CENTRE FOR
ECONOMIC
POLICY
RESEARCH

CEPR PRESS

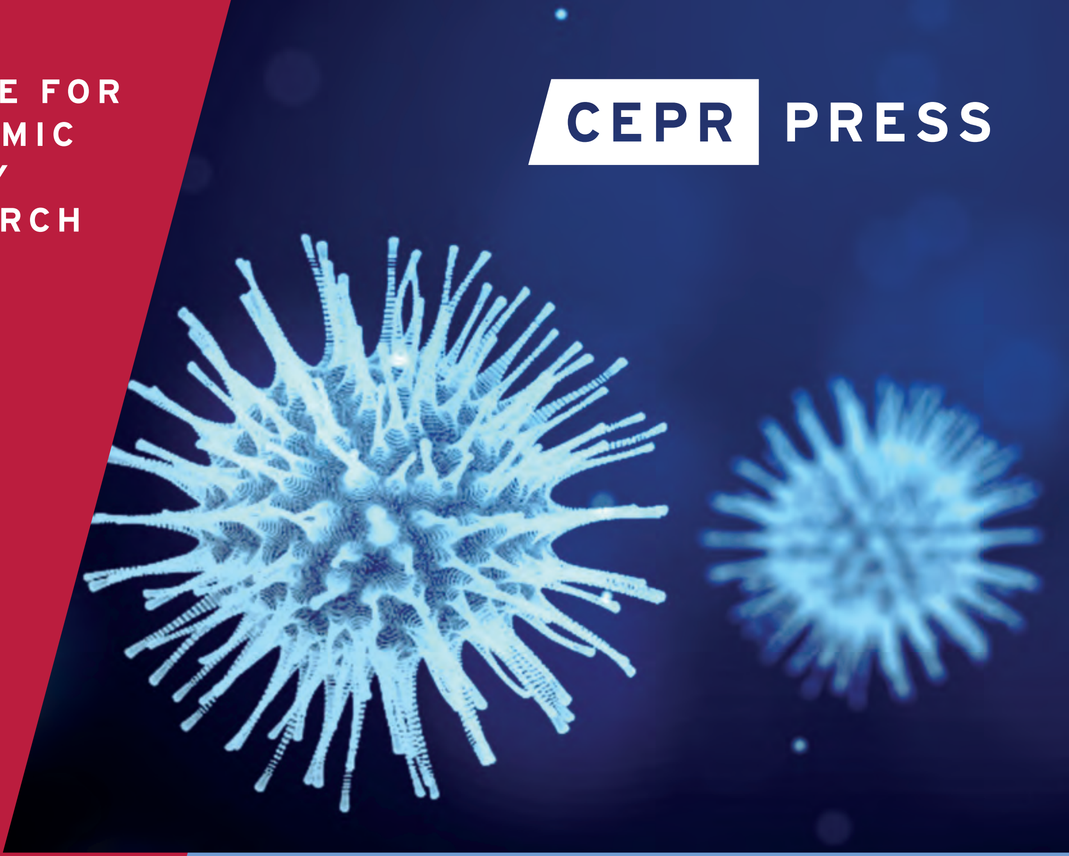

COVID ECONOMICS

VETTED AND REAL-TIME PAPERS

ISSUE 21
22 MAY 2020

DISASTER RESILIENCE AND ASSET PRICES
Marco Pagano, Christian Wagner and Josef Zechner

CATASTROPHIC DECLINES IN EMPLOYMENT IN THE US
Olivier Coibion, Yuriy Gorodnichenko and Michael Weber

TRADE OF ESSENTIAL GOODS
Fernando Leibovici and Ana Maria Santacreu

SOCIAL NETWORKS AND CONSUMPTION
Christos A. Makridis and Tao Wang

BANKS AS LENDERS OF FIRST RESORT
Lei Li, Philip E. Strahan and Song Zhang

OPTIMAL SOCIAL DISTANCING
Thomas Kruse and Philipp Strack

Covid Economics
Vetted and Real-Time Papers

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

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

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

The Centre for Economic Policy Research (CEPR)

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

Editorial Board

Ethics

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

Submission to professional journals

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

American Economic Review

American Economic Review, Applied Economics

American Economic Review, Insights

American Economic Review, Economic Policy

American Economic Review, Macroeconomics

American Economic Review, Microeconomics

American Journal of Health Economics

Economic Journal

Economics of Disasters and Climate Change

International Economic Review

Journal of Development Economics

*Journal of Econometrics**

Journal of Economic Growth

Journal of Economic Theory

*Journal of the European Economic Association**

Journal of Finance

Journal of Financial Economics

Journal of International Economics

*Journal of Labor Economics**

Journal of Monetary Economics

Journal of Public Economics

Journal of Political Economy

Journal of Population Economics

*Quarterly Journal of Economics**

Review of Economics and Statistics

*Review of Economic Studies**

Review of Financial Studies

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

Covid Economics
Vetted and Real-Time Papers

Issue 21, 22 May 2020

Contents

CEPR PRESS

1

Covid Economics Issue 21, 22 May 2020

Disaster resilience and asset prices[1]

Marco Pagano,[2] Christian Wagner[3] and Josef Zechner[4]

Date submitted: 17 May 2020; **Date accepted:** 20 May 2020

This paper investigates whether security markets price the effect of social distancing on firms' operations. We document that firms that are more resilient to social distancing significantly outperformed those with lower resilience during the COVID-19 outbreak, even after controlling for the standard risk factors. Similar cross-sectional return differentials already emerged before the COVID-19 crisis: the 2014-19 cumulative return differential between more and less resilient firms is of similar size as during the outbreak, suggesting growing awareness of pandemic risk well in advance of its materialization. Finally, we use stock option prices to infer the market's return expectations after the onset of the pandemic: even at a two-year horizon, stocks of more pandemic-resilient firms are expected to yield significantly lower returns than less resilient ones, reflecting their lower exposure to disaster risk. Hence, going forward, markets appear to price exposure to a new risk factor, namely, pandemic risk.

Covid Economics 21, 22 May 2020: 1-39

1 Marco Pagano gratefully acknowledges financial support from the Italian Ministry for Education, University and Research (MIUR) and the Einaudi Institute for Economics and Finance (EIEF). Christian Wagner acknowledges support from the Center for Financial Frictions (FRIC) grant no. DNRF102.

2 Professor of Finance, University of Naples Federico II; Director, Centre for Studies in Economics and Finance (CSEF); Fellow, Einaudi Institute for Economics and Finance (EIEF); CEPR Research Fellow.

3 Professor of Finance, WU Vienna University of Economics and Business; faculty member, Vienna Graduate School of Finance (VGSF).

4 Professor of Finance and Investments, WU Vienna University of Economics and Business; faculty member, Vienna Graduate School of Finance (VGSF); CEPR Research Fellow.

Covid Economics 21, 22 May 2020: 1-39

1 Introduction

The COVID-19 pandemic and the resulting lockdown are currently inflicting massive harm on the economy, in the form of unprecedented output losses, massive redundancies and countless bankruptcies. The effects of the pandemic are widely perceived not to be purely transient, as witnessed by current changes in expectations (Coibion et al., 2020b; Hanspal et al., 2020; Coibion et al., 2020a) and in asset prices (Gormsen and Koijen, 2020). But its effects are also proving to be quite heterogeneous: some firms, especially in high-tech industries, appear to have adapted quite well to social distancing requirements, for instance by resorting extensively to teleworking, while others, such as food catering, travel and tourism, could not do so, as the nature of their business requires close contact with customers and/or among employees. This heterogeneity is also visible in the diverging stock price performance of companies: in the first quarter of 2020, stocks such as Apple, Microsoft and Google outperformed the market, yielding market-adjusted returns of 19%, 12% and 33% respectively, while others such as Marriott, United Airlines and Royal Caribbean massively underperformed, with market-adjusted returns of −38%, −53% and −66%.

Hence, the COVID-19 shock has unearthed a hitherto hidden economic watershed, namely, that between disaster-resilient activities and non-resilient ones. Insofar as the COVID-19 pandemic persists or revives in the near future, or similar disasters may occur further in the future, resilience may become a key firm attribute, one which will be relevant to investors' portfolio choices, banks' lending policies, and managers' investment decisions.

In this paper, we investigate whether asset prices reveal growing investors' awareness that pandemic resilience, defined as reliance on technologies and/or organizational structures that are robust to social distancing, is priced by security markets. To measure firms' resilience to pandemics, we rely on measures recently introduced in labor economics by Dingel and Neiman (2020), Hensvik et al. (2020) and Koren and Pető (2020), intended to capture the extent to which firms' operations are compatible with social distancing. These measures quantify the degree to which jobs can be done from home and do not rely on human interaction in physical proximity.

We then test whether the stocks of more resilient companies have generated excess returns, after controlling for market risk and other established risk factors, *before* and/or *during* the COVID-19 shock. We find this to be the case: pandemic-resilient stocks outperformed less resilient ones not only between late February and March

2020, i.e. during the outbreak of the pandemic, but also in the previous six years, in which their cumulative excess return was of similar magnitude as during the crisis.

Furthermore, we investigate the returns that investors expect the two stock classes to generate *after* the COVID-19 shock, by extracting the expected stock returns implied by option prices. We find that, going forward, pandemic-resilient stocks are expected to generate lower excess returns, indicating that, since the COVID-19 pandemic materialized, the market has priced a disaster premium.

Thus, our research question is not only *whether*, but also *when* financial markets started pricing disaster risk and resilience to it. Interestingly, our evidence indicates that investors became gradually aware of such risk even before the COVID-19 shock, and they consider such risk to be still price-relevant, even after it has materialized.

Given the largely unanticipated nature of the current pandemic, it may appear unrealistic that investors became cognizant to its threat in advance. But it should be recalled that as early as five years before the COVID-19 outbreak, high-profile business and political leaders already issued public warnings of the risk of devastating epidemics. For instance, in a speech delivered on 2 December 2014 at the NIH about the response to the Ebola epidemic, U.S. President Barack Obama stated:

> “There may and likely will come a time in which we have both an airborne disease that is deadly. And in order for us to deal with that effectively, we have to put in place an infrastructure – not just here at home, but globally – that allows us to see it quickly, isolate it quickly, respond to it quickly. [...] So that if and when a new strain of flu, like the Spanish flu, crops up five years from now or a decade from now, we've made the investment and we're further along to be able to catch it. It is a smart investment for us to make. It's not just insurance; it is knowing that down the road we're going to continue to have problems like this – particularly in a globalized world where you move from one side of the world to the other in a day.”

Similarly, in 2015, Microsoft co-founder Bill Gates gave a TED Talk about pandemics that attracted widespread attention. In this talk he warned that in 2014 the world had barely avoided a global outbreak of Ebola, largely because of pure luck, and, just like Obama, alerted the audience to the need to prepare for future pandemics, from scenario planning to vaccine research and health worker training. Hence, it cannot be ruled out that the most alert investors may have started taking

Covid Economics 21, 22 May 2020: 1-39

Covid Economics 21, 22 May 2020: 1-39

into account such concerns in their portfolio choices well in advance of the current pandemic, shying away from the stocks of companies that would be less resilient to it and starting to overweight those likely to be more resilient.

One may also wonder why even after the occurrence of COVID-19 the market still prices pandemic risk to some extent, being willing to accept lower expected returns on more resilient stocks, as revealed by option prices according to our estimates. However, currently there is still high uncertainty regarding the duration of the COVID-19 pandemic. Medical experts have repeatedly warned about the risk of a second wave of contagion as the lockdown measures are gradually relaxed (see Xu and Li, 2020, among others). Indeed, top health officials do not rule out that the disease may become endemic. On 13 May 2020, Dr. Mike Ryan, executive director of emergencies at the World Health Organization (WHO) stated:

> "I think it's important to put this on the table. This virus may become just another endemic virus in our communities. And this virus may never go away. HIV has not gone away, we've come to terms with the virus [...] I think it is important that we're realistic and I don't think anyone can predict when or if this disease will disappear."

Such uncertainty, possibly coupled with heightened awareness that similar disasters may occur again in the future, could explain why, since COVID-19, pandemic risk is priced in the cross-section of returns, as shown by our evidence based on option prices. Hence, going forward, asset pricing models will have to include exposures to this additional risk factor among those used to explain the cross-section of returns, and asset managers will have to take such exposures into account in portfolio selection.

Our analysis is related to the asset pricing literature on rare disasters, starting with Rietz (1988), who extends the Lucas (1978) model to include a rare disaster state and shows that this leads to high equity risk premia and low risk-free returns, even with reasonable time discounting and risk preferences. Barro (2006) and Barro (2009) extend this model and show that empirically calibrated disaster probabilities may suffice to explain the observed high equity premium, low risk-free rate and stock return volatility. The theoretical literature on disaster risk has also been extended to allow for learning (see, e.g.,Veronesi (2004), Wachter and Zhu (2019), Gillman et al. (2014), Lu and Siemer (2016)), and/or for stochastic disaster risk (see, e.g., Wachter (2013) and Gabaix (2012)).

CEPR PRESS

One common feature of these models is that risk premia that would appear abnormally high conditioning on no disaster occurring, are in fact justified, being merely an equilibrium compensation for the expected loss in a disaster plus a risk premium as this loss occurs when the marginal utility of consumption is high. In particular, the model by Gabaix (2012) implies that, in the cross-section of stocks, those that are expected to be less resilient to disasters, should carry a higher risk premium, conditioning on no disaster occurring, but also unconditionally. Our empirical analysis speaks to these predictions.

Our work is also related to a fast growing literature on the stock market response to the COVID-19 pandemic in the first quarter of 2020. The early comprehensive study by Ramelli and Wagner (2020) documents that during the 'outbreak' period (which they define as 20 January to 21 February 2020), U.S. firms with high exposure to China and, more generally, to international trade, as well as firms with high leverage and low cash holdings experienced the sharpest stock price declines. The leverage- and cash holding effects also persist through the 'fever' period (from 24 February to 20 March 2020). Bretscher et al. (2020) provide evidence for supply chain effects in the cross-section of stocks during COVID-19. Moreover, Albuquerque et al. (2020) find that firms with high environmental and social (ES) ratings offered comparably higher returns and lower return volatility in the first quarter of 2020.

Some studies relate the price response of different stocks to the pandemic to the corresponding firms' exposure to the disease. Ramelli and Wagner (2020) and Hassan et al. (2020) analyze conference call data, which the latter use to construct text-based, firm-level measures for exposures to epidemic diseases, and find that stock returns are significantly and negatively related to disease exposures, with demand- and supply-chain related concerns being primary drivers. Alfaro et al. (2020) analyse the effect of unanticipated changes in infections during the SARS and the COVID-19 epidemics on stock returns, and show that stock prices drop in response to high unexpected infections. In the cross-section, exposure to pandemic risk turns out to be greater for larger and more capital intensive firms, and, consistently with Ramelli and Wagner (2020), more levered and less profitable ones.

Some of the results obtained for the response of U.S. stock returns to the pandemic appear to extend to non-U.S. stock returns. Ding et al. (2020) show for a sample of over 50 countries that firms with better financials, less supply chain exposures and more corporate social responsibility (CSR) activities experienced milder stock price reactions in the first quarter of 2020. Other studies focus on the response

Covid Economics 21, 22 May 2020: 1-39

COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

of country-level stock market indices to COVID-19: Ru et al. (2020) find that stock markets in countries with 2003 SARS experience reacted more quickly to the outbreak than countries without prior experience, while Gerding et al. (2020) document that market declines were more severe in countries with lower fiscal capacity, defined as higher debt/GDP ratio. Finally, Croce et al. (2020) use Twitter news to study the (real-time) COVID-19 caused contagion in global equity markets.

