

Table 3. Simulated policies introduced by the Decree Law 18/2020

Measure	Estimated cost (billion euros)	Target
Wage Supplementation Schemes (i.e. CIG)	3.4 + 1.7 (CIG cost + figurative contributions)	Salary workers excluding temporary workers and housekeeping workers
Lump sum transfer (600€)	3.1	Self-employed (if enrolled in professional body, subject to income limit equal to 35.000€)
Lump sum transfer (100€)	0.8	Employees working on company premises, subject to income limit equal to 40.000€)
Mortgage suspension		Self-employed

5. Empirical evidence

In our simulations we assume that all individuals working in sectors subject to the shutdown benefit from the discretionary policy measures described above.

Table 4 reports the simulated costs and the number of entitled individuals for each measure, considering only one month of application of the different schemes.

The Wage Supplementation Scheme would cost around 5.6 billion euros (plus 2.8 billion euros of credit contributions) with 7 million workers benefitting from it. The lump sum for the self-employed would cost 1.4 billion euros involving 2.4 million individuals. Five million workers would benefit from the lump sum of 100€ with a total cost of 0.5 billion euros.

The simulated costs are somehow different from those estimated by the government and ratified by the Parliamentary Fiscal Council (UPB, 2020), reported in Table 3. This has to do with how we define the individuals entitled, which we related to the take-up of benefits. The government assumes an average take-up rate of around 80% uniform across economic sectors, while

Covid Economics 8, 22 April 2020: 106-133

the Fiscal Council assume differentiated take-up rates across sectors with an overall average of around 60%. In our simulation we assume that 100% of individuals working in the sectors affected by the shutdown down are entitled to the Wage Supplementation Schemes and the Lump sum transfer (600€) and they do take-up these benefits. We assume that 50% of those employed in sectors not subject to the lockdown are still working on company premises (Fondazione Studi Consulenti del Lavoro, 2020) and they receive the lump sum transfer (100€).

Depending on how reliable our identification of the sectors subject to the shutdown is, our scenario can be considered as an upper-bound scenario in terms of the individuals entitled to receive the benefits and the overall cost of the measures. We assume that all individuals working in the sectors subject to the shutdown are negatively affected (i.e. they lose their earnings) but there could be individuals still working due to specific waivers. On the other hand, there could be individuals working in the sectors not subject to the shutdown who are negatively affected and we are not able to identify them.

Table 4. Policies introduced by the Decree Law 18/2020: simulated costs and entitled individuals

Policy	Simulated cost: billion euros	Simulated cost: % of annual GDP	Entitled: thousands
Wage Supplementation Schemes	5.6	0.31	7,013
- Figurative Social Contributions	2.8	0.16	
Lump sum transfer (600€)	1.4	0.08	2,360
Lump sum transfer (100€)	0.5	0.03	4,962
Mortgage subsidy	0.15	0.01	363

Notes: Costs refer to a one-month application of the different schemes. Workers entitled to Wage Supplementation Schemes are individuals with positive employment income, working in sectors subject to the shutdown and not in the public sector. Workers entitled to a lump sum transfer (600€) are individuals with positive self-employment income, working in sectors subject to the shutdown and not receiving employment income. Workers entitled to the lower lump sum transfer (100€) are 50% of the individuals with positive employment income, working in sectors not subject to the lockdown (randomly selected and arbitrarily assumed). Source: EUROMOD I2.0+

Covid Economics 8, 22 April 2020: 106-133

Overall, a one-month shutdown imposed by the government would imply a loss of original income of around 20 billion euros, representing 1.1% of annual GDP and around 33% of observed original income before the shutdown. With such a loss of original income, the government would lose 2.7 billion euros of income tax revenue and 5.9 billion euros of social security contributions (including both employer and employee contributions). Despite additional 7.6 billion euros of transfers (i.e. Wage Supplementation Scheme and lump sum transfers), the loss of disposable income for the families affects by the economic shutdown is around 8 billion euros or 12% of the observed disposable income before the shock.

Table 5. Income changes due to the economic shutdown

Income source	billion euros	% of annual GDP	% change
Original income	-20.2	-1.13	-32.75
Social security contribution employer	-4.0	-0.22	-31.02
Social security contribution employee	-1.9	-0.11	-32.23
Income tax	-2.7	-0.15	-16.38
Transfers	7.6	0.43	27.39
Disposable income	-7.9	-0.44	-11.86

Notes: Income changes refer to one-month shutdown. Source: EUROMOD I2.0+

Figure 1 shows the unequal distribution of income losses along quintile groups. Original income losses are more pronounced at the bottom of the distribution: those in the first quintile group would lose more than 40% of their original income while those in the top quintile group less than 30%. This is due to the fact the one-earner families are more concentrated at the bottom of the distribution and the shutdown causes the loss of their main income sources. Along the income distribution, families are characterised by more earners and other income sources (e.g. property and capital income) not affected, in the short term, by the economic shutdown.

Covid Economics 8, 22 April 2020: 106-133

CEPR PRESS

Figure 1. Income losses due to the economic shutdown, by household income quintile groups.

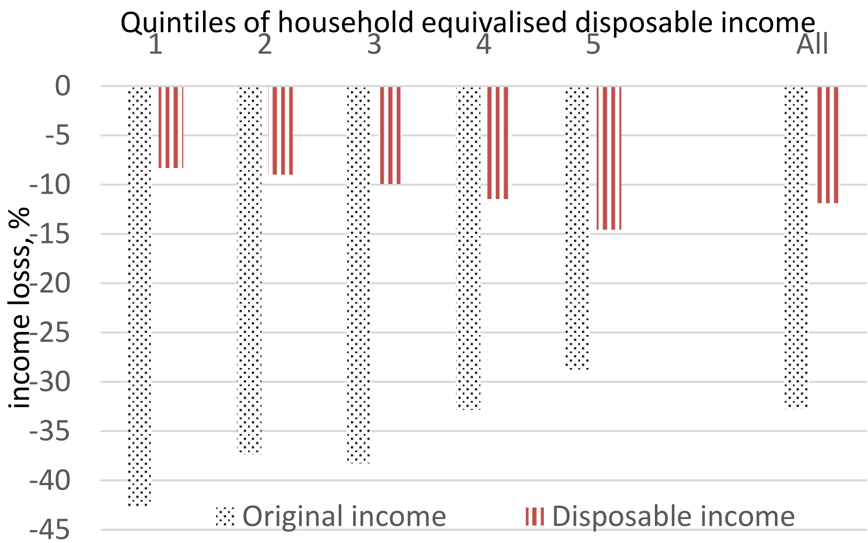

Source: EUROMOD I2.0+

Due to these income changes that also hide re-rankings of individuals moving to the bottom part of the distribution when they lose their earnings, one can expect a different level of inequality in the income distribution after the shock. The Gini of the disposable income distribution is equal to 0.31 before the shutdown and 0.33 after the shutdown, highlighting a non-negligible increase in inequality, explained by a larger role of between population groups inequality, namely those affected and those not affected by the shutdown. Without the policy measures introduced by the government the inequality level in disposable income would have been higher , with Gini equal to 0.42.