Our analysis differs from that in all these papers since it focuses on the asset pricing implications of companies' exposure to social distancing, which is the main economic consequence of the epidemic, and studies such implications not only for the period of the COVID-19 outbreak, but also prior and after its occurrence. This enables us to investigate whether learning about pandemic risk occurred in advance of the outbreak, and whether investors kept pricing it since the outbreak.

The rest of the paper is structured as follows. Section 2 provides a framework to interpret the relationship between disaster risk and the stock returns of firms featuring different disaster resilience. Section 3 provides a brief discussion of alternative measures of firms' disaster resilience. The data and results are presented in Section 4. Section 5 concludes.

2 Disaster Awareness and Risk Premia

In this section we sketch three distinct and mutually exclusive models of how financial markets may respond to disaster risk and to its materialization. As we shall see, each model has different predictions about the stock return differential of resilient *vs.* non-resilient firms prior, during and after the occurrence of a disaster, as shown in Table 1. Our empirical analysis in Section 4 will investigate which of the three sets of predictions is most consistent with the data.

Table 1: **Predicted return differential of resilient vs. non-resilient firms**

Theory	Before COVID-19	During COVID-19	After COVID-19
Unpriced disaster risk	Zero	Positive	Zero
Priced disaster risk	Negative	Positive	Negative
Pre-disaster learning	Positive	Positive	Negative

Unpriced disaster risk. We start with a model where disaster is completely unex-

Covid Economics 21, 22 May 2020: 1-39

CEPR PRESS

pected. In the case of COVID-19, this amounts to assuming that before the first quarter of 2020 financial market participants were unaware of the danger posed by the virus and of its consequences in terms of social distancing. Going forward, a new disaster is again regarded as a zero-probability event, or anyway as being price irrelevant (for instance, in the case of COVID-19, due to development of vaccines and/or effective drugs). In this model, disaster risk would not be priced before COVID-19 nor after it. Such market expectations could result either from bounded rationality, i.e. investors assigning a zero probability weight to a positive probability event, or from a disaster truly having a negligible probability both before and after its occurrence. In the latter case, unpriced disaster risk would be consistent with rational expectations.

If this is the correct model in the case of the COVID-19 pandemic, one should observe (i) no return pattern related to firm pandemic resilience prior to 2020; (ii) a sharper price drop for less resilient firms than for more resilient ones at the time of the outbreak, as investors take into account that the cash flows of the former will be hurt more than those of the latter; (iii) no differential response of the expected returns of the two classes of firms after the crisis, since COVID-19 does not lead to any updating of the return-generating process, i.e. disaster risk remains unpriced. As disaster is considered as a one-time event, its occurrence leaves the stochastic discount factors of all stocks unchanged.

Priced disaster risk. The second model is one where disasters, however rare, are rationally anticipated, so that more resilient firms are priced at a premium relative to less resilient ones, i.e. offer a lower expected return in no-disaster periods, as predicted by Barro (2006), Barro (2009) and Gabaix (2012). Such models predict that, in equilibrium, securities more exposed to a disaster, i.e. those issued by less resilient firms, pay a risk premium to compensate investors for this risk. Hence, they predict that (i) prior to the disaster, stocks' excess returns are related to firms' disaster resilience; (ii) during the disaster, investors take into account that the cash flows of less resilient firms drop by more than those of the more resilient ones, so that their disaster-time stock performance is worse; (iii) after the disaster, the pre-disaster excess return pattern re-emerges, i.e. disaster risk remains priced. In principle, this hypothesis does not preclude that investors may update the probability of disasters upon the occurrence of one: if they have increased this probability as a result of COVID-19, the expected return differential between non-resilient and resilient firms should become more pronounced after the pandemics than it was before it.

Covid Economics 21, 22 May 2020: 1-39

Prior learning about disaster risk. Finally, we consider a model where investors learn about the probability of a disaster occurring or about its implications before its occurrence. In the case of COVID-19, as mentioned in the introduction, they may have revised upwards the probability of a pandemic occurring or become more keenly aware of its social distancing implications, for instance as a result of SARS, Ebola, and/or the statements by Bill Gates, Barack Obama and other opinion leaders. By the same token, investors may have become more aware of the characteristics that make firms more resilient to pandemics. Any of these forms of learning implies a demand shift by investors from less to more resilient stocks before the pandemic, leading the latter to outperform the former, once their respective exposures to "traditional" risk factors are controlled for. Once disaster strikes, one would again observe the stocks of more resilient firms outperforming less resilient ones. But this pattern should reverse in the post-disaster phase: as at that stage learning would be complete, the stocks of more resilient firms will be priced at a premium relative to less resilient ones, i.e. should offer a lower expected return. This scenario has parallels with learning about the "importance" of Environmental, Social and Governance (ESG) scores by investors. In a rational expectations equilibrium without learning, portfolios with strict ESG rules should underperform, but in the presence of gradual learning about the importance of ESG ratings for investors, the stocks of firms with high ESG ratings appreciate, and ESG mutual funds outperform (see Pastor et al., 2020).

Hence, this model predicts that (i) excess returns of resilient relative to non-resilient firms should be observed prior to the disaster; (ii) when disaster strikes, investors take into account that the cash flows of less resilient firms drop by more than those of more resilient ones; (iii) after the disaster, resilient firms offer lower expected returns that non-resilient ones, as learning about disaster risk has taken place.

3 Measuring Firm Resilience to Pandemics

This section describes social distancing measures that may be relevant for the pricing of stocks, as firms with operations requiring less direct physical interaction and more easily performed from home should be more resilient to social distancing rules than other firms. To gauge the effect of social distancing on firms, recent research in labor economics has developed measures of the extent to which jobs can be done from home and rely on human interaction in physical proximity. Some studies have developed

Covid Economics 21, 22 May 2020: 1-39

Covid Economics 21, 22 May 2020: 1-39

such measures starting from worker-level survey responses, while others have done so by characterizing the tasks required by each occupation based on the Occupational Information Network (O*Net) and on the authors' own judgement.

Hensvik et al. (2020, HLR) use the first approach: they rely on the 'American Time Use Survey' (2011-2018) to estimate the prevalence of working at 'home' and at the 'workplace' and, starting from worker-level survey responses, estimate the fraction of employees that work at home and at the workplace, as well as the hours worked at home and at the workplace at the industry-level.[1] Instead, both Dingel and Neiman (2020, DN) and Koren and Pető (2020, KP) use data from O*Net surveys. DN use this data, and their own judgement, to assess the teleworkability of occupations and provide industry-level estimates for the percentage of jobs that can be done at home as well as for the percentage of wages associated with teleworkable occupations. KP construct three types of industry-level measures of face-to-face interactions, depending on whether these are due to internal communication ('teamwork'), external communication ('customers'), or physical proximity to others ('presence'). They also aggregate 'teamwork' and 'customers' to a measure of 'communication' intensity and construct an industry-level measure of the percentage of employees 'affected' by social distancing regulations due to their occupations being communication-intensive and/or requiring close physical proximity to others.[2]

In our estimates, we rely primarily on the measures proposed by KP, but also check whether our results are robust to the use of those produced by HLR and DN. The results presented in the next section focus mostly on KP's 'affected_share' variable. We choose this as our main variable because, beside teleworkability, it also explicitly accounts for physical proximity to others, i.e. exactly what social distancing rules aim to avoid. Additionally, we consider DN's measure of the fraction of wages accounted for by jobs that can be performed at home ('teleworkable_manual_wage') and HLR's measure of daily work hours at the workplace ('dur_workplace'). However, we also discuss results obtained using all other variables suggested by KP, DN, and HLR. Table A.1 in the Appendix lists all the measures and presents their definitions.

[1] As the authors mention, their measures for working at home should provide a lower bound for the current situation, as the COVID-19 crisis is likely to have prompted additional substitution of work at the workplace with work at home.

[2] The estimates of these studies are available for NAICS industry classifications, at the 2- and 3-digit level for DN, at the 4-digit level for HLR, and at the 3-digit-level for KP. For details on the data, see Table A.1 in the Appendix.

CEPR PRESS

4 Empirical Results

We use the NAICS industry classification of stocks to assign the DN, HLR, and KP metrics to firms, and code each firm as a 'High' or 'Low' resilience one, depending on how its industry scores relative to the median value of the relevant metric. Then, we analyze whether and how the resulting variation in U.S. firms' pandemic resilience affects the cross-section of their stock returns at the time of the COVID-19 shock (Subsection 4.2), before the shock (Subsection 4.3) and after its occurrence (Subsection 4.4).

4.1 Data and Methodology

We obtain prices for all common stocks listed at the NYSE, AMEX and NASDAQ from the Compustat Capital IQ North America Daily database and compute daily returns, accounting for price-adjustments and dividends. We also retrieve data on daily risk-free, market and standard factor returns from Kenneth French's website. Our estimates in Subsection 4.2 require daily data from January 2019 to March 2020: after estimating firms' exposures to common factors from 2019 data, we use these exposures to compute factor model-adjusted stock returns in the first quarter of 2020.[3]

For the regressions of daily stock returns on factor returns, we require a minimum of 127 daily observations, and estimate the following specifications: we alternatively regress stock returns on market excess returns (CAPM), on the returns of the three Fama and French (1993) factors (FF3, i.e. market, size, and value) and the five Fama and French (2015) factors (FF5, i.e. market, size, value, investment, and profitability). In addition, we augment the FF3 and FF5 specifications by the momentum factor (Carhart, 1997) so as to obtain FF4 and FF6 exposures. The 2019 exposures are then used to measure factor model-adjusted stock returns in the first quarter of 2020 as the difference between a stock's daily excess return and its CAPM beta multiplied by the daily market excess return; we proceed analogously for the FF specifications.

Next, stock return data are matched with the resilience proxies based on KP, DN, and HLR metrics by industry, based on firms' 2-, 3-, or 4-digit NAICS codes. Only industries for which resilience measures are available are retained in the data set.

[3]This approach follows Ramelli and Wagner (2020) and other related papers.

Covid Economics 21, 22 May 2020: 1-39

CEPR PRESS

Covid Economics 21, 22 May 2020: 1-39

Moreover, firms with equity market capitalization below USD 10 million are dropped from the sample. For the first quarter of 2020, this results in a sample with a total of 227,812 observations for 3,614 firms in 75 industries at the NAICS 3-digit level for KP and DN, and 222 industries at the NAICS 4-digit level for HLR.

For the analysis in Subsection 4.4, we additionally obtain S&P 500 index option and individual stock options data from OptionMetrics for the first quarter of 2020. We use the volatility surface data to compute SVIX-measures of the risk-neutral variance, keeping data for firms for which we can compute SVIX-measures for horizons of 30, 91, 182, 365, and 730 days, and match this with the stock data. This results in a sample of 160,951 observations for 2,721 firms in 74 industries at the NAICS 3-digit level for KP and DN and 212 industries at the NAICS 4-digit level for HLR.

4.2 Returns and Firm Resilience during the Disaster

We study stock returns in the first quarter of 2020, and specifically from February 24, the day after Italy introduced its lockdown, to March 20, the last trading day before the Fed announced aggressive action intended to soften the blow of the pandemic.[4] As shown by Panel A of Figure 1, in this time window there was a surge in the public's attention to the COVID-19 epidemic (as measured by Google trends), while the prices of U.S. stocks (as measured by the Fama-French market factor) fell sharply.

A first look at the data suggests that the stocks of more pandemic-resilient firms, i.e. those included in the 'High' resilience portfolio (based on the 'affected_share' metric by KP) performed better in this time window than those in the 'Low' resilience portfolio. Panel B in Figure 1, which plots cumulative excess returns for the value-weighted portfolios of firms with high and low resilience, shows that both portfolios dropped sharply in price, but that of high-resilience firms depreciated far less: from February 24 to March 20, their shares outperformed those of the other group by approximately 10%.

Since the different performances of the two portfolios shown in Figure 1 may stem

[4]This period corresponds to the 'fever'-period in Ramelli and Wagner (2020); see their paper for a detailed account of the sequence of events. On Monday March 23, the Fed unveiled its plan to buy an unlimited amount of bonds with government guarantees, including some commercial mortgage debt. It also established the Secondary Market Corporate Credit Facility (SMCCF), in order to purchase existing investment-grade corporate debt, including exchange-traded funds, as well as the Primary Market Corporate Credit Facility (PMCCF), to purchase newly issued corporate bonds, so as to prevent companies facing pandemic fallout from dismissing employees and terminating business relationships.

Figure 1. Attention to Covid19 and US equity returns

Panel A illustrates the attention to Covid19 in the United States, as measured by the Google trends index for the term "Coronavirus" in the US, and the cumulative returns of the US stock market, as measured by the Fama-French market factor, during the first quarter of 2020. Panel B plots the cumulative returns of portfolios sorted by firms' resilience to disaster risk. On any given day, we assign a firm to the 'High' portfolio if its 'affected_share' (as defined by Koren and Pető, 2020) is below the median value and to the 'Low' portfolio if it is above. We plot the cumulative value-weighted portfolio returns for the 'High' portfolio (in green) and the Low portfolio (in red) as well as the High-Low differential return (in blue). The dashed vertical lines mark February 24, the day after Italy introduced its lockdown, and March 20, the last trading day before the Fed announced its intervention.

Panel A. Google trends index and aggregate US stock market

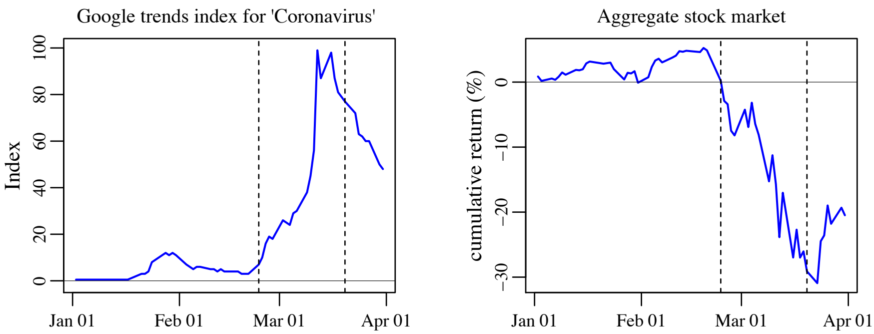

Panel B. Excess returns of firms with high and low resilience to social distancing

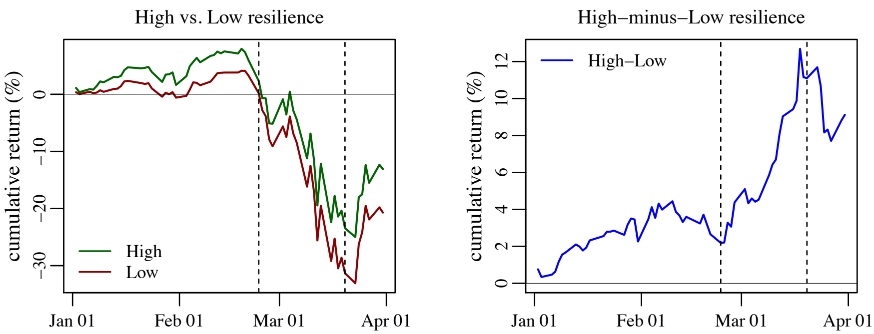

Covid Economics 21, 22 May 2020: 1-39

from their different exposure to standard risk factors, we estimate CAPM- and FF-exposures from daily data in 2019. We then use these exposures to compute daily CAPM-adjusted and FF-adjusted returns, i.e., excess returns net of such exposures multiplied by the respective factor returns (see Subsection 4.1 for details). Figure 2 presents the results: controlling for market factor risk (CAPM-adjusted returns in Panel A), the cumulative return of the High- and Low-resilience portfolio are about +10% and −15%, respectively, from February 24 to March 20, i.e. the cumulative CAPM-adjusted High-minus-Low return is approximately 25%. Panels B and C show that accounting for the FF-factors does not change the results qualitatively: the resulting risk-adjusted return differentials are in the range between 15% and 20%. The plots suggest that the differential return between the two portfolios has been negligible from early January until late February, and that the results in the COVID-19 time window are mostly driven by the sharp decline of the Low-resilience portfolio. Once the Fed announced its intervention (on March 23) the High-resilience portfolio recovered slightly and, as a result, the Low-minus-High differential dropped. Interestingly, the time-series of the cumulative Low-minus-High returns resembles that of the Google trends index. Indeed, regressing daily returns on changes in the Google index confirms a statistically significant relation, with R^2 between 0.19 and 0.22, depending on which factor-model adjustment is used.