5.1. Relative resilience

The average Net Replacement Rat is illustrative of the relative resilience due to differences in tax-benefit systems, characteristics of the individuals affected by the shutdown and household composition.

Household income on average falls to as much as 78% of its pre-shock level considering all households with at last one individual affected by the lockdown.

Covid Economics 8, 22 April 2020: 106-133

The protective role played by Original Income (including earnings of other household members) is illustrated in Figure 2 which shows the Net Replacement Rates by its components (with Taxes and Contributions reducing the Replacement Rates and hence negative) and by household income quintile groups. Income from other benefits (i.e. mainly pensions, disability benefits and income-tested benefits) plays a similar but smaller role. The sum of these two components makes up around 60% of post-shock household income, almost constant along the income distribution, with the original incomes less relevant at the bottom of the income distribution and vice versa for the other benefits.

Earnings of other household members are progressively more important as household income increases: the average Net Replacement Rates are likely to be pushed up by the presence of these incomes at the top of the income distribution, but this is partly compensated by progressive income tax. Wage Supplementation Benefits play a large role ranging from 20% to 28% of post-shock household income, with an inverted U shape along the income distribution. COVID benefits are clearly relevant at the bottom of the distribution where they represent almost 20% of post-shock household income.

The general lesson of this analysis is that it is necessary to consider the social protection system as a whole and how it interacts with household composition and incomes received by other household members. Focusing exclusively on discretionary measures is not enough to have a comprehensive picture.

Covid Economics 8, 22 April 2020: 106-133

Covid Economics 8, 22 April 2020: 106-133

Figure 2. Decomposition (by income sources) of Net Average Replacement Income for those affected by the lockdown, by household income quintile groups

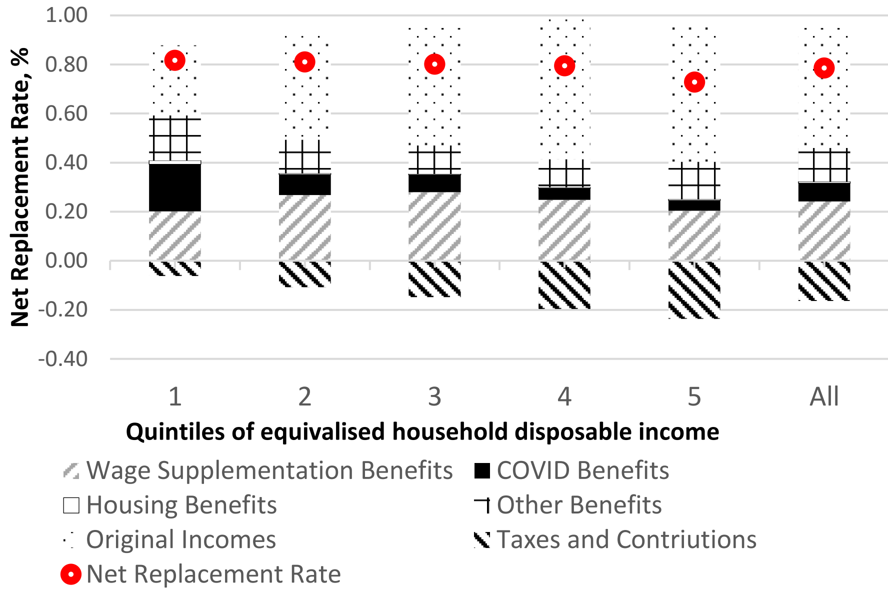

Notes: Net Replacement Rate is the ratio of household disposable income after and before the earning shock. "COVID Benefit" include newly discretionary policies such as lump sum transfers to self-employed and employees; "Housing Benefits" include the amount equivalent to the mortgage instalment for the main residence; "Other Benefits" include pension and invalidity benefits, minimum income schemes, family benefits; "Taxes and Contributions" include personal income tax, employee social insurance contributions and other direct taxes. Source: EUROMOD I2.0+.

To focus on the income protection offered by public support, we adopt the Compensation Rate approach. It shows that the average net public contribution to the disposable income as a proportion of the net earnings lost because of the lockdown is around 55% with a decreasing pattern along the income distribution (Figure 3).

Most public support is channelled through the Wage Supplementation Scheme of benefits (the shaded area with a forward-sloping line pattern) only slightly reduced by the progressive income tax (the shaded area with a backward-sloping line pattern) payable on these benefits. Benefits received due to COVID-19 make up the largest share of public support at the bottom of the

Covid Economics 8, 22 April 2020: 106-133

distribution but represent a non-negligible compensation for those in the upper part of the distribution as well.

Families in the first quintile group benefit relatively more from COVID benefits as individuals entitled to these lump-sum transfers (i.e. self-employed and occasional workers) have more representation in this group, with original income relatively low compared to the 600 € lump-sum transfer. The Compensation Rate decreases with income because the Wage Supplementation Schemes represent a decreasing income replacement, given that it is capped at 1,130 euro.

Figure 3. Decomposition (by income sources) of average Compensation Rates for those affected by the lockdown, by household income quintile groups

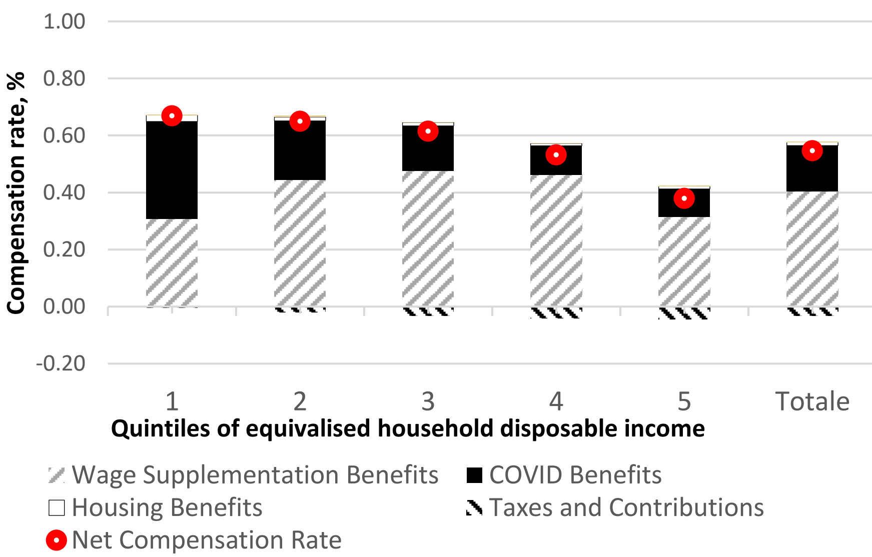

Note. See Figure 1. Quintile groups based on disposable income before the pandemic. The lump sum of 100€ to the employees is not included in the Compensation Rate because it is given to employees who are not subject to a reduction in their original income. In order to avoid the impact of outliers, the sample is restricted to employees with a Compensation rate between 0 and 1 and to self-employed with income larger than 50€ per month. The Figures reports individual averages which are not strictly comparable with numbers behind Figure 1 which are aggregates at quintile levels. Source: EUROMOD version I2.0+.

5.2. Absolute resilience

The extent to which the tax-benefit instruments allow those affected by the shutdown to avoid falling below a given level of income depends on the generosity of the system, whether workers are entitled to receiving wage supplementation benefits and COVID benefits, the income position of the individuals before losing their earnings and their household circumstances.

Table 6 shows the poverty rates, for different groups of the population, in three different scenarios: (1) before the shutdown due to the COVID-19 pandemic, (2) after the shut-down without considering the compensation policies implemented by the government and (3) after the shut-down considering the discretionary policies introduced by the government. The poverty line is always constant as in the scenario before the shutdown.

Focusing on the workers active in sectors subject to the shutdown, the share of those at risk of poverty before the shock is around 13%. The impact of the shutdown alone is disruptive with the poverty rate that would have reached 68% of workers without any compensation measure. The policies implemented by the government are able to limit such an impact, limiting the poverty rate at 28%.

The individuals living in one-earner families are, as expected, more exposed to poverty risk: 22% are poor already before the COVID-19 pandemic, 80% would have been in a poverty status without compensation measures and 44% are below the poverty threshold with the discretionary policies in operation.

When extending the analysis to the overall population and considering the compensation measures implemented by the government, the breakthrough impact of the pandemic on the poverty status is evident, with an increase in the poverty rate of more than 8 percentage points, and of more than 13 percentage points when we focus children.

Table 6. Poverty rates before and after the COVID-19 pandemic

	Before COVID-19	Shut-down, without compensation policies	Shut-down, with compensation policies
Workers in sectors subject to shut down	12.53%	67.97%	28.15%
Workers in sectors subject to shut-down and living in one-earner families	22.13%	80.49%	43.71%
Overall population	19.07%	38.41%	27.28%
Children	23.27%	49.63%	36.34%

Notes: The poverty threshold is fixed at 60% of baseline median household disposable equivalised income. Poverty rates based on household equivalised disposable income. Source: EUROMOD version I2.0+

6. Conclusions

We have analysed the extent to which the Italian tax-benefit system provides income support to those affected by the economic shutdown at the beginning of the COVID-19 pandemic.

In order to assess the impact of both the existing and the newly designed benefits on household income, counterfactual scenarios are simulated with EUROMOD, the EU-wide microsimulation model, integrated with information from the activity sectors subject to the economic shutdown.

In interpreting our results there are some caveats to be borne in mind. Most importantly, our paper offers a scenario rather than a forecast and it provides a reference point by which one can evaluate the economic unfolding of the situation and the new policies that will be implemented.

Moreover, our analysis entails potential economic effects of the first month of the COVID-19 pandemic and examines the extent of the intended effects of the schemes, though in reality the transfer payments (i.e. wage supplementation and the emergency lump-sum transfers) were inevitably delayed and this lag might constrain the liquidity of families. In order to limit the delay, the government reached an agreement with commercial banks that anticipate the transfers corresponding to the Wage Supplementation Schemes

Covid Economics 8, 22 April 2020: 106-133

and disburse the owed amounts to the entitled workers. With that said, our analysis abstracts from any possibility of income and consumption smoothing that individuals can exploit over a longer period of time. Individual preferences for consumption smoothing lead, for instance, to a decrease in current consumption in the presence of economic insecurity. Consequently, the overall effects of the crisis would be exacerbated if the government does not provide immediately an income stabilisation for those who actually experience earning loss, which can potentially translate into detrimental effects on the aggregated demand.

Based on our scenario, one can expect a loss of market income as related to individuals of more than 30%, only partially compensated by new policy measures which tend to guarantee to a larger extent the income of those at the bottom of the distribution. Nevertheless, an increase in the overall inequality and poverty risk is expected, amounting to 15 percentage points among individuals affected by the shut-down and to more than 8 percentage point considering the overall population.

It is clear that the effects of the COVID-19 pandemic are asymmetric and particularly relevant from an economic perspective for some families and less for others, despite the compensation measures implemented by the government. It is crucial to take into account such unequal distribution of the shock if the economic consequences are expected to last long.

As clearly pointed out by Sacchi (2018) while reforms occurred since 2012 have modernised the Italian welfare system "this does not mean that it is necessarily ready for the challenges has to face". In particular, the first month of the COVID-19 pandemic highlights important deficiencies of the Italian welfare system.

That is, the most important automatic stabilizers embedded in the tax-benefit system (i.e. Minimum guaranteed income - RdC, Family allowances – ANF and in-work bonus – Bonus IRPEF) depend on past year's incomes and do not react to a sudden loss of earnings such as those experienced in March 2020. Moreover, some of the welfare tools deployed during the emergency, such as the lump sum transfer of 600€ to self-employed, do not seem to be well-thought in terms of size and design as they provide equal transfers to all entitled while ignoring the possibility of individuals having historically declared lower incomes than the one transferred in March 2020 and preventing full coverage, with domestic workers being excluded.

Covid Economics 8, 22 April 2020: 106-133

At the time of writing this paper, the Italian government has decided that (i.e. bookshops, baby clothes shops, …) some commercial activities previously subject to the shutdown (i.e. en-detail retail such as book and stationary shops, children's clothing, etc.) can reopen starting from mid of April 2020 and is currently writing a new Decree Law with new and more generous compensation measures, including a new "emergency income" which should help protect individuals from income losses.