Figures A.1 and A.2 in the Appendix provide evidence on the robustness of the results shown in Figure 2 with respect to other resilience measures. They plot the risk-adjusted returns and the differentials of the High- and Low-resilience portfolios based on DN's 'teleworkable_manual_wage' and HLR's 'dur_workplace'. The results are qualitatively similar to those in Figure 2.

To analyze the statistical significance of these findings, Table 2 reports the averages of daily excess returns (in Column 1) and risk-adjusted returns (in Columns 2 to 6) for High-resilience and Low-resilience stocks from February 24 to March 20, as well as the difference between the two, using the resilience measures based on KP, DN and HLR. Panel A shows that the differential return based on SKP's 'affected_share' metric is statistically significant, whether it is based on raw excess returns or adjusted for market exposure (CAPM), for the classic Fama-French factors (FF3 and FF5) or for those that also control for momentum (FF4 and FF6). For resilience measures based on DN (Panel B) and HLR (Panel C), the average CAPM-adjusted differential returns are statistically significant as well, but there is some variation in the significance of the FF-adjusted returns. A common feature across all resilience

Figure 2. Risk-adjusted returns of stocks with high and low resilience to social distancing

This figure plots the cumulative risk-adjusted returns of portfolios sorted by firms' resilience to disaster risk during the first quarter of 2020. On any given day, we assign a firm to the 'High' portfolio if its 'affected_share' (as defined by Koren and Pető, 2020) is below the median value and to the 'Low' portfolio if it is above. In Panel A, we present CAPM-adjusted returns, i.e. controlling for exposure to market risk. Panels B and C present results controlling for the Fama-French three factor model exposures (i.e. market, size, value) and five factor model exposures (i.e. market, size, value, investments, profitability), respectively. We plot the cumulative value-weighted portfolio returns for the 'High' portfolio (in green) and the Low portfolio (in red) as well as the High-Low differential return (in blue). The dashed vertical lines mark February 24, the day after Italy introduced its lockdown, and March 20, the last trading day before the Fed announced its intervention.

Panel A. CAPM-adjusted returns

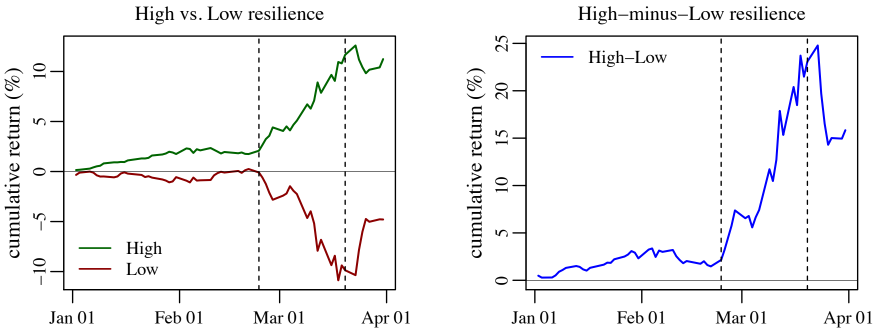

Panel B. FF3-adjusted returns

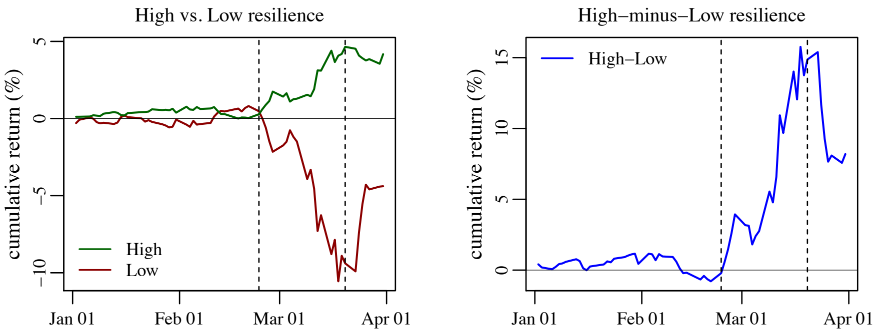

Panel C. FF5-adjusted returns

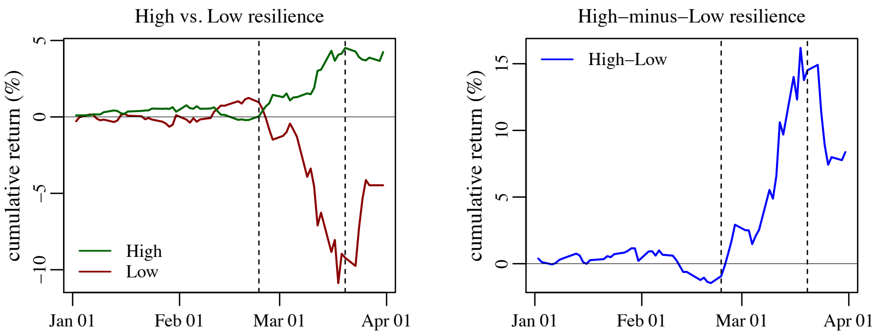

Table 2: Excess and risk-adjusted returns of high and low resilience stocks

This table reports averages of daily returns for value-weighted portfolios of low and high resilience stocks from February 24 to March 20, 2020, i.e. from the day after Italy introduced its lockdown to the last trading day before the Fed announced its intervention. On any given day, we assign a firm to the 'High' ('Low') portfolio, in Panel A, if its 'affected_share' (as defined by Koren and Pető, 2020) is below (above) the median value; in Panel B, if its 'teleworkable_manual_wage' (as defined by Dingel and Neiman, 2020) is above (below) the median value; in Panel C, if its 'dur_workplace' (as defined by Hensvik et al., 2020) is below (above) the median value. and to the 'Low' portfolio if it is above. In addition to rae excess returns (ret), we report CAPM-adjusted returns, i.e. controlling for exposure to market risk, returns adjusted for the Fama-French three factor model exposures (ff3, i.e. market, size, value) and five factor model exposures (ff5, i.e. market, size, value, investments, profitability), and the Fama-French models augmented by the momentum factor (ff4 and ff6). Values in square brackets are t-statistics based on standard errors following Newey and West (1987), where we choose the optimal truncation lag as suggested by Andrews (1991). *, ** and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Panel A. Resilience based on 'affected_share' as in Koren and Pető (2020)

	ret	capm	ff3	ff4	ff5	ff6
High resilience	-1.48^{***}	0.47^{***}	0.23^{***}	0.16^{**}	0.23^{***}	0.17^{**}
	[-3.02]	[6.65]	[2.69]	[2.06]	[3.04]	[2.29]
Low resilience	-1.88^{***}	-0.52^{***}	-0.52^{***}	-0.42^{***}	-0.53^{***}	-0.44^{***}
	[-3.84]	[-3.00]	[-2.89]	[-2.65]	[-2.90]	[-2.70]
High-minus-Low	0.40^{***}	0.99^{***}	0.75^{***}	0.59^{***}	0.77^{***}	0.61^{***}
	[3.22]	[4.26]	[2.94]	[2.65]	[3.12]	[2.85]

Panel B. Resilience based on 'teleworkable_manual_wage' as in Dingel and Neiman (2020)

	ret	capm	ff3	ff4	ff5	ff6
High resilience	-1.58^{***}	0.35^{***}	0.08	0.04	0.09	0.07
	[-3.65]	[3.36]	[0.59]	[0.32]	[1.24]	[1.06]
Low resilience	-1.77^{***}	-0.39^{**}	-0.40^{**}	-0.33^{*}	-0.41^{**}	-0.38^{**}
	[-3.43]	[-2.03]	[-2.11]	[-1.86]	[-2.36]	[-2.27]
High-minus-Low	0.19	0.74^{***}	0.48^{*}	0.37	0.49^{**}	0.45^{**}
	[1.16]	[2.69]	[1.66]	[1.28]	[2.22]	[2.09]

Panel C. Resilience based on 'dur_workplace' as in Hensvik et al. (2020)

	ret	capm	ff3	ff4	ff5	ff6
High resilience	-1.42^{***}	0.46^{***}	0.10	0.08	0.11	0.11
	[-3.34]	[4.64]	[0.77]	[0.63]	[1.11]	[1.12]
Low resilience	-1.89^{***}	-0.35^{**}	-0.31^{**}	-0.31^{**}	-0.31^{*}	-0.35^{**}
	[-3.73]	[-2.53]	[-2.10]	[-2.09]	[-2.01]	[-2.11]
High-minus-Low	0.46^{***}	0.81^{***}	0.41	0.39	0.42^{*}	0.46^{**}
	[3.40]	[3.90]	[1.67]	[1.61]	[1.92]	[1.99]

Covid Economics 21, 22 May 2020: 1-39

measures is that Low-resilience stocks generate (at least marginally) significant negative excess returns in all return specifications. For KP, we additionally find that all risk-adjusted returns of the High-resilience portfolio are significantly positive in all specifications.

Tables A.2 to A.4 in the Appendix present results for the other metrics proposed by KP, DN, and HLR, respectively. For KP (Table A.2), most return differentials are significantly different from zero when controlling for CAPM- and FF-exposures. For 'presence_share', the measure that aims to capture the necessity of working in close proximity to others, the results are even stronger than those based on 'affected_share'. For DN and HLR, relying on the other proxies generally reduces the significance of the results, but CAPM-adjusted return differentials remain statistically significant (see Tables A.3 and A.4, respectively).

To better understand the source of the High-minus-Low differential returns, we study the cumulative risk-adjusted returns in the cross-section of (value-weighted) industry portfolios and present results for the 25 industries with the highest number of firms, in total 2,974 firms. Table 3 presents summary statistics.

Figure 3 plots the cumulative risk-adjusted returns of value-weighted industry portfolios, ranked by their resilience to pandemics, based on KP's 'affected_share' variable. The figure shows that less resilient industries feature substantially lower cumulative risk-adjusted returns during the COVID-19 crisis: the stocks of the least resilient industries (such as those in NAICS-industry 212: Mining, except oil and gas; 483: Water transportation) generated returns 40% to 50% lower than the most resilient ones (224: Computer and electronic products; 511: Publishing industries, except Internet), depending on the risk adjustment. The cross-sectional relationship between pandemic resilience and cumulative returns is not only highly statistically significant in all three panels of the figure, but also economically significant: for instance, a decrease of 10 percentage points in the resilience metric is associated with a drop of 7.2% in the CAPM-adjusted cumulative return. In the Appendix, Figure A.3 shows qualitatively similar results when resilience is measured on the basis of DN's 'teleworkable_manual_wage' variable.

4.3 Returns and Firm Resilience before the Disaster

The evidence in the previous section indicates that *when* the public became aware of the COVID-19 outbreak, stock returns reacted differently depending on companies'

Covid Economics 21, 22 May 2020: 1-39

Table 3: Summary statistics for industry portfolios

This table presents summary statistics for the returns of value-weighted industry portfolios from February 24 to March 20, 2020, i.e. from the day after Italy introduced its lockdown to the last trading day before the Fed announced its intervention. We define resilience as 100 (%) minus the 'affected_share' defined by Koren and Pető (2020) and present results for the 25 industries with the highest number of firms (in total 2,974). We report the industries' 3-digit NAICS code, their description, the number of firms in the respective industries and their cumulative return. Specifically, we present CAPM-adjusted returns, i.e. controlling for exposure to market risk as well as results controlling for the Fama-French three factor model exposures (ff3, i.e. market, size, value) and five factor model exposures (ff5, i.e. market, size, value, investments, profitability).

NAICS	description	firms	resilience	capm	ff3	ff5
211	Oil and gas extraction	88	70	−26.14	−7.87	−6.44
212	Mining, except oil and gas	87	29	−29.04	−34.09	−33.88
213	Support activities for mining	37	46	−30.18	−8.02	−6.36
221	Utilities	94	54	−24.79	−28.64	−29.01
311	Food manufacturing	53	77	−3.80	−6.13	−6.73
325	Chemicals	639	79	7.56	2.30	2.24
332	Fabricated metal products	54	79	−7.50	−3.72	−3.44
333	Machinery	118	80	3.37	12.06	12.45
334	Computer and electronic products	327	87	19.52	15.50	15.59
335	Electrical equipment and appliances	44	83	0.01	11.05	11.37
336	Transportation equipment	97	81	−17.28	−14.74	−14.40
339	Miscellaneous durable goods manufacturing	89	84	−1.32	−9.29	−9.26
423	Wholesale trade: Durable goods	58	68	−3.44	2.79	3.19
424	Wholesale trade: Nondurable goods	49	71	−14.89	−13.12	−12.80
483	Water transportation	52	26	−31.38	−24.68	−23.87
511	Publishing industries, except Internet	92	84	20.02	7.47	7.48
515	Broadcasting, except Internet	71	65	−11.69	−9.90	−10.10
518	Data processing, hosting and related services	71	81	12.38	2.19	2.73
519	Other information services	161	76	11.39	−1.19	−0.71
523	Securities, commodity contracts, investments, and funds and trusts	136	71	−2.62	5.68	5.70
524	Insurance carriers and related activities	127	72	−12.24	−11.42	−11.57
531	Real estate	222	48	−26.41	−32.80	−32.93
541	Professional and technical services	104	77	−3.58	−9.49	−9.52
561	Administrative and support services	57	65	−3.13	−8.49	−8.56
722	Food services and drinking places	47	47	−19.93	−24.79	−25.38

Covid Economics 21, 22 May 2020: 1-39

Covid Economics 21, 22 May 2020: 1-39

Figure 3. Resilience to social distancing and industry portfolio returns

This figure plots the cumulative risk-adjusted returns of value-weighted industry portfolios against the industries' resilience to disaster risk. The sample period is from February 24 to March 20, 2020, i.e. from the day after Italy introduced its lockdown to the last trading day before the Fed announced its intervention. We define resilience as 100 (%) minus the 'affected_share' defined by Koren and Pető (2020) and present results for the 25 industries with the highest number of firms (in total 2,974). In Panel A, we present CAPM-adjusted returns, i.e. controlling for exposure to market risk. Panels B and C present results controlling for the Fama-French three factor model exposures (i.e. market, size, value) and five factor model exposures (i.e. market, size, value, investments, profitability), respectively. The plot labels indicate the industries' 3-digit NAICS codes. The plot legends report results for cross-sectional regressions with t-statistics based on White (1980) standard errors in square brackets.