In order to avoid an increase in inequality and poverty two national think-tanks, Forum Diseguaglianze Diversità and ASviS, suggest implementing two extraordinary and temporary instruments: (1) the so called *Sostegno di Emergenza per il Lavoro Autonomo* (SEA – Emergency Support of Self-Employment) - an income support that takes into account the economic conditions of the household of the self-employed who lose their job – and (2) *Reddito di Cittadinanza per l'Emergenza* (REM - Emergency Citizenship Income) – a last safety net for those not covered by other instruments based on the design of the Citizenship Income (Forum Diseguagliane Diversità and ASviS, 2020). These measures would allow the country to have a systematic set of instruments to support incomes in the short term and allow the government to focus on the actions needed for the medium- and long-term economic recovery

In general terms, our analysis has demonstrated the importance of the income of other household members in determining the economic resilience of those affected by the shutdown. The sharing of risks within the household can be seen in general terms as a complement to the insurance function of the Welfare State. However, as it is usual in distributive analysis, we have assumed complete income pooling within the household. The possibility that incomes are not in fact pooled serves to remind us of the non-equivalence of income received in the form of Wage Supplementation Schemes as an individual entitlement on the one hand, and income support schemes, usually assessed on the economic situation of the family as a whole, on the other.

Finally, we believe that the stress test approach applied to tax-benefit schemes offers some potential opportunities for further research.

First, we will trace the evolution of the effects of the shutdown on the labour market in the context of the COVID-19 pandemic and will monitor the effects of the compensation schemes enacted by Italian fiscal authorities on household incomes.

Covid Economics 8, 22 April 2020: 106-133

Second, we will extend our analysis to the most important European economies to capture the heterogeneous effects of the COVID-19 asymmetric shock across other European welfare systems. In a cross-country perspective, it will be important to understand how well-suited existing institutional arrangements are for compensating income loss during the pandemic. Moreover, such evidence will raise normative issues on the protection level that the tax-benefit system should guarantee to the population and backs up the idea that unconditional Basic Income instruments would have made comprehensive compensation possible during the pandemic, without the need of discretionary and temporary policies (Atkinson, 2015).

Covid Economics 8, 22 April 2020: 106-133

References

Atkinson A. B., "Stress-Testing the Welfare State," in B. Ofstad, O. Bjerkholt, K. Skrede and A. Hylland (eds), *Rettferd og Politik Festskrift til Hilde Bojer*, Emiliar Forlag, Oslo, 31-39, 2009.

Atkinson, A. B., "Inequality: What can be done?", Harvard University Press, Cambridge, 2015.

Boadway R. and M. Keen, "Redistribution," in A.B. Atkinson and F. Bourguignon (eds), *Handbook of Income Distribution*, Elsevier Vol. 1, chapter 12, 677-789, 2000.

Ceriani L., F. Figari, and C. Fiorio, "EUROMOD Country Report, Italy 2016-2019", 2019, https://www.euromod.ac.uk/using-euromod/country-reports

Dolls M., C. Fuest, and A. Peichl, "Automatic Stabilizers and Economic Crisis: US vs. Europe," *Journal of Public Economics,* 96, 279-294, 2012.

Dorn F., C. Fuest, M. Göttert, C. Krolage, S. Lautenbacher, S. Link, A. Peichl, M. Reif, S. Sauer, M. Stöckli, K. Wohlrabe and T. Wollmershäuser, "The Economic Costs of the Coronavirus Shutdown for Selected European Countries: A Scenario Calculation", EconPol Policy Brief 25, April 2020, https://www.econpol.eu/publications/policy_brief_25

Figari F., A. Paulus and H. Sutherland, "Microsimulation and Policy Analysis", in Handbook of Income Distribution Volume 2B, edited by A. B. Atkinson and F. Bourguignon, Elsevier, 2015.

Fernandez Salgado, M., F. Figari, H. Sutherland and A. Tumino, "Welfare compensation for unemployment in the Great Recession", *Review of Income and Wealth* 60, S177-S204, 2014.

Fiorio, Carlo V., and Francesco D'Amuri. "Workers' Tax Evasion in Italy." *Giornale degli Economisti e Annali di Economia* 64 (Anno 118), no. 2/3 (2005): 247–70.

Fondazione Studi Consulenti del Lavoro, "Emergenza COVID-19: l'impatto su aziende e lavoratori secondo i Consulenti del Lavoro", Roma, 2020

Forum Diseguagliane Diversità and ASviS, "Curare l'Italia di oggi, guardare all'Italia di domani", https://tinyurl.com/uarjkd7, 2020, Istat, “Monthly Report on the Italian Economy – March 2020”, Istat, Rome, 2020

Immervoll H. and C. O’Donoghue, “What Difference Does a Job Make? The Income Consequences of Joblessness in Europe” in D. Gallie (eds), *Resisting Marginalisation: Unemployment Experience and Social Policy in the European Union*, Oxford University Press, Oxford, 105-139, 2004.

Jenkins S. P., A. Brandolini, J. Micklewright and B. Nolan, *The Great Recession and the Distribution of Household Income*, Oxford University Press, Oxford, 2013.

Jones M.T, P. Hilbers and G. Slack, “Stress Testing Financial Systems: What to Do When the Governor Calls,” IMF Working Paper, 04/127, 2004.

Matsaganis M. and C. Leventi, “The Distributional Impact of the Crisis in Greece,” EUROMOD Working Paper Series, 3/11, 2011.

OECD, “Evaluating the initial impact of COVID-19 containment measures on economic activity”, Economics Department, OECD, Paris, 2020a

OECD, “Supporting people and companies to deal with the Covid-19 virus: options for an immediate employment and social-policy response”, ELS Policy Brief on the Policy Response to the Covid-19 Crisis, OECD, Paris, 2020b

Saez E. and G. Zucman, “Keeping business alive: the government will pay”, 18 march 2020 www.socialeurope.eu, 2020.

Sacchi S., The Italian Welfare State in the Crisis: Learning to Adjust?, South European Society and Politics, 23:1, 29-46, 2018.

Sutherland H., and F. Figari “EUROMOD: the European Union Tax-Benefit Microsimulation Model,” International Journal of Microsimulation, 6(1) 4-26, 2013.

Covid Economics 8, 22 April 2020: 106-133

Covid Economics **Issue 8, 22 April 2020**

National containment policies and international cooperation[1]

Thorsten Beck[2] and Wolf Wagner[3]

Date submitted: 13 April 2020; **Date accepted:** 17 April 2020

Covid Economics 8, 22 April 2020: 134-148

Policies that curtail social and economic activities during a pandemic are predominantly decided upon at the national level, but have international ramifications. In this paper we examine what type of inefficiencies this may create and how cooperation across countries may improve outcomes. We find that inefficiencies arise even among completely identical countries. We show that countries are likely to choose excessively lenient policies from the perspective of world welfare in later stages of the pandemic. This provides a rationale for setting minimum containment standards internationally. By contrast, in early and intermediate stages of the pandemic, national containment policies may also be excessively strict. Whether or not this is the case depends on a country's degree of economic integration relative to (outward and inward) mobility of people. Analyzing the stringency of containment policies during the current epidemic confirms that countries with higher economic integration adopt stringent containment policies more quickly whereas countries subject to high mobility do so later.

1 We thank Robin Döttling and Dion Bongaerts for dicussions on the topic.

2 Professor of Banking and Finance, Cass Business School and CEPR Research Fellow.
3 Professor of Finance, Rotterdam School of Management and CEPR Research Fellow.

CEPR PRESS

Covid Economics 8, 22 April 2020: 134-148

1 Introduction

Governments around the world have responded with unprecedented measures to the spread of the COVID-19. Borders are now closed for travel among most major economies. Drastic policies have been introduced to curtail economic and social activity. The early evidence so far suggests that they have indeed helped to limit the spread of the virus, but at the same time there are likely to be significant economic consequences.

What is striking though is that even though the pandemic is clearly an international problem (as the spread of the virus from one country through nearly the entire world has shown), the policy responses are up to that moment entirely national. This suggests a potential tension. In an integrated world, uncoordinated national decisions of such severity are unlikely to produce a desirable outcome. Recently, policy makers have started to become aware of the issue. For example, the European Commission is starting to work on an exit-plan across the EU. However, little is known so far of what type of problems national containment policies create, if any. Even less is known about appropriate policy responses to such problems. How could a coordinated policy look like? Should it strive to increase the severity of containment or rather encourage countries to reverse them quickly? Should policies differ across countries, and the dynamics of the pandemic?

In this paper we analyze national containment policies in an integrated world. We examine two countries that independently choose their containment policies. We consider two stages of the pandemic, a *lockdown stage* with an initially severe spread of the virus and closed borders, and a *new normal* stage where the virus is fairly contained (but still alive and kicking!), and borders have been reopened. The two stages interact because containment policies in the first period affect the initial (pandemic) conditions in the second stage, and through this potentially optimal policies and welfare

In our model, countries independently choose their "activity-level" which can be given an economic or social interpretation. Activities provide net benefits to a country, but also facilitate the spread of the virus. Crucially, a country's activity choice has international repercussions, through economic channels and potentially also through spread of the virus across borders. We show that there are generally distortions in national decisions, but that the direction and the intensity of the distortions depends on the stage of the pandemic as well as country characteristics. At an early stage, economic externalities from containment policies may dominate as countries incur significant losses due to disruptions in the supply chain. This implies that individual countries may choose containment policies that are excessive from the international perspective, as they do not internalize their negative consequences on the economies abroad. However, due to the dynamic impact of containment policies (arising because containing the virus domestically also affects the *new normal*, and through this also the other country), we show that countries may also choose excessively lenient policies, even in the early stage.

Over time, economic externalities decline as firms in other countries adjust to new conditions. In addition, the process of opening borders makes a second dimension of the externality more important, arising from international travel. This externality works in the opposite direction. A country that implements stricter containment measures benefits other countries because, by reducing the number of infecting people in its own country, there is less likely a spread of the virus to other countries.

Overall, we show that national containment policies are more likely to be excessive in initial stages of the pandemic, and for countries that are well internationally integrated. In later stages of the pandemic, by contrast, containment policies may be excessively lenient, and in particular so in countries that have a high mobility (inward and outward), such as popular tourist destinations. Based on these results we discuss various policy responses, such as minimum containment standards at the international level in later stages of the pandemic, as well stimulating higher stringency at individual countries through subsidies in earlier stages of the pandemic.

The central ingredients in our model follow closely the literature on the benefits and cost of international cooperation, which has emphasized cross-border externalities as a rationale for cooperation. For example, there is a clear potential benefit for macroeconomic policy coordination as both fiscal and monetary policy have effects beyond the country where they are instituted (see, for example, Cooper, 1969, and Hamada, 1976, for fiscal and monetary policy coordination, respectively). An example for such coordination is the Plaza accord of 1995. Cooperation efforts have more recently also been stepped up

Covid Economics 8, 22 April 2020: 134-148

Covid Economics 8, 22 April 2020: 134-148

following the Global Financial Crisis (for a review of the extant literature on macroeconomic policy coordination, see Frankel (2015)). The economic externality present in our paper is similar to the externalities arising from macroeconomic policies; the international "virus-externality", however, is not commonly considered in economics.[1]

Cooperation has also costs, arising from country heterogeneity. The literature on fiscal decentralization (see, for example, Oates, 1972) argues that the comparative advantage of centralization increases with the size of interjurisdictional externalities but decreases with preference heterogeneity.[2] A similar trade-off also applies to optimal currency areas and trade-blocs. Following Mundell (1961), a common currency can reduce spillovers from *beggar-thy-neighbour* policies. However, a cost of having a common currency is that countries are subject to different shocks (Mundell, 1961), hence their "optimal" exchange rate differs (Mundell, 1961, and Maloney and Macmillen, 1999). Dell'Arricia and Marquez (2006) consider cooperation in banking supervision. They show that the gains from delegating supervisory decisions to a supranational agency increase in cross-border externalities but decrease in heterogeneity across countries arising from preferences. Beck, Silva Buston and Wagner (2016) provide (indirect) evidence for such a trade-off by showing that countries are more likely to cooperate in the supervision of their banks when there are large bilateral externalities and when countries are similar to each other. In contrast to the majority of the literature, in this paper we mostly abstract from issues arising from country-heterogeneity.[3]

Within a very short time frame, the advent of the Corona-pandemic has spurred important contributions from policy makers and academics (see for example, the VoxEU-book on the "Economics in the Time of Covid-19", and the "Covid Economics" journal of CEPR). Perhaps most closely related to our paper is the analysis in Eichenbaum, Rebelo and Tra-

[1] Beyond economic cooperation, there is also a large tradition of successful cooperation in public health policies; for a discussion see Cooper (2001).

[2] Rogoff (1985) shows that cooperation also has a cost arising because it can reduce the credibility of national central banks.