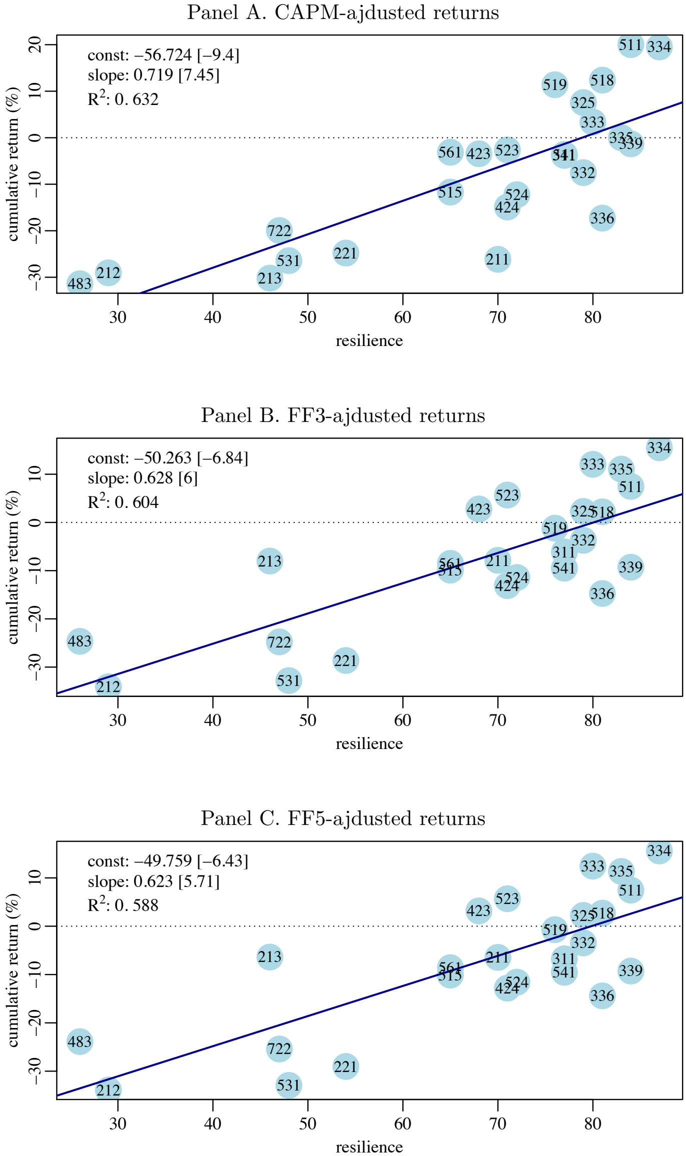

Covid Economics 21, 22 May 2020: 1-39

resilience to the social distancing rules triggered by the pandemic. However, in principle investors may have been aware of pandemic risk even before the COVID-19 shock, so that this risk may have to some extent been priced by the stock market in advance. As explained in Section 2, if investors were fully aware of such risk in advance, they should have required lower expected returns on the stocks of pandemic-resilient companies than on those of non-resilient ones, controlling for their respective exposures to other risks. If instead investors had become gradually aware of such risk, one should observe the opposite pattern, namely, the stocks of pandemic-resilient companies outperforming non-resilient ones. Finally, if investors were completely unaware of such risk, one should not detect any difference in the stock market performance of the two types of companies, prior to the pandemic.

Figure 4 provides evidence on this point, by displaying the time series pattern of risk-adjusted cumulative returns for High- and Low-resilience stocks for six years before the COVID-19 crisis, as well as for a High-minus-Low-resilience portfolio.[5] Irrespective of the risk adjustment considered (CAPM, FF3 or FF5), the figure shows that High-resilience stocks vastly outperformed Low-resilience ones, with most of the differential return stemming from the outperformance of the former rather than the underperformance of the latter. Moreover, about half of the cumulative risk-adjusted return differential between the two portfolios over the whole interval from 2014 to early 2020 materialized *before* the COVID-19 crisis, the spike occurring *during* the crisis accounting for the other half.

Hence, this evidence appears consistent with the third hypothesis outlined in Section 2, namely, that investors have become gradually aware of disaster risk before the current pandemics, and therefore have sought to acquire the stocks of pandemic-resilient stocks at increasingly high prices, to insulate their portfolios against this previously unknown type of risk.

4.4 Option-Implied Expected Returns after the Disaster

To study how the market prices resilience to disaster risk going forward, we rely on equity options data. Options prices are observable in real time and inherently forward-looking. These features are especially useful in the current crisis, in which

[5]The empirical approach is the same as before, i.e. we estimate exposure from daily excess returns over a calendar year and use these exposures to compute risk-adjusted returns in the next year.

Covid Economics 21, 22 May 2020: 1-39

Figure 4. Risk-adjusted returns of high and low resilience stocks prior to the Covid19-crisis

This figure plots the cumulative risk-adjusted returns of portfolios sorted by firms' resilience to disaster risk during from January 2014 through March 2020. On any given day, we assign a firm to the 'High' portfolio if its 'affected_share' (as defined by Koren and Pető, 2020) is below the median value and to the 'Low' portfolio if it is above. In Panel A, we present CAPM-adjusted returns, i.e. controlling for exposure to market risk. Panels B and C present results controlling for the Fama-French three factor model exposures (i.e. market, size, value) and five factor model exposures (i.e. market, size, value, investments, profitability), respectively. We plot the cumulative value-weighted portfolio returns for the 'High' portfolio (in green) and the Low portfolio (in red) as well as the High-Low differential return (in blue).

Panel A. CAPM-adjusted returns

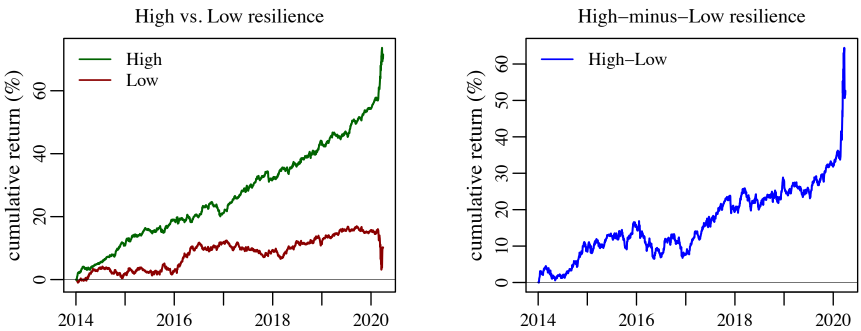

Panel B. FF3-adjusted returns

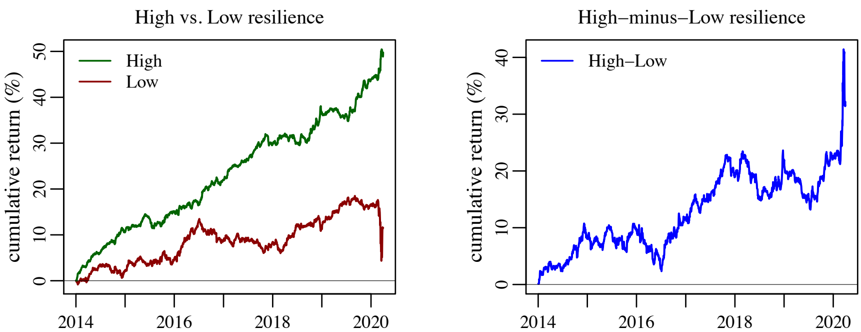

Panel C. FF5-adjusted returns

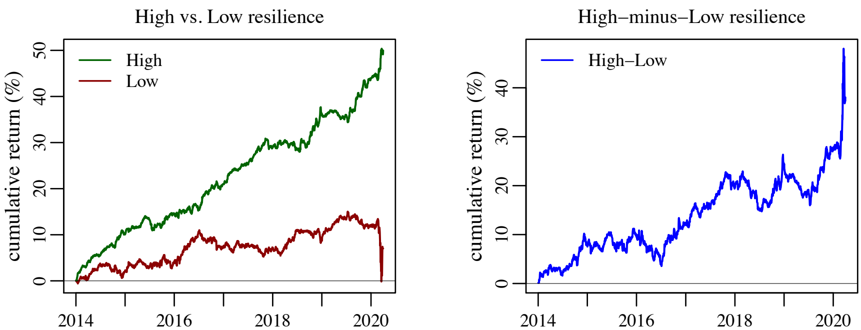

Covid Economics 21, 22 May 2020: 1-39

– due to its unprecedented nature – relying on historical data appears particularly questionable.

Recent research shows how prices of index options and stock options can be used to compute measures of expected market returns and expected stock returns. In our analysis, we follow the approaches suggested by Martin (2017) and Martin and Wagner (2019). Martin (2017) argues that the risk-neutral variance of the market provides a lower bound on the equity premium. He also argues that, empirically, the lower bound is approximately tight, so that the risk-neutral variance of the market directly measures the equity premium. Martin and Wagner (2019) derive a formula for the expected return on a stock in terms of the risk-neutral variance of the market and the stock's excess risk-neutral variance relative to that of the average stock.

The three measures of risk-neutral variance – for the market, a particular stock and the average stock – can be computed from option prices using the approach of Breeden and Litzenberger (1978). The market risk-neutral variance, SVIX_t^2, is determined by the prices of index options:

$$\mathrm{SVIX}_t^2 = \frac{2}{R_{f,t+1}S_{m,t}^2}\left[\int_0^{F_{m,t}} \mathrm{put}_{m,t}(K)\,dK + \int_{F_{m,t}}^{\infty} \mathrm{call}_{m,t}(K)\,dK\right],$$

where $R_{f,t+1}$ is the gross riskfree rate, $S_{m,t}$ and $F_{m,t}$ denote the spot and forward (to time $t+1$) prices of the market, and $\mathrm{put}_{m,t}(K)$ and $\mathrm{call}_{m,t}(K)$ denote the time t prices of European puts and calls on the market that expire at time $t+1$ with strike K. The length of the time interval from t to $t+1$ corresponds to the maturity of the options used in the computation.

The risk-neutral variance at the individual stock level, $\mathrm{SVIX}_{i,t}^2$, is defined in terms of individual stock option prices:

$$\mathrm{SVIX}_{i,t}^2 = \frac{2}{R_{f,t+1}S_{i,t}^2}\left[\int_0^{F_{i,t}} \mathrm{put}_{i,t}(K)\,dK + \int_{F_{i,t}}^{\infty} \mathrm{call}_{i,t}(K)\,dK\right],$$

where the subscripts i indicate the underlying stock i.

Finally, using $\mathrm{SVIX}_{i,t}^2$ for all firms available at time t, we calculate the risk-neutral average stock variance index as $\overline{\mathrm{SVIX}}_t^2 = \sum_i w_{i,t}\,\mathrm{SVIX}_{i,t}^2$.

Using these three risk-neutral variances, one can compute the expected return on

Covid Economics 21, 22 May 2020: 1-39

a stock using the formula derived by Martin and Wagner (2019):

$$\frac{\mathbb{E}_t\, R_{i,t+1} - R_{f,t+1}}{R_{f,t+1}} = \text{SVIX}_t^2 + \frac{1}{2}\left(\text{SVIX}_{i,t}^2 - \overline{\text{SVIX}}_t^2\right).$$

where $R_{i,t+1}$ denotes the one period gross return on stock i. Hence, in the cross-section, differences in expected returns reflect variation in $\text{SVIX}_{i,t}^2$.

In our analysis of expected returns, we use the measures proposed by KP, DN, and HLR to construct risk-neutral variance indices for firms with high resilience to disaster risk, $\overline{\text{SVIX}}_H^2$, and for firms with low resilience $\overline{\text{SVIX}}_L^2$. For every day in our sample, we compute these indices for each resilience measure as the value-weighted sum of the underlying firms' $\text{SVIX}_{i,t}^2$, for maturities ranging from 30 to 730 days. We present results for all stocks for which options are available in OptionMetrics, but to face possible concerns about limited trading in options on small stocks we also present results for the subset of S&P 500 firms in the Appendix.

Figure 5 plots the time series of $\overline{\text{SVIX}}_L^2$ and $\overline{\text{SVIX}}_H^2$, measuring resilience on the basis of KP's 'affected_share' measure, as well as their difference for maturities of 30, 91, 365 and 730 days. The results show that, until early March, the options-implied variance of High-resilience stocks slightly exceeded that for Low-resilience stocks. At the time of the Italian lockdown (February 24), the approximate 2% (p.a.) difference $\overline{\text{SVIX}}_L^2 - \overline{\text{SVIX}}_H^2$ implies that High-resilience firms had higher expected returns than Low-resilience ones. During the Covid-19 crisis, the sign of this premium reverses and its magnitude surges to more than 30% (p.a.) at the peak, as implied by one-month options decreasing with longer maturities to approximately 8% (p.a.) for a forecast horizon of two years. Until the end of March these premia gradually declined to half (interestingly, starting shortly before the reversal in stock prices) and imply, as of March 31, that the stocks of Low-resilience firms are expected to carry a premium of about 5.5% (p.a.) over the next year and about 4% (p.a.) over the next two years. The results are qualitatively very similar when only large firms (i.e., constituents of the S&P 500) are retained in the sample, although these firms feature lower SVIX-levels, as shown in Figure A.4 of the Appendix.

To illustrate these results with reference to some well-known stocks, Figure 6 presents the expected returns of a selected group of S&P firms, respectively featuring high and low resilience to the pandemic, plotting all of them on the same scale, for all stocks and maturities. As examples of very high-resilience firms, the figure plots the option-implied expected returns of Apple, Google, and Microsoft. At the opposite

Covid Economics 21, 22 May 2020: 1-39

Figure 5. Stock options-implied risk-neutral variances

This figure plots stock options-implied risk-neutral variance indices for firms with high and low resilience to social distancing during the first quarter of 2020. On any given day, we assign a firm to the high resilience index, $\overline{\text{SVIX}}^2_{H,t}$, if its 'affected_share' (as defined by Koren and Pető, 2020) is below the median value and to the low resilience index, $\overline{\text{SVIX}}^2_{L,t}$, if it is above. The indices are computed as the value-weighted sums of individual firms' risk-neutral variances, $\text{SVIX}^2_{i,t}$. The difference $\overline{\text{SVIX}}^2_{L,t} - \overline{\text{SVIX}}^2_{L,t}$ measures the expected return of low resilience in excess of high resilience stocks. Panels A to D present results using options maturities of 30, 91, 365 and 730 days, respectively. The dashed vertical lines mark February 24 and March 20.

Panel A. 30-day horizon

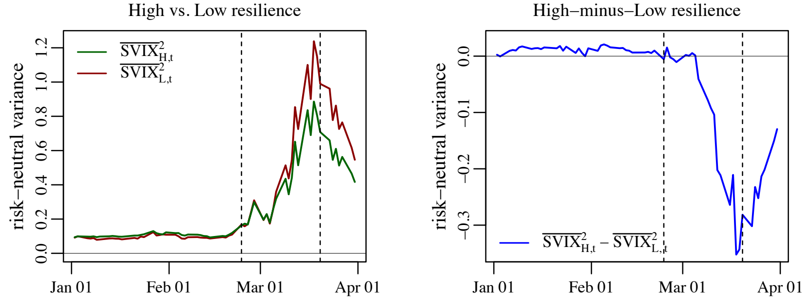

Panel B. 91-day horizon

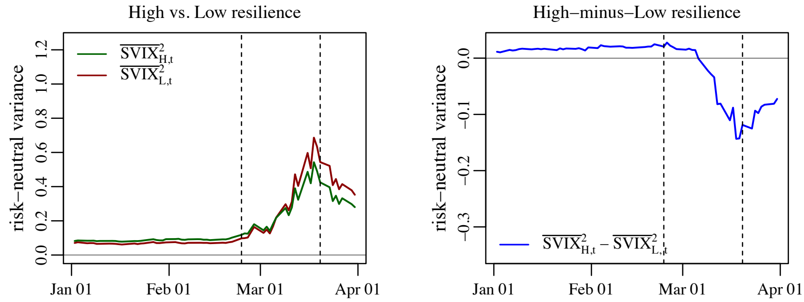

Panel C. 365-day horizon

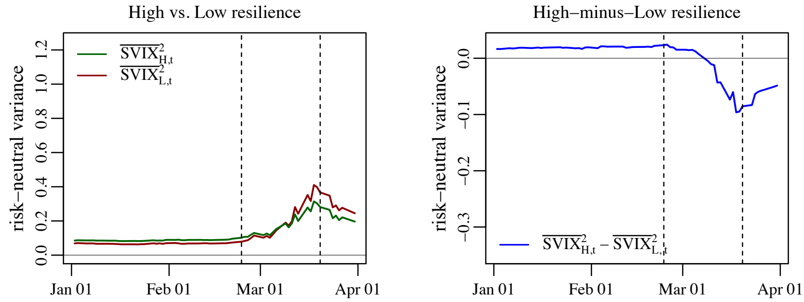

Panel D. 730-day horizon

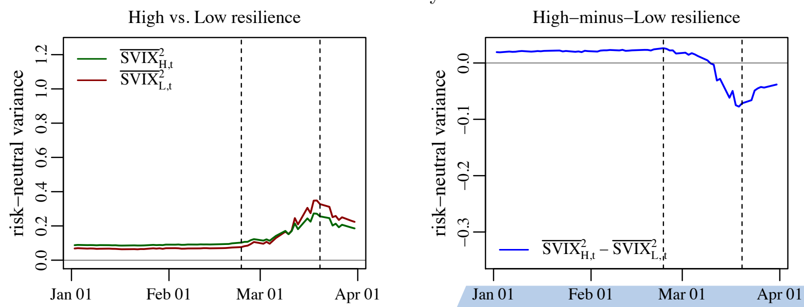

CEPR PRESS

end of the resilience range, as examples of very low-resilience stocks, the figure plots the expected returns of Marriott, United Airlines and Royal Caribbean: travel and tourism have been among the industries hardest hit by the stay-at-home orders and social distancing rules.

Two results emerge strikingly from Figure 6. First, at the outbreak of the COVID-19 crisis, all the option-implied expected returns rose, but those of low-resilience stocks (right-side charts) increased by an order of magnitude more than those of low-resilience ones (left-side charts). At the peak of the crisis, the expected return implied by short-term options became an enormous 300% (p.a.) for United Airlines and Royal Caribbean, reflecting unprecedented uncertainty about the immediate future of their businesses. Second, this increase is much more persistent for low-resilience stocks: at the end of our sample, on March 31, their expected returns are still elevated, while for high-resilience stocks they revert back to pre-COVID-19 levels, especially for the two-year horizon. Third, for all stocks expected returns decrease in levels as maturities increase, indicating that there is a term structure to pandemic risk: it is perceived to decrease substantially as the horizon lengthens, though it far from vanishes for low-resilience stocks.

Figure 7 provides a clearer view of the time-series patterns of the expected returns of the six stocks, as it adapts the vertical scale of the plot to their range of variation. The figure allows in particular to appreciate that, even at the end of our sample, almost two months after the outbreak, investors still require much higher expected returns from Marriott, United Airlines and Royal Caribbean than from Apple, Google and Microsoft, and that even at the two-year horizon the expected return is a multiple of what it was at the beginning of the year before the COVID-19 crisis. The most extreme case is Royal Caribbean, whose two-year options imply, as of March 31, an expected return of 60% (p.a.) for a two-year horizon.

Taken together, these results indicate that disaster resilience is priced in equity options and that the COVID-19 crisis has greatly affected *how* financial markets price resilience to disaster risk: While the two-year expected returns have reverted to their pre-crisis levels for the firms least affected by social distancing requirements, the expected returns for the firms most severely affected by the pandemic are much higher than before the crisis. Hence, it appears that, going forward, markets consider disaster risk, and specifically the resilience against a pandemic, to be much more important than they did before COVID-19.

Covid Economics 21, 22 May 2020: 1-39

CEPR PRESS

Figure 6. Expected returns of selected S&P 500 firms with high and low resilience (same y-axis)
This figure plots stock options-implied expected returns for selected S&P 500 firms during the first quarter of 2020. The high resilience stocks we consider are Apple (AAPL), Google (GOOG), and Microsoft (MSFT), the low resilience stocks are Marriott (MAR), United Airlines (UAL) and Royal Caribbean (RCL). We compute the expected return on a stock in from the risk-neutral variance of the market and the stock's excess risk-neutral variance relative to that of the average stock following the approach of Martin and Wagner (2019). Panels A to D present results for forecast horizons of 30, 91, 365 and 730 days, respectively. The dashed vertical lines mark February 24 and March 20.

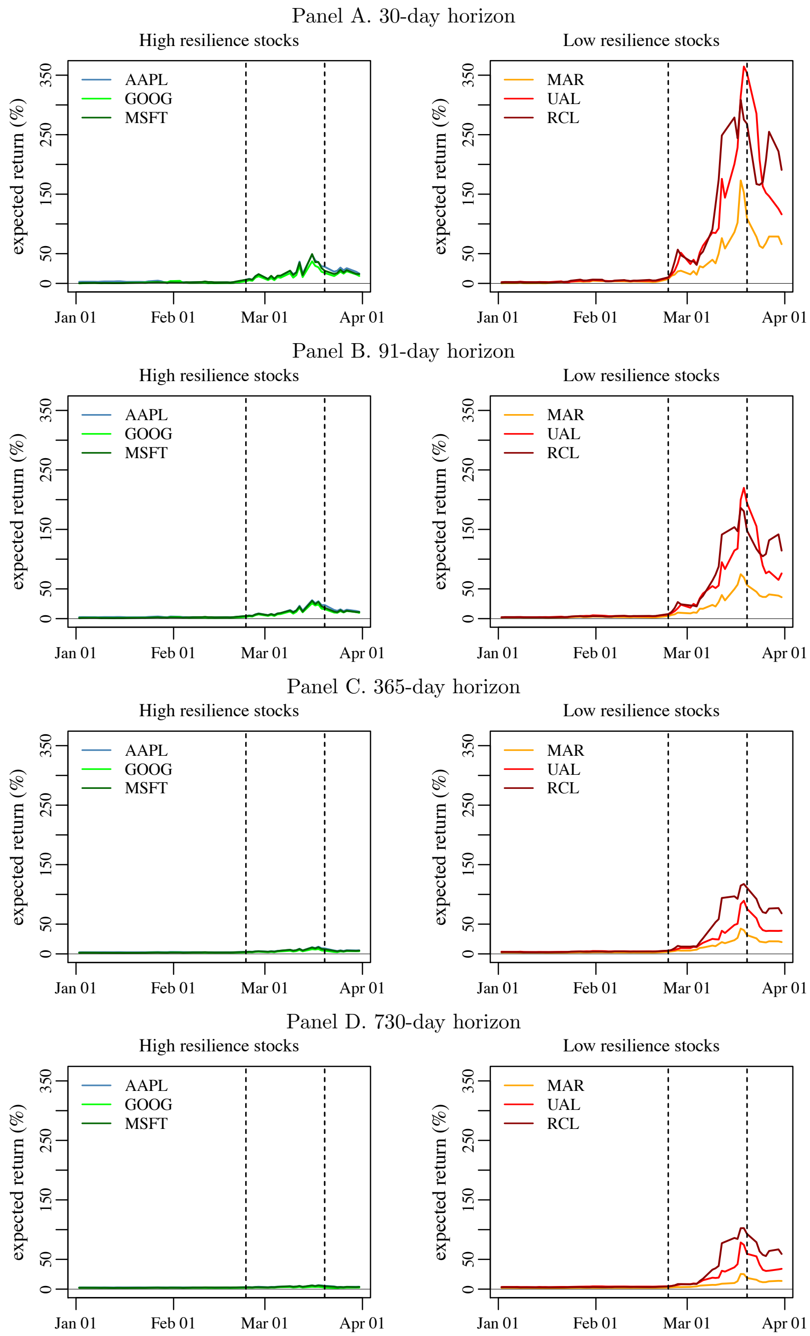

Covid Economics 21, 22 May 2020: 1-39

Figure 7. Expected returns of selected S&P 500 firms with high and low resilience

This figure plots stock options-implied expected returns for selected S&P 500 firms during the first quarter of 2020. The high resilience stocks we consider are Apple (AAPL), Google (GOOG), and Microsoft (MSFT), the low resilience stocks are Marriott (MAR), United Airlines (UAL) and Royal Caribbean (RCL). We compute the expected return on a stock in from the risk-neutral variance of the market and the stock's excess risk-neutral variance relative to that of the average stock following the approach of Martin and Wagner (2019). Panels A to D present results for forecast horizons of 30, 91, 365 and 730 days, respectively. The dashed vertical lines mark February 24 and March 20.

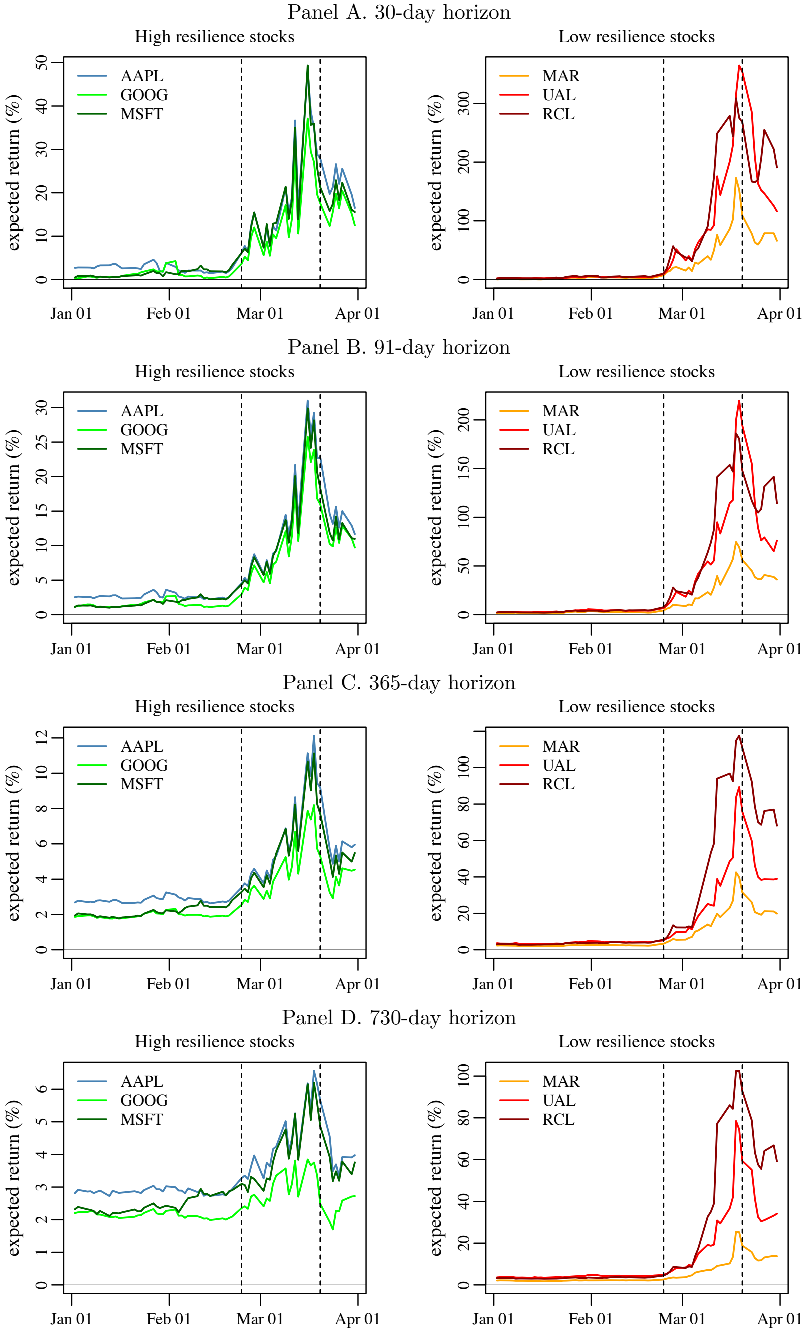

Covid Economics 21, 22 May 2020: 1-39

5 Conclusions

The contribution of this paper is threefold. First, it investigates whether the COVID-19 outbreak triggered a different stock return response depending on companies' resilience to social distancing, which is the most severe constraint that the pandemic has imposed on firms' operations. On this score, we find that more resilient companies greatly outperformed less resilient ones, even after controlling for all conventional measures of risk premia.

Second, the paper explores whether similar cross-sectional return differentials already emerged before the COVID-19 outbreak. Indeed this is the case: in the 2014-19 interval, the cumulative return differential between more and less pandemic-resilient firms is of about the same magnitude as during the outbreak, i.e. between late February and early April 2020. We interpret this as evidence of learning by investors, i.e., of their growing awareness of the potential threat posed by pandemics well in advance of its materialization.

Finally, we exploit option price data to infer whether, after the COVID-19 outbreak, investors price pandemic risk over different horizons, and find that they do: even on a 2-year horizon, stocks of more pandemic-resilient firms are expected to yield significantly lower returns than less resilient ones, reflecting lower exposure to disaster risk. Such differences are massive in the case of some stocks: for example, as late as early April 2020, the expected return on low-resilience stocks such as Royal Caribbean and United Airlines is around 60% and 40% respectively, while those of high-resilience stocks such as Apple and Microsoft are between 3% and 4%.

Hence, going forward, markets appear to price exposure to a new risk factor, namely, pandemic risk. In future development of this work, we plan to investigate whether such risk is part of a wider sustainability risk factor, or at least whether the two types of risk are correlated. We also plan to investigate whether resilience to social distancing has not only direct effects on stock prices, but also indirect effects via demand and supply linkages, i.e. whether for instance the stocks of firms that depend heavily on low-resilience firms are themselves more exposed to pandemic risk, other things equal.

Covid Economics 21, 22 May 2020: 1-39

References

Albuquerque, R. A., Y. Koskinen, S. Yang, and C. Zhang. 2020. Resiliency of environmental and social stocks: an analysis of the exogenous COVID-19 market crash. *Available at SSRN 3583611* .

Alfaro, L., A. Chari, A. Greenland, and P. Schott. 2020. Aggregate and firm-level stock returns during pandemics. *Available at SSRN 3562034* .