[3] Generally, such heterogeneity works as an impediment to cooperation, so the cooperation gains identified in our analysis should be seen as an upper limit that may in practice not be reached due to political and other constraints.

bandt (2020), due to its complementary nature. Eichenbaum et al. model the macroeconomic implications of the pandemic and derive optimal policy responses from a domestic perspective. A key element in their paper is that economic activities create (negative) health externalities through interactions, by spreading the virus. Those externalities may arise in the production process, but also when products or services are consumed. The externalities provide a clear rationale for (domestic) policies to neutralize their negative effects, such as "taxing" economic activity (the shutdown of a sector can be seen as a prohibitive tax on this sector). Our analysis fully abstracts from this domestic dimension (we implicitly assume that domestic inefficiencies have been already solved through appropriate policies) but rather focus on the international aspect.

The following section sets up our model. Section 3 analyzes national containment policies. Section 4 contrasts with optimal international policies, and derives policy recommendations. Section 5 provides suggestive empirical evidence on cross-country variation on containment policies in Europe. The final section concludes.

2 Setup

We consider two identical countries, A and B (we will discuss asymmetry later) and examine two phases of the pandemic:

1. *Lockdown-phase*: There is a serious spread of the virus at the start of this period. Countries are implementing severe lockdown policies (curtailing both economic and social activities). All borders are closed.

2. *New-normal:* The virus is under control, but not fully eradicated. Countries have opened their borders again. National containment policies are still in place, but are now fairly light.

We thus do not model the initial spread of the virus, but directly enter a world where a significant number of people is infected. Arguably, our lockdown-phase could be further broken down into two parts: a strict lockdown-phase, and a period where countries are

Covid Economics 8, 22 April 2020: 134-148

starting to partially loosen their lockdown policies. As we will see later, the *direction* of the inefficiencies created in both stages are similar (their *intensity* is different though), so we analyze this as one stage. We also fully abstract from modelling explicitly the evolution of the virus spread (we do this in very simple reduced-form though), but focus here on the international ramifications of different stages of the pandemic (for a full analysis of the dynamics of a pandemic in a domestic context, see Eichenbaum, Rebelo and Trabandt 2020).

At each phase (date 1: lockdown; date 2: new normal) each country chooses its *activity-level*, x. The activity can be interpreted in an economic sense (how much to produce) or in a social sense (how much to engage in social interaction). Choosing an activity level of x brings about (net) benefits $b(x)$ (in absence of the virus). In the case of production, this can be interpreted as benefits from consuming good and services, or from exporting them. In the case of social interactions, its simply the utility derived from them. Note that these are already the *net* benefits. For example, for production this would amount to the profits (revenue minus costs). We assume that the relationship between the activity and the net benefits is concave ($b'(x) > 0, b''(x) < 0$). Let us denote with $\overline{x}$ the activity-level that maximizes net benefits (implicitly defined by $b'(\overline{x}) = 0$). We can interpret this as the country's pre-virus activity-level, i.e. the activity level that prevailed prior to the lockdown phase (an "imaginary" date 0).

When the virus is present in the economy, the action has the additional cost of spreading the virus (domestically). Specifically, let us denote the *severity* of the virus pandemic at the end of the prior period with s_{t-1} (≥ 0). We assume that the activity contributes to the spread of the virus, that is $s_t = f(x_t, s_{t-1})$ with $\frac{\partial f}{\partial x_t} > 0$ (and f bounded from above by the country's population number). In the case of production, the spread of the virus can either result in the production process itself (people working together to assemble a product), or when the good or service is consumed (as in Eichenbaum et al. 2020). We also assume that $\frac{\partial f}{\partial s_{t-1}} > 0$, that is, there is "memory" in the epidemic and higher severity in the previous period contributes to severity in the current period.[4] The prevalence of the

[4]Epidemiological models, however, suggest that this relationship could also be a negative one. In particular, by building up "herd immunity" early on, the consequences of the virus in later periods may

CEPR PRESS

virus causes costs $v(s)$ ($v'(s) > 0$) to the country. This may be because of deaths, but also due to increased costs for the health care system

A country's choice of activity level also has externalities on the other country. These externalities depend crucially on whether borders are open and whether the activity level changes in an unexpected fashion. Consider first the lockdown-phase, during which borders are closed for travel. During this phase, externalities arise predominantly on the economic side. If a country reduces its (economic) activity level below $\overline{x}$, this will lead to an unexpected disruption in the production in other countries, through supply chain linkages. Shrinkage in production will also have negative consequence through aggregate demand spillovers in a recession, and because foreigners may hold claims on domestic firms. We capture these date-1 externalities by the function $e_1(x)$, with $e_1'(x) > 0$ when $x < \overline{x}$ (that is, higher domestic activity causes positive externalities abroad).

By contrast, externalities in the new-normal will be predominantly coming through spreading the virus. At this stage, borders are open, allowing people to travel internationally.[5] Direct disruptions in the production process from curtailing production in the other country are thought to be of less relevance then. This is because the production process will have adjusted; firms will have modified their supply chain and countries will have become more autarkic. We thus take the date-2 externalities to be decreasing in the activity level of the country: $e_2'(x) < 0$. For example, a less severe lockdown in a country (higher activity level) will mean that more people will become infected (in the country), and due to travel, this will lead to more infected people abroad.[6]

be mitigated. However, most countries in the world are not (or are no longer) following this strategy.

[5] We do not consider potential coordination problems that may arise from border openings. This is because border openings are *two-sided*; a country can always protect itself from a negative externality from another country opening its borders by keepings its own border closed.

[6] Formally, containment policies in our model refer to domestic activities. However, in a reality less strict containment policies will also enable (or encourage) international travel, further contributing to the spread of the virus.

Covid Economics 8, 22 April 2020: 134-148

Covid Economics 8, 22 April 2020: 134-148

3 National Policies

We now analyze how national policies will be chosen. Specifically, we consider governments that maximize the welfare of their citizen by optimally choosing activity-levels at date 1 and date 2. A higher activity level can be interpreted as a more lenient containment policy (for example, a government shuts down less sectors or relaxes lockdown restrictions).

A country's welfare consists of the combined (domestic) surplus from both periods, which for country A is given by

$$W^A(x_1^A, x_2^A, x_1^B, x_2^B) = b(x_1^A) - v(f(x_1^A, s_0^A)) + e_1(x_1^B) + b(x_2^A) - v(f(x_2^A, s_1^A)) + e_2(x_2^B). \quad (1)$$

The government maximizes domestic welfare choosing x_1^A and x_2^A, taking as given the foreign policy choices x_1^B and x_2^B.[7]

The FOCs for date 1 and date 2 are:

$$x_1^* \;:\; b'(x_1) = v'(s_1)\frac{\partial f(x_1, s_0)}{\partial x_1} + v'(s_2)\frac{\partial f(x_2, s_1)}{\partial s_1}\frac{\partial f(x_1, s_0)}{\partial x_1}, \quad (2)$$

$$x_2^* \;:\; b'(x_2) = v'(s_2)\frac{\partial f(x_2, s_1)}{\partial x_2}. \quad (3)$$

where we have suppressed the country-index due to symmetry.