Andrews, D. W. K. 1991. Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica* 59:817–858.

Barro, R. J. 2006. Rare Disasters and Asset Markets in the Twentieth Century. *The Quarter Journal of Economics* 121:823–866.

Barro, R. J. 2009. Rare Disasters, Asset Prices, and Welfare Costs. *American Economic Review* 99:243–64.

Breeden, D. T., and R. H. Litzenberger. 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* pp. 621–651.

Bretscher, L., A. Hsu, and A. Tamoni. 2020. The Supply Channel of Uncertainty Shocks and the Cross-Section of Returns: Evidence From the COVID-19 Crisis. *Available at SSRN 3588418* .

Carhart, M. M. 1997. On persistence in mutual fund performance. *The Journal of finance* 52:57–82.

Coibion, O., Y. Gorodnichenko, and M. Weber. 2020a. The Cost of the COVID-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending. *Chicago Booth Research Paper* .

Coibion, O., Y. Gorodnichenko, and M. Weber. 2020b. Labor Markets During the COVID-19 Crisis: A Preliminary View. Tech. rep., National Bureau of Economic Research.

Croce, M. M. M., P. Farroni, and I. Wolfskeil. 2020. When the markets get covid: Contagion, viruses, and information diffusion. *Available at SSRN 3560347* .

Ding, W., R. Levine, C. Lin, and W. Xie. 2020. Corporate Immunity to the COVID-19 Pandemic. *NBER Working Paper 27055* .

Dingel, J., and B. Neiman. 2020. How many jobs can be done at home? *NBER Working Paper 26948* .

Fama, E. F., and K. R. French. 1993. Common Risk Factors in the Returns on Stocks and Bonds. *Journal of Financial Economics* 33:3–56.

Fama, E. F., and K. R. French. 2015. A Five-Factor Asset Pricing Model. *Journal of Financial Economics* 116:1–22.

Gabaix, X. 2012. Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance. *The Quarterly Journal of Economics* 127:645–700.

Gerding, F., T. Martin, and F. Nagler. 2020. The value of fiscal capacity in the face of a rare disaster. *Available at SSRN 3572839* .

Gillman, M., M. Kejak, and M. Pakoš. 2014. Learning about Rare Disasters: Implications For Consumption and Asset Prices. *Review of Finance* 19:1053–1104.

Gormsen, N. J., and R. S. Koijen. 2020. Coronavirus: Impact on stock prices and growth expectations. *University of Chicago, Becker Friedman Institute for Economics Working Paper* .

Hanspal, T., A. Weber, and J. Wohlfart. 2020. Income and Wealth Shocks and Expectations during the COVID-19 Pandemic. *CEBI Working Paper Series* .

Hassan, T. A., S. Hollander, L. van Lent, and A. Tahoun. 2020. Firm-level Exposure to Epidemic Diseases: Covid-19, SARS, and H1N1. *NBER Working Paper 26971* .

Hensvik, L., T. Le Barbanchon, and R. Rathelot. 2020. Which jobs are done from home? Evidence from the American Time Use Survey. *CEPR Discussion Paper14611* .

Koren, M., and R. Pető. 2020. Business disruptions from social distancing. *Covid Economics 2* pp. 13–31.

Lu, Y. K., and M. Siemer. 2016. Learning, Rare Disasters, and Asset Prices. Tech. Rep. 1977867, SSRN.

Lucas, R. E. 1978. Asset prices in an exchange economy. *Econometrica: Journal of the Econometric Society* pp. 1429–1445.

Martin, I. 2017. What is the Expected Return on the Market? *The Quarterly Journal of Economics* 132:367–433.

Martin, I. W., and C. Wagner. 2019. What is the Expected Return on a Stock? *The Journal of Finance* 74:1887–1929.

Newey, W. K., and K. D. West. 1987. A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55:703–708.

Pastor, L., R. F. Stambaugh, and L. A. Taylor. 2020. Sustainable Investing in Equilibrium. *Working Paper* .

Ramelli, S., and A. F. Wagner. 2020. Feverish stock price reactions to Covid-19 .

Rietz, T. A. 1988. The equity risk premium a solution. *Journal of monetary Economics* 22:117–131.

Ru, H., E. Yang, and K. Zou. 2020. What Do We Learn from SARS-CoV-1 to SARS-CoV-2: Evidence from Global Stock Markets. *Available at SSRN 3569330* .

Veronesi, P. 2004. The peso problem hypothesis and stock market returns. *Journal of Economic Dynamics and Control* 28:707–725.

Wachter, J. A. 2013. Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance* 68:987–1035.

Wachter, J. A., and Y. Zhu. 2019. Learning with Rare Disasters. Tech. Rep. 3407397, SSRN.

White, H. 1980. A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroscedasticity. *Econometrica* 48:817–838.

Xu, S., and Y. Li. 2020. Beware of the Second Wave of COVID-19. *The Lancet* 395:1321–1322.

Covid Economics 21, 22 May 2020: 1-39

CEPR PRESS

Appendix for

"Disaster Resilience and Stock Returns"

Marco Pagano Christian Wagner Josef Zechner

This Appendix provides additional results referred to in the paper.

Covid Economics 21, 22 May 2020: 1-39

COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

Covid Economics 21, 22 May 2020: 1-39

Figure A.1. Risk-adjusted returns of stocks with high and low resilience to social distancing (DN)

This figure plots the cumulative risk-adjusted returns of portfolios sorted by firms' resilience to disaster risk during the first quarter of 2020. On any given day, we assign a firm to the 'High' portfolio if its 'teleworkable_manual_wage' (as defined by Dingel and Neiman, 2020) is above the median value and to the 'Low' portfolio if it is below. In Panel A, we present CAPM-adjusted returns, i.e. controlling for exposure to market risk. Panels B and C present results controlling for the Fama-French three factor model exposures (i.e. market, size, value) and five factor model exposures (i.e. market, size, value, investments, profitability), respectively. We plot the cumulative value-weighted portfolio returns for the 'High' portfolio (in green) and the Low portfolio (in red) as well as the High-Low differential return (in blue). The dashed vertical lines mark February 24, the day after Italy introduced its lockdown, and March 20, the last trading day before the Fed announced its intervention.

Panel A. CAPM-adjusted returns

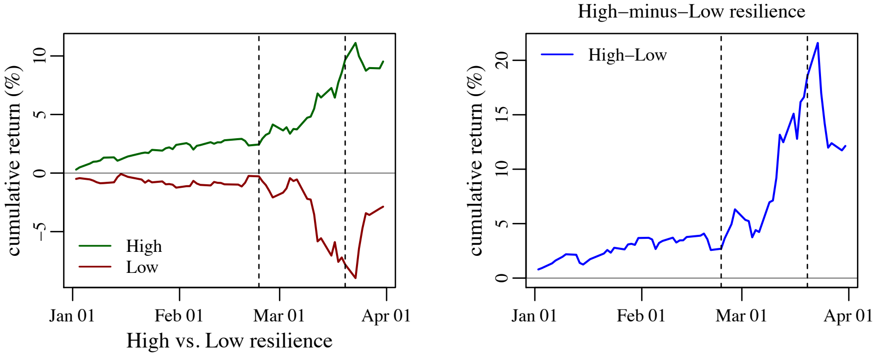

Panel B. FF3-adjusted returns

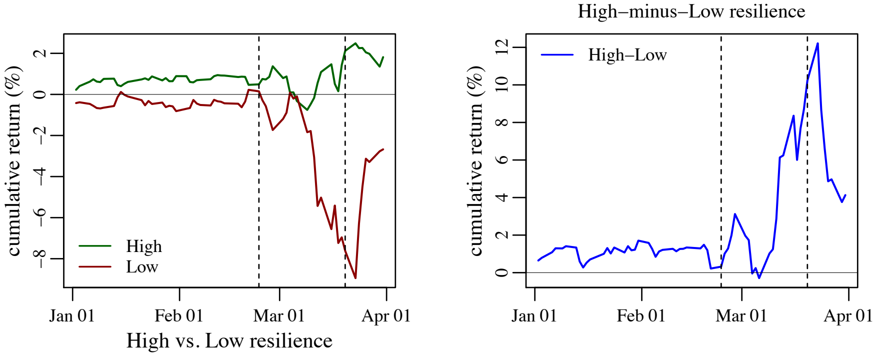

Panel C. FF5-adjusted returns

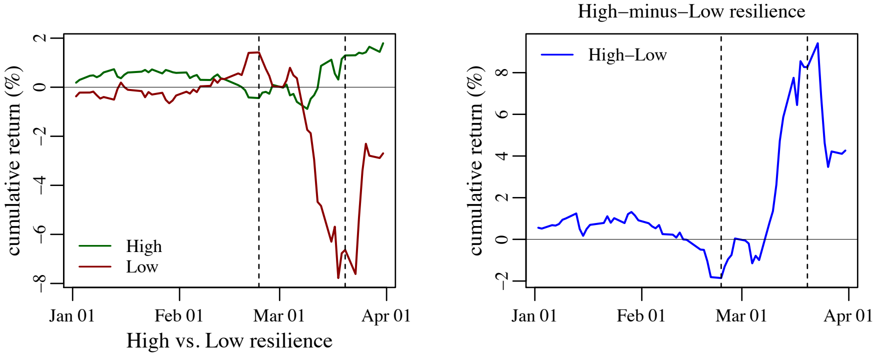

Covid Economics 21, 22 May 2020: 1-39

Figure A.2. Risk-adjusted returns of stocks with high and low resilience to social distancing (HLRN)

This figure plots the cumulative risk-adjusted returns of portfolios sorted by firms' resilience to disaster risk during the first quarter of 2020. On any given day, we assign a firm to the 'High' portfolio if its 'dur_workplace' (as defined by Hensvik et al., 2020) is below the median value and to the 'Low' portfolio if it is above. In Panel A, we present CAPM-adjusted returns, i.e. controlling for exposure to market risk. Panels B and C present results controlling for the Fama-French three factor model exposures (i.e. market, size, value) and five factor model exposures (i.e. market, size, value, investments, profitability), respectively. We plot the cumulative value-weighted portfolio returns for the 'High' portfolio (in green) and the Low portfolio (in red) as well as the High-Low differential return (in blue). The dashed vertical lines mark February 24, the day after Italy introduced its lockdown, and March 20, the last trading day before the Fed announced its intervention.

Panel A. CAPM-adjusted returns

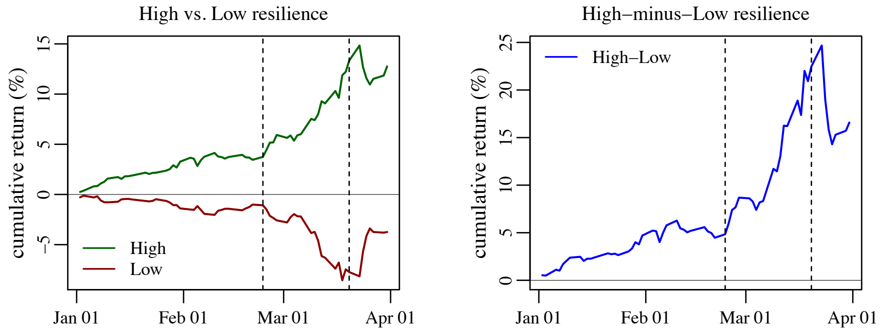

Panel B. FF3-adjusted returns

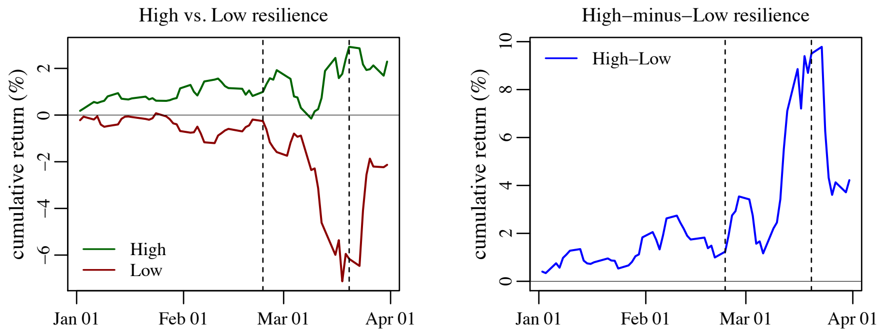

Panel C. FF5-adjusted returns

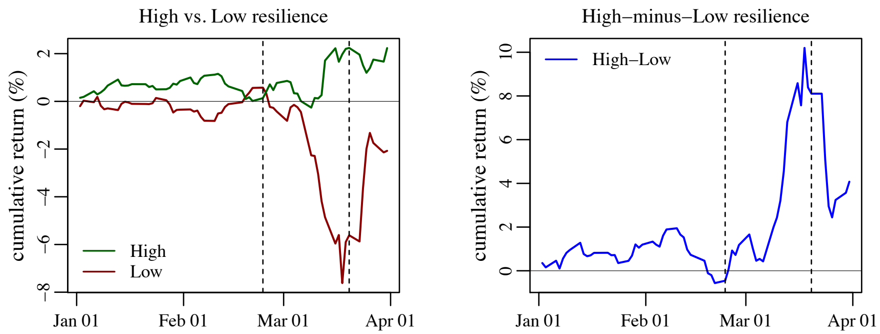

Covid Economics 21, 22 May 2020: 1-39

Figure A.3. Resilience to social distancing and industry portfolio returns (DN)

This figure plots the cumulative risk-adjusted returns of value-weighted industry portfolios against the industries' resilience to disaster risk. The sample period is from February 24 to March 20, 2020, i.e. from the day after Italy introduced its lockdown to the last trading day before the Fed announced its intervention. We define resilience as 100 (%) minus the 'teleworkable_manual_wage' defined by Dingel and Neiman (2020) and present results for the 25 industries with the highest number of firms (in total 2,974). In Panel A, we present CAPM-adjusted returns, i.e. controlling for exposure to market risk. Panels B and C present results controlling for the Fama-French three factor model exposures (i.e. market, size, value) and five factor model exposures (i.e. market, size, value, investments, profitability), respectively. The plot labels indicate the industries' 3-digit NAICS codes. The plot legends report results for cross-sectional regressions with t-statistics based on White (1980) standard errors in square brackets.