Let us first consider the date-2 choice. At this date, the government trades-off higher benefits from the activity ($b'(x_2) > 0$) with resulting costs from a higher spread of the virus in this period ($v'(s_2)\frac{\partial f(x_2,s_1)}{\partial x_2} > 0$). The trade-off at date-1 is the same, except that there is now an additional *dynamic* cost from increasing the activity, captured by the term $v'(s_2)\frac{\partial f(x_2,s_1)}{\partial s_1}\frac{\partial f(x_1,s_0)}{\partial x_1} > 0$. It arises because a higher activity at date-1 leads to an increase in the virus spread at date-1, causing the economy to enter date 2 with a higher virus severity. For a given date-2 policy, the country would thus also end up with higher date-2 virus costs (the impact on the date-2 policy for welfare can be ignored, as per the envelope theorem).[8]

[7] In particular, we assume that a government also takes future foreign policy as given, that is, it does not perceive that when it changes its date-1 policy, the date-2 policy of the other country may be affected. The motivation is that in reality we have a large number of countries, and each country on its own is too small to perceive a meaningful influence of its own actions on the containment policies of other countries.

[8] This provides a reason for lower optimal activity levels (stricter lockdown) at date 1, compared to date

Covid Economics 8, 22 April 2020: 134-148

4 International Cooperation

How do the domestic policies differ from the ones that are efficient from the international perspective? The problem of optimal policies from the world perspective can be seen as choosing $(x_1^A, x_2^A, x_1^B, x_2^B)$ to maximize the combined welfare in the countries, $W^A + W^B$.

The FOC are given by

$$\begin{aligned} x_1^W \quad &: \quad b'(x_1) + e_1'(x_1) = v'(s_1)\frac{\partial f(x_1, s_0)}{\partial x_1} + v'(s_2)\frac{\partial f(x_2, s_1)}{\partial s_1}\frac{\partial f(x_1, s_0)}{\partial x_1} \qquad (4) \\ &\quad -e_2'(x_2)\frac{\partial f(x_2, s_1)}{\partial s_1}\frac{\partial f(x_1, s_0)}{\partial x_1}, \\ x_2^W \quad &: \quad b'(x_2) = v'(s_2)\frac{\partial f(x_2, s_1)}{\partial x_2} - e_2'(x_2). \qquad (5) \end{aligned}$$

These conditions differ from the domestic ones, given by (2) and (3). Starting again from date-2, we can see that the international solution perceives higher costs of activities than the domestic government (due to $e_2'(x_2) < 0$). Given the concavity of the problem, this means that a domestic government will choose a higher activity level than what is optimal from the world perspective: $x_2^* > x_2^W$. At date-1, there are two reasons why domestic and international solutions differ. First, a domestic government ignores that a higher level of date-1 activity leads to less economic disruptions in the other country at this date ($e_1'(x_1) > 0$). Second, it also ignores that a higher activity level will mean that there is a higher virus intensity in the next period, which will lower welfare abroad due to international travel ($e_2'(x_2) < 0$). It is thus not clear whether the domestic activity benefits exceed the international ones. In fact, they may also be lower than the international ones. Where or not this is the case depends on the ratio of the externalities, with corresponding consequences for the direction of the domestic activity bias.

We can summarize

Proposition 1 *Domestically chosen activity levels generally differ from the (globally) efficient ones:*

2. A second reason is that reducing the activity-level may be more effective (in absolute terms) when the prevailing virus severity is high (this is the case if $\frac{\partial^2 f(x_1,s)}{\partial x_1 \cdot \partial s} > 0$).

(i) In the lockdown-phase, domestic activity levels are excessive ($x_1^* > x_1^W$) *when* $|e_1'(x_1^*)| < |e_2'(x_1^*)| \frac{\partial f(x_2^*, s_1^*)}{\partial s_1} \frac{\partial f(x_1^*, s_0^*)}{\partial x_1}$ *and insufficient* ($x_2^* < x_2^W$) *when*
$|e_1'(x_1^*)| > |e_2'(x_1^*)| \frac{\partial f(x_2^*, s_1^*)}{\partial s_1} \frac{\partial f(x_1^*, s_0^*)}{\partial x_1}$;

(ii) In the new-normal, domestic activity levels are excessive ($x_2^* > x_2^W$).

The proposition is derived for a symmetric setup that only allows common variation in externalities among countries (e.g., we can consider sets of countries with either both high or both low date-1 externalities). However, it is easy to see that the insights also carry over to assymmetric settings. In particular, when we have $\left|e_1^{A\prime}(x_1^*)\right|$ is sufficiently low (relative to $\left|e_2^{A\prime}(x_1^*)\right|$) for country A but sufficiently high for country B, date-1 activity levels will be excessive in A, but insufficient in B.

Proposition 1 suggests that in the new-normal, there is scope for policy coordination among countries, with the objective of avoiding that countries end up with too lax policies. This could take the form of minimum containment standards across countries. For example, countries may decide to discourage larger gatherings, such as sport events of festivals that exceed a threshold number of participants. Alternatively, this may take the form of countries making wearing masks for certain infection-prone activities (such as services like hairdressing) compulsory. Notably, give that the externalities in the new-normal phase (arising from travel) are essentially worldwide, one would need a global approach for this, for example orchestrated by the WHO. In the absence of global cooperation, regional cooperation, such as within the EU is called for. However, one may also expect to see individual approaches. For example, a country that sees a lot of its residents travelling to another country (tourism!) may put pressure on the other country to adhere to strict policies in order to avoid its citizen being infected while abroad (and bringing the virus back home).

In the lockdown-phase, policies can either be too lax or too lenient. We would expect them to be too strict for countries that display a high degree of economic integration, as in this case the date-1 externality, running among others through the supply chain, will be dominating. By contrast, for countries that have high "people integration" (that is, countries with a lot of mobility), the second externality may dominate and we may expect

Covid Economics 8, 22 April 2020: 134-148

CEPR PRESS

Covid Economics 8, 22 April 2020: 134-148

such countries to be too lax in their policies (resulting in too high activity levels). We would thus anticipate excessive activity levels in fairly closed economies that have a high amount of travelling. Examples would be here typically tourist destinations, such as Thailand, Turkey or Greece. There is consequently international interest to curtail activities in such countries, in order to avoid a new spread of the virus.[9] By contrast, in countries with high economic integration relative to their international mobility, such as Germany or China (supply chain!), there may be international interest in relaxing their domestic restrictions, resulting in higher domestic activities that are less disruptive to global supply chains.