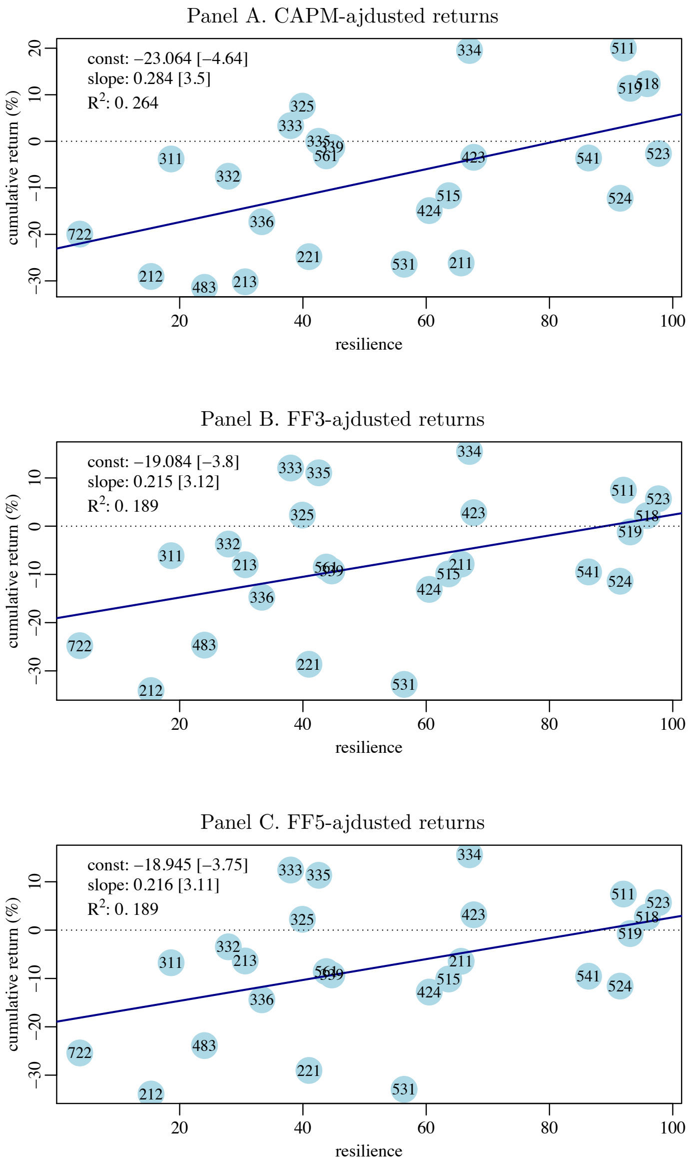

COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

Figure A.4. Stock options-implied risk-neutral variances of S&P 500 firms

This figure plots stock options-implied risk-neutral variance indices for S&P 500 firms with high and low resilience to social distancing during the first quarter of 2020. On any given day, we assign a firm to the high resilience index, $\overline{\text{SVIX}}^2_{H,t}$, if its 'affected_share' (as defined by Koren and Pető, 2020) is below the median value and to the low resilience index, $\overline{\text{SVIX}}^2_{L,t}$, if it is above. The indices are computed as the value-weighted sums of individual firms' risk-neutral variances, $\text{SVIX}^2_{i,t}$. The difference $\overline{\text{SVIX}}^2_{L,t} - \overline{\text{SVIX}}^2_{L,t}$ measures the expected return of low resilience in excess of high resilience stocks. Panels A to D present results using options maturities of 30, 91, 365 and 730 days, respectively. The dashed vertical lines mark February 24 and March 20.

Panel A. 30-day horizon

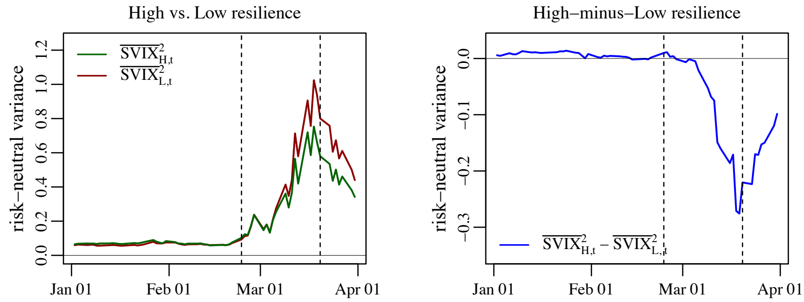

Panel B. 91-day horizon

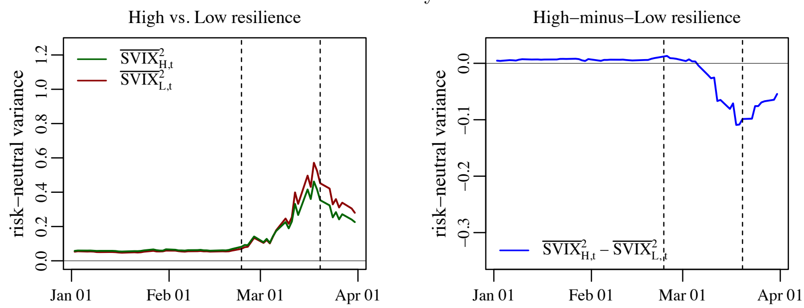

Panel C. 365-day horizon

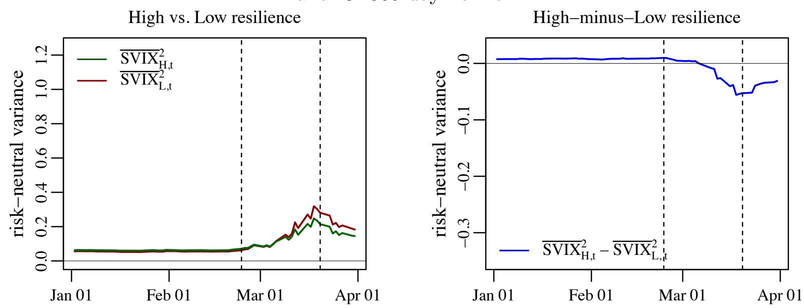

Panel D. 730-day horizon

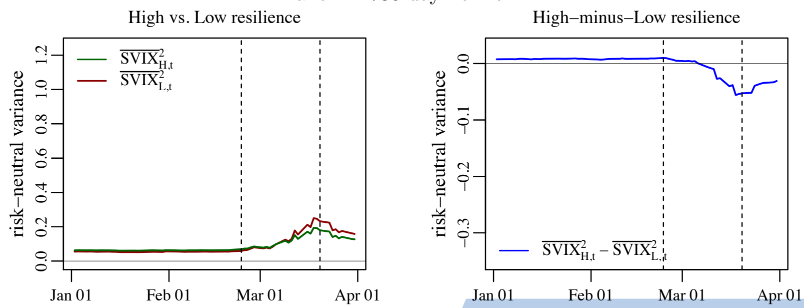

CEPR PRESS

Covid Economics 21, 22 May 2020: 1-39

Table A.1: Measures of teleworkability, working at home and at the workplace, and business face-to-face interactions

This Table provides an overview of the empirical measures on which we base our analysis of stocks' disaster resilience. Panel A lists the teleworkability measures provided by Dingel and Neiman (2020) for 24 industries at the NAICS 2-digit level and for 88 industries at the NAICS 3-digit level. Panel B lists the work at home and work at the workplace measures provided by Hensvik et al. (2020) for 310 industries at the NAICS 4-digit level. Panel C lists the communication-intensity and physical proximity measures suggested by Koren and Pető (2020) for 84 industries at the NAICS 3-digit level.

Panel A. Dingel and Neiman (2020):

'teleworkable_emp'	fraction of jobs that can be done from home estimated from O*Net data
'teleworkable_wage'	fraction of wages to jobs that can be done from home estimated from O*Net data
'teleworkable_manual_emp'	fraction of jobs that can be done from home based on manual classification by the authors
'teleworkable_manual_wage'	fraction of wages to jobs that can be done from home based on manual classification by the authors

Panel B. Hensvik et al. (2020)

'home'	fraction of respondents that work at home
'workplace'	fraction of respondents that work at workplace
'dur_home'	hours worked at home per day
'dur_workplace'	hours worked at workplace per day
'share_home'	hours worked at home divided by hours worked at home and at workplace

Panel C. Koren and Pető (2020)

'teamwork_share'	percentage of workers in teamwork-intensive occupations, i.e. internal communication
'customer_share'	percentage of workers in customer-facing occupations, i.e. external communication
'communication_share'	percentage of workers in teamwork-intensive and/or customer-facing occupations
'presence_share'	percentage of workers whose jobs require physical presence in close proximity to others
'affected_share'	percentage of workers in occupations that are communication-intensive and/or require physical presence in close proximity to others

Table A.2: Other measures of Koren and Pető (2020)

This table presents results analogous to Panel A in Table 2 but using the other measures constructed by Koren and Pető (2020) instead of their 'affected_share' variable.

Panel A. Portfolios sorted by 'teamwork_share'

	ret	capm	ff3	ff4	ff5	ff6
High resilience	−1.67***	0.30***	0.09	0.05	0.10	0.07
	[-3.79]	[2.72]	[0.75]	[0.44]	[1.48]	[1.58]
Low resilience	−1.86***	−0.51**	−0.59**	−0.57**	−0.59**	−0.58**
	[-3.38]	[-2.00]	[-2.42]	[-2.35]	[-2.37]	[-2.33]
High-minus-Low	0.19	0.81***	0.68**	0.61**	0.69**	0.66**
	[1.01]	[2.58]	[2.21]	[2.06]	[2.35]	[2.27]

Panel B. Portfolios sorted by 'customer_share'

	ret	capm	ff3	ff4	ff5	ff6
High resilience	−1.60***	0.21***	0.16***	0.02	0.16***	−0.01
	[-3.33]	[4.51]	[4.10]	[0.26]	[3.24]	[-0.14]
Low resilience	−1.68***	−0.05	−0.29***	−0.20***	−0.30***	−0.18***
	[-3.36]	[-0.74]	[-3.49]	[-2.87]	[-4.14]	[-3.94]
High-minus-Low	0.08	0.26***	0.45***	0.22***	0.46***	0.17***
	[0.96]	[15.17]	[6.51]	[2.69]	[9.94]	[5.07]

Panel C. Portfolios sorted by 'communication_share'

	ret	capm	ff3	ff4	ff5	ff6
High resilience	−1.59***	0.23***	0.17***	0.05	0.17***	0.01
	[-3.39]	[5.46]	[15.28]	[0.89]	[3.61]	[0.19]
Low resilience	−1.77***	−0.20**	−0.40***	−0.30***	−0.40***	−0.29***
	[-3.63]	[-2.06]	[-3.74]	[-3.67]	[-3.85]	[-3.96]
High-minus-Low	0.18**	0.43***	0.56***	0.35***	0.57***	0.30***
	[2.02]	[5.62]	[5.61]	[6.72]	[6.14]	[12.44]

Panel D. Portfolios sorted by 'presence_share'

	ret	capm	ff3	ff4	ff5	ff6
High resilience	−1.57***	0.41***	0.16	0.11	0.16**	0.14***
	[-3.63]	[3.09]	[1.07]	[0.85]	[2.49]	[2.80]
Low resilience	−2.06***	−0.79***	−0.74***	−0.68***	−0.75***	−0.73***
	[-3.78]	[-3.31]	[-3.05]	[-2.98]	[-3.16]	[-3.14]
High-minus-Low	0.49***	1.20***	0.90***	0.79**	0.92***	0.87***
	[2.90]	[3.57]	[2.68]	[2.50]	[3.27]	[3.23]

Covid Economics 21, 22 May 2020: 1-39

CEPR PRESS

Table A.3: Other measures of Dingel and Neiman (2020)

This table presents results analogous to Panel B in Table 2 but using the other measures constructed by Dingel and Neiman (2020) instead of their 'teleworkable_manual_wage' variable.

Panel A. Portfolios sorted by "teleworkable_wage"

	ret	capm	ff3	ff4	ff5	ff6
High resilience	-1.65^{***}	0.29^{***}	0.05	-0.01	0.05	0.02
	[-3.67]	[2.80]	[0.40]	[-0.12]	[1.46]	[0.65]
Low resilience	-1.66^{***}	-0.31	-0.34^{*}	-0.26	-0.35^{**}	-0.31^{*}
	[-3.24]	[-1.59]	[-1.77]	[-1.39]	[-2.08]	[-1.89]
High-minus-Low	0.01	0.60^{**}	0.39	0.24	0.41^{**}	0.33^{*}
	[0.06]	[2.20]	[1.36]	[0.81]	[1.99]	[1.67]

Panel B. Portfolios sorted by "teleworkable_manual_emp"

	ret	capm	ff3	ff4	ff5	ff6
High resilience	-1.58^{***}	0.25^{***}	-0.02	-0.03	-0.02	-0.01
	[-3.34]	[3.74]	[-0.23]	[-0.32]	[-0.37]	[-0.14]
Low resilience	-1.77^{***}	-0.26^{*}	-0.24^{*}	-0.23	-0.24	-0.28^{*}
	[-3.60]	[-1.80]	[-1.74]	[-1.70]	[-1.60]	[-1.78]
High-minus-Low	0.18	0.51^{***}	0.22	0.20	0.22	0.27
	[1.56]	[2.93]	[1.02]	[0.96]	[1.21]	[1.41]

Panel C. Portfolios sorted by "teleworkable_emp"

	ret	capm	ff3	ff4	ff5	ff6
High resilience	-1.64^{***}	0.20^{***}	-0.04	-0.07	-0.04	-0.05
	[-3.25]	[3.14]	[-0.51]	[-0.81]	[-0.94]	[-1.00]
Low resilience	-1.68^{***}	-0.20	-0.21^{*}	-0.17	-0.21	-0.22^{*}
	[-3.53]	[-1.44]	[-1.66]	[-1.41]	[-1.65]	[-1.67]
High-minus-Low	0.04	0.40^{***}	0.16	0.10	0.17	0.17
	[0.54]	[2.56]	[0.85]	[0.50]	[1.24]	[1.21]

Covid Economics 21, 22 May 2020: 1-39

COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

Table A.4: Other measures of Hensvik et al. (2020)

This table presents results analogous to Panel C in Table 2 but using the other measures constructed by Hensvik et al. (2020) instead of their 'dur_workplace' variable.

Panel A. Portfolios sorted by 'workplace'

	ret	capm	ff3	ff4	ff5	ff6
High resilience	−1.47***	0.44***	0.13	0.06	0.13	0.08
	[-3.45]	[5.16]	[0.92]	[0.47]	[1.40]	[0.86]
Low resilience	−1.84***	−0.35**	−0.34**	−0.28**	−0.35**	−0.30**
	[-3.68]	[-2.40]	[-2.22]	[-1.91]	[-2.22]	[-1.99]
High-minus-Low	0.37***	0.80***	0.47*	0.33	0.49**	0.38*
	[3.17]	[3.77]	[1.87]	[1.28]	[2.11]	[1.75]

Panel B. Portfolios sorted by 'dur_home'

	ret	capm	ff3	ff4	ff5	ff6
High resilience	−1.63***	0.19***	−0.10	−0.05	−0.10	−0.03
	[-3.46]	[2.53]	[-0.93]	[-0.43]	[-1.54]	[-0.40]
Low resilience	−1.82***	−0.19	−0.16	−0.23*	−0.17	−0.25
	[-3.66]	[-1.58]	[-1.32]	[-1.76]	[-1.12]	[-1.56]
High-minus-Low	0.19*	0.38***	0.06	0.18	0.07	0.22
	[1.88]	[2.71]	[0.33]	[0.89]	[0.40]	[1.16]

Panel C. Portfolios sorted by 'home'

	ret	capm	ff3	ff4	ff5	ff6
High resilience	−1.58***	0.27***	−0.01	0.01	−0.01	0.02
	[-3.49]	[5.89]	[-0.11]	[0.13]	[-0.20]	[0.37]
Low resilience	−1.94***	−0.39**	−0.32*	−0.37*	−0.32	−0.38*
	[-3.77]	[-2.12]	[-1.80]	[-1.98]	[-1.64]	[-1.84]
High-minus-Low	0.36**	0.66***	0.31	0.38	0.31	0.41
	[2.01]	[3.31]	[1.23]	[1.52]	[1.22]	[1.56]

Panel D. Portfolios sorted by 'share_home'

	ret	capm	ff3	ff4	ff5	ff6
High resilience	−1.62***	0.27***	−0.01	−0.03	−0.01	−0.01
	[-3.73]	[3.69]	[-0.08]	[-0.28]	[-0.14]	[-0.08]
Low resilience	−1.84***	−0.30*	−0.28*	−0.26	−0.28*	−0.28*
	[-3.64]	[-1.88]	[-1.79]	[-1.67]	[-1.66]	[-1.63]
High-minus-Low	0.22	0.56***	0.27	0.23	0.27	0.28
	[1.63]	[2.85]	[1.09]	[0.90]	[1.23]	[1.19]

Covid Economics 21, 22 May 2020: 1-39

CEPR PRESS

Covid Economics **Issue 21, 22 May 2020**

Labor markets during the Covid-19 crisis: A preliminary view[1]

Olivier Coibion,[2] Yuriy Gorodnichenko[3] and Michael Weber[4]

Date submitted: 16 May 2020; **Date accepted:** 18 May 2020

Covid Economics 21, 22 May 2020: 40-58

We use a repeated large-scale survey of households in the Nielsen Homescan panel to characterize how labor markets are being affected by the covid-19 pandemic. We document several facts. First, job loss has been significantly larger than implied by new unemployment claims: we estimate 20 million lost jobs by April 8th, far more than jobs lost over the entire Great Recession. Second, many of those losing jobs are not actively looking to find new ones. As a result, we estimate the rise in the unemployment rate over the corresponding period to be surprisingly small, only about 2 percentage points. Third, participation in the labor force has declined by 7 percentage points, an unparalleled fall that dwarfs the three percentage point cumulative decline that occurred from 2008 to 2016. Early retirement almost fully explains the drop in labor force participation both for those survey participants previously employed and those previously looking for work. We find increases in the fraction of those being retired across the whole age distribution with women and blacks driving a large part of the accelerated retirement.