We may also expect the relative importance of the two externalities to vary during the lockdown-phase. In early stages of the lockdown phase, the economic externalities are expected to be severe, as the arrival of the lockdown comes as a full surprise. However, as the lockdown progresses, domestic economies will adjust. This suggests that as the lockdown phase progresses there will be a tendency for domestic activities levels to move from (possible) initial excessive strictness to (excess) leniency. This implies that the focus of international cooperation should change correspondingly during the lockdown policy.

5 Some suggestive empirical evidence

To provide suggestive evidence for our theoretical findings, we relate containment policies across 27 European countries to death rates and their reliance on merchandise trade and tourism. Specifically, we use data on the geographic distribution of COVID-related deaths from the European Centre for Disease Prevention and Control to identify the date when in each of the 27 countries the COVID-19 death toll has reached or passed 10 deaths.[10] We use data from a database put together by Olivier Lejeune[11] on containment policies across the globe to identify the date when non-essential shops, restaurants and bars closed as this is a measure that most European countries have taken now (some countries have gone even

[9] As the countries in question would see their economies contracting as a consequence, this may require subsidies from other countries. These subsidies could come from countries that loosen their excessively strict policies, and hence benefit economically.

[10] We include EU countries except Latvia and Sweden, as well as the UK and Switzerland.

[11] https://github.com/OlivierLej/Coronavirus_CounterMeasures

Covid Economics 8, 22 April 2020: 134-148

further). We calculate the number of days between the date of 10 or more COVID-related deaths and this containment measure, which ranges between -20 (Bulgaria) to 12 (Italy), with lower numbers indicating a quicker adoption of containment policies. As countries were hit by the virus at different points in time, thus allowing for learning effects, we first regress this difference on the date of 10 or more COVID-related deaths and find a significant and negative relationship, i.e., countries that experienced the outbreak later imposed containment policies more quickly.[12] We then compare the residuals from this crisis across four groups of countries, countries with above and below median merchandise trade to GDP and countries with above and below international tourism receipts as share of exports (data from the World Development Indicators and for 2018).

Figure 1 provides suggestive evidence consistent with our model. Countries with above median merchandise exports were quicker in imposing containment policies, not taking into account externalities on other countries through supply chains or demand externalities. Countries with above median reliance on tourism, on the other hand, were slower in imposing containment policies. When considering the two groups of countries in the two extremes of our two variables high (low) merchandise trade and low (high) reliance on tourism the difference is even stronger.

Figure 1: Variation in containment policies

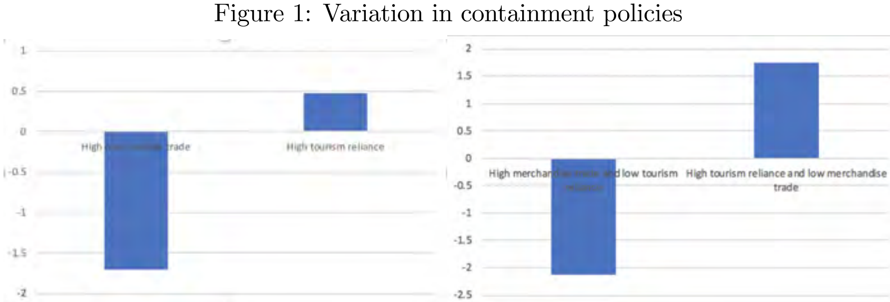

[12] For countries that have not reached ten deaths yet, we set the date at 100 days after Italy and the difference between this date and the adoption of containment policy at -50. Two countries - Latvia and Sweden, which had not adopted robust containment policies as of 6 April - are dropped from the analysis.

6 Conclusions

This paper has analyzed the question of whether national containment policies lead to international inefficiencies. The answer is yes, even in a fully symmetric world. The direction of the efficiency is shown to depend both on the stage of the epidemic, as well as country characteristics, such as economic integration and mobility. Based on this we have derived policy recommendations for countries can eliminate (or at least reduce) the inefficiencies. Importantly, given the global nature of the virus pandemic (both in terms of the pandemic itself, but also in terms of spillovers from containment policies), measures that aim to be effective have to be taken at the truly global level (for example, instigated by the WHO).

Covid Economics 8, 22 April 2020: 134-148

Covid Economics 8, 22 April 2020: 134-148

References

[1] Baldwin, Richard and Beatrice Weder di Mauro (editors) (2020), Economics in the Time of Covid-19, CEPR ebook.

[2] Beck, Thorsten, Silva Buston, Consuelo and Wolf Wagner (2018), The Economics of Supranational Bank Supervision, CEPR Discussion Paper No. 12764.

[3] Cooper, Richard (1969), Macroeconomic Policy Adjustment in Interdependent Economies. Quarterly Journal of Economics 83(1), pp. 1–24.

[4] Cooper, Richard (2001), International Cooperation in Public Health as a Prologue to Macroeconomic Cooperation, In Can Nations Agree? Issues in International Economic Cooperation, eds. Richard N. Cooper, Barry Eichengreen, Gerald Holtham, Robert D. Putnam, and C. Randall Henning. Washington, DC: Brookings Institution.

[5] Dell'Ariccia, Giovanni and Robert Marquez (2006), Competition Among Regulators and Credit Market Integration, Journal of Financial Economics 79, 401-30.

[6] Eichenbaum, Martin, Rebelo, Sergio and Mathias Trabandt (2020), The Macroeconomics of Epidemics, NBER Working Paper No. 26882

[7] Frankel, Jeffrey (2015), International Coordination, Policy Challenges in a Diverging Global Economy , Asia Economic Policy Conference, organized by Reuven Glick and Mark Spiegel (Federal Reserve Bank of San Francisco), Nov. 19-20.

[8] Hamada, Koichi. 1976. "A Strategic Analysis of Monetary Interdependence." Journal of Political Economy 84(4), pp. 677–700

[9] Maloney, John, and Malcolm Macmillen (1999). 'Do Currency Areas Grow Too Large for their Own Good?' Economic Journal, 109, 572–587.

[10] Mundell, Robert (1961). Theory of Optimum Currency Areas. American Economic Review, 51, pp. 657-665

[11] Oates, Wallace (1972), Fiscal Federalism, Harcourt-Brace, New York.

[12] Rogoff, Kenneth, 1985. "Can international monetary policy cooperation be counter-productive?," Journal of International Economics, Elsevier, vol. 18(3-4), pages 199-217, May.

Covid Economics 8, 22 April 2020: 134-148