1 We thank the National Science Foundation for financial support in conducting the surveys. We also thank Shannon Hazlett and Victoria Stevens at Nielsen for their assistance with the collection of the PanelViews Survey. Results in this article are calculated based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. Information on availability and access to the data is available at http://research.chicagobooth.edu/nielsen.

2 Professor of Economics at UT Austin.
3 Professor of Economics at UC Berkeley.
4 Professor of Finance at the University of Chicago.

The arrival of the covid-19 virus and the policy responses have led to unprecedented numbers of initial claims for unemployment since early 2020: over 16.5 million by April 4th, 2020, with new claims arriving at a rate of 6-7 million per week. But concerns about state governments' inability to process so many claims in such a short period, combined with the fact that many workers are ineligible for unemployment benefits, has led to concerns that total job losses are being understated by these numbers. Furthermore, because official labor market indicators compiled by the Bureau of Labor Statistics (BLS) take time to be released, the current state of the U.S. labor market remains unclear.

Using new ongoing large-scale surveys of U.S. households much like the ones run by the BLS, we provide some preliminary evidence on the response of labor markets in the U.S. to the current crisis. We focus on three key variables typically measured by the BLS: the employment-to-population ratio, the unemployment rate, and the labor force participation rate. Historically, the employment-to-population ratio and the unemployment rate are near reverse images of one another during recessions as workers move out of employment and into unemployment (or workers in unemployment find it harder to move into employment). More severe recessions also sometimes lead to a phenomenon of "discouraged workers," in which some unemployed workers stop looking for work. This leads them to be reclassified as "out of the labor force" by the BLS definitions, so the unemployment rate can decline along with the labor force participation rate while the employment-to-population ratio shows little recovery, not because the unemployed are finding work but rather because they stop trying to find one. Jointly, these three metrics therefore provide a succinct and informative summary of the state of labor markets.

Using surveys prior to and at the height of the covid-19 crisis, we provide new estimates of how these variables have changed in the last two months. Our most recent numbers are from individuals surveyed April 2nd-8th, 2020, and therefore provide a sneak preview at what the equivalent BLS numbers will likely show when they are ultimately released. Our findings are striking. First, the employment-to-population ratio (the fraction of the adult population reporting that they had a paid job) has declined by about 7.5 percentage points. With an adult (civilian non-institutional) U.S. population of 260 million, this corresponds to nearly 20 million jobs lost as of April 8th. This estimate is in line with (albeit even higher than) new unemployment claims through this time period and confirms the widespread job loss.

Twenty million jobs lost relative to the pre-crisis labor force would correspond to an increase in the unemployment rate of 12.2 percentage points, so to a level of around 16%, were all the newly unemployed looking for work. This is, however, not what we find. When we construct an unemployment rate in an analogous manner as the BLS (i.e., define the unemployed as not working but looking for work),[1] we find an increase in the unemployment rate of only 2 percentage points. This reflects the fact that most of the newly unemployed surveyed are not looking for new work at this time, so they are defined as out of the labor force rather than unemployed.

[1] BLS classifies laid-off workers as unemployed even if they are not looking for a job. Our survey does not differentiate laid-off workers from other form of non-employment and so our measures of unemployment and labor force do not include laid-off workers who are not looking for a job.

Covid Economics 21, 22 May 2020: 40-58

Correspondingly, we document an extraordinary decline in the labor force participation rate of nearly 8 percentage points. In short, we find labor market changes that differ markedly from those of a typical recession.

I Related Literature

We relate to the fast-growing literature studying the economic consequences of the covid-19 pandemic. Binder (2020) shows that 30% - 40% of American are very concerned about the corona crisis, postponed travel and delayed purchases of larger ticket items as early as March 2020 but became more optimistic about the unemployment situation and revised downward their inflation expectations once being told about the cut in the federal funds target rate on March 3rd. Fetzer et al. (2020) show the arrival of the corona virus in a country leads to a large increase in internet searches around the world. In a survey experiment on a U.S. population, they find survey participants vastly overestimate the mortality rate and extent of contagion. Hanspal et al. (2020) study the income and wealth loss in a survey and the impact on expectations about the economic recovery. Barrios and Hochberg (2020) and Allcott et al. (2020) use internet searches, survey data, and travel data from smartphones to document that political partisanship determines the perception of risk associated with covid-19 and non-essential travel activity. In contemporaneous work, Bick and Blandin (2020) and Foote et al. (2020) provide real-time, high-frequency measures of the U.S. labor market following the CPS protocol, Beland et al. (2020) focus on changes in hours worked, wages and the unemployment rate in the U.S., and von Gaudecker et al. (2020) study labor market outcomes in the Netherlands and heterogeneity by education and employment, whereas Adams-Prassl et al. (2020) study the within and across country implications of covid-19 on labor market outcomes. We differ from these studies in that we study provide historical comparisons as benchmark, in our focus on early retirement and the implications for the economic recovery, and the fact that we can study labor market dynamics across several survey waves. Baek et al. (2020) disentangle the effects of stay at home orders from the effect of the pandemic on initial claims and argue these orders can only explain a small fraction of the claims filed in March. Dingel and Neiman (2020) use data from responses to two Occupational Information Network surveys and estimate that about 37% of jobs can be performed from home, whereas Mongey (2020) documents that employees that are less likely to be able to work from home are mainly non-white and without a college degree. On the quantitative side, a growing literature jointly models the dynamics of the pandemic and the economy to quantify the economic costs and benefits of different policies (Atkeson (2020), Barro et al. (2020), Eichenbaum et al. (2020), Guerrieri et al. (2020), Alvarez et al. (2020), and Dietrich et al. (2020)). Finally, our Nielsen survey builds on previous work using the Nielsen panelists to study the formation and updating of economic expectations (Coibion et al. (2019, 2020) and D'Acunto et al (2020a, b)).

II Measuring Labor Markets using the Nielsen Survey

We start by describing our customized Nielsen Homescan surveys and how they can be used to construct measures of labor market outcomes. The pre-crisis wave of the survey was run between January 6th and January 27th 2020

Covid Economics 21, 22 May 2020: 40-58

prior to much of the covid-19 outbreak. Potential participants were those households participating in the Nielsen Homescan, which is a panel of 80,000-90,000 households who track their purchases on a daily basis. Nielsen allows for surveys to be emailed to those households and respondents receive points and prizes for participating in Nielsen surveys. The panel of households used by Nielsen is meant to be representative of the U.S. population in age, size, income, etc. Possible imbalances are corrected using weights provided by Nielsen. We received 18,344 responses to the first wave of the survey.

In this and subsequent waves of the survey, we asked respondents several job-related questions. First, they were asked whether they have a paid job, with answers being either yes or no. Anyone answering yes we define as being employed. Note that this is slightly different from the BLS, which asks respondents whether they have worked in the survey reference week, and those who "did any work at all for pay or profit" are classified as employed. This means some respondents who would be classified as employed by the BLS are classified as non-employed using our question. Consistent with this, we find somewhat lower employment rates (as a share of adult population) in our pre-crisis data than was the case in corresponding BLS surveys. Table 1, for example, shows that the BLS was reporting an employment to population ratio of 61.1% in February 2020 while our survey yielded an employment-to-population ratio of 57.7%.

Respondents who said they did not have a paid job were then asked if they were actively looking for a job, with possible answers being yes or no. We define those who answer yes as unemployed while those who answer no are classified as out of the labor force. Again, this is slightly different from the BLS questionnaire, which asks individuals to select ways in which they had looked for jobs during the prior 4 weeks and only classifies individuals as unemployed if they select answers which indicate a sufficiently active search such as posting resumes, contacting potential employers or filling out applications (i.e., not just scanning newspaper ads). Given that we allow individuals to specify themselves if they are "actively" searching, one might expect that this would also lead to a higher prevalence of unemployment in our Nielsen survey than in the corresponding BLS survey. Consistent with this, our estimated aggregate unemployment rate prior to the crisis is 8.6% while the corresponding BLS estimates for January and February of 2020 were 3.6% and 3.5% respectively. The labor force participation rate, however, is very similar to that estimated by the BLS: 63.1% vs 63.3% respectively. We also want to note that the BLS issued additional guidance to Census Bureau interviewers starting in March due to the special situation and the occurrence of many outliers because of the corona virus. Specifically, individuals that did not work during the reference week because of the corona virus are classified as unemployed on temporary layoff if the survey respondents think that they will be recalled to their jobs within a 6 months period. If they are uncertain, the BLS assumes that they will be recalled and classify them as unemployed.[2] Given the uncertainty about the speed of the economic recovery and the nature of jobs allowed to operate with lockdown restrictions only gradually lifted, these BLS procedures may overstate the degree to which people are in the labor force.

[2] Please see https://www.bls.gov/cps/employment-situation-covid19-faq-march-2020.pdf for details.

Covid Economics 21, 22 May 2020: 40-58

Covid Economics 21, 22 May 2020: 40-58

Table 1. Employment statistics at the state and aggregate levels, Nielsen survey.

	BLS	Nielsen			
		Raw moments		Adjusted moments	
	pre-crisis	pre-crisis	crisis	pre-crisis	crisis
	(1)	(2)	(3)	(4)	(5)
Panel A. State-level data (average)					
Unemployment rate	3.52	7.84	10.93	3.50	6.30
	(0.84)	(2.80)	(5.07)	(2.55)	(4.61)
Labor force participation rate	63.91	62.89	56.94	63.97	57.20
	(3.94)	(7.19)	(8.58)	(8.17)	(9.75)
Employment to population ratio	61.67	58.00	50.77	61.74	54.11
	(4.03)	(7.27)	(8.63)	(7.68)	(9.12)
Panel B. Aggregate level					
Unemployment rate	3.5	8.62	10.95	4.21	6.32
		[0.25]	[0.43]	[0.23]	[0.39]
Labor force participation rate	63.4	63.13	56.57	64.23	56.79
		[0.33]	[0.51]	[0.38]	[0.58]
Employment to population ratio	61.1	57.68	50.38	59.94	52.22
		[0.34]	[0.52]	[0.36]	[0.54]

Notes: Panel A report statistics across states. Averages are in the top row of each subsection and standard deviations are in the bottom rows (in parentheses) of each subsection. Panel B reports statistics for the aggregate level. Point predictions are reported in the top row of each subsection and standard errors are reported in the bottom rows (in square brackets) of each subsection. Column (1) reports pre-crisis data from the Bureau of Labor Statistics (BLS), average values for January-February 2020. These data are based on the Current Population Survey (CPS). Columns (2) and (3) report moments for the Nielsen survey unadjusted for possible differences in design between the CPS and Nielsen surveys. Columns (4) and (5) report moments for the Nielsen survey adjusted for possible differences in design between the CPS and Nielsen surveys. Adjustment is based on state-level regressions and is given by equations (1)-(2).

While some difference in levels is to be expected given that questions are not identical across the two surveys, it is important to verify that they are still capturing broadly similar features. We do so by comparing state-level estimates of all variables from both the pre-crisis BLS and Nielsen surveys. We find a strong positive correlation between BLS and Nielsen data at the state level for unemployment rates, employment-to-population ratios, and labor force participation rates (see Appendix Figure 1)

We can control for the initial difference in levels of unemployment that stems from the different questions asked in the surveys in the following way. Let $UE_{i,t}^{Nielsen}$ be the unemployment rate in state i in the Nielsen survey in year t and $UE_{i,t}^{BLS}$ be the unemployment rate in state i in the Current Population Survey (CPS) survey in month t. We run the following regression on the pre-crisis data (January 2020)

$$UE_{i,t}^{Nielsen} = b_0 + b_1 UE_{i,t}^{BLS} + error \tag{1}$$

and then use the estimated coefficients to adjust the Nielsen statistic for month s (January or April 2020) as

$$\widetilde{UE}_{i,s}^{Nielsen} = (UE_{i,s}^{Nielsen} - \hat{b}_0)/\hat{b}_1. \tag{2}$$

By construction, the average value of $\widetilde{UE}_{i,s}^{Nielsen}$ is now equal to the average value of $UE_{i,s}^{BLS}$ across states, which may still differ from the aggregate unemployment rate (which is a population-weighted average across states). We apply a similar procedure for employment-to-population ratios and labor force participation rates. The resulting adjusted measures of aggregate unemployment, labor force participation, and employment-to-population ratios pre-crisis from the Nielsen survey are presented in Table 1, column 4. For easier comparison to BLS numbers, we focus primarily on these adjusted measurements (applied to both waves of our survey) from now on.

III Labor Markets since the Covid-19 Crisis

A second survey was run on households participating in the Nielsen Homescan panel between the afternoon of April 2nd-8th, 2020. The response rate was 18.6% with 9,445 responses. The same labor market questions were asked in both waves of the survey, so we can directly compare the two waves to get a sense of how labor markets have evolved since the onset of the covid-19 virus and the associated policy responses. Table 1 summarizes results from the second wave of the survey. Our main results are as follows.

First, the employment-to-population ratio has declined sharply. Using the adjusted metrics described above, we find that the employment ratio fell from 60% of the population down to 52.2%, a nearly 8% point decline (s.e. 0.7%). As illustrated in Panel C of Figure 1, this decline in employment is enormous by historical standards and is larger than the entire decline in the employment to population ratio experienced during the Great Recession. Given that the U.S. civilian non-institutional population is approximately 260 million, this drop in employment to population ratio is equivalent to 20 million people losing their jobs (the 90 percent confidence interval is 16.9 million to 23.1 million). This drop is even larger than the 16.5 million new unemployment claims over this time period.

Covid Economics 21, 22 May 2020: 40-58

Covid Economics 21, 22 May 2020: 40-58

Figure 1. Time series of key employment statistics.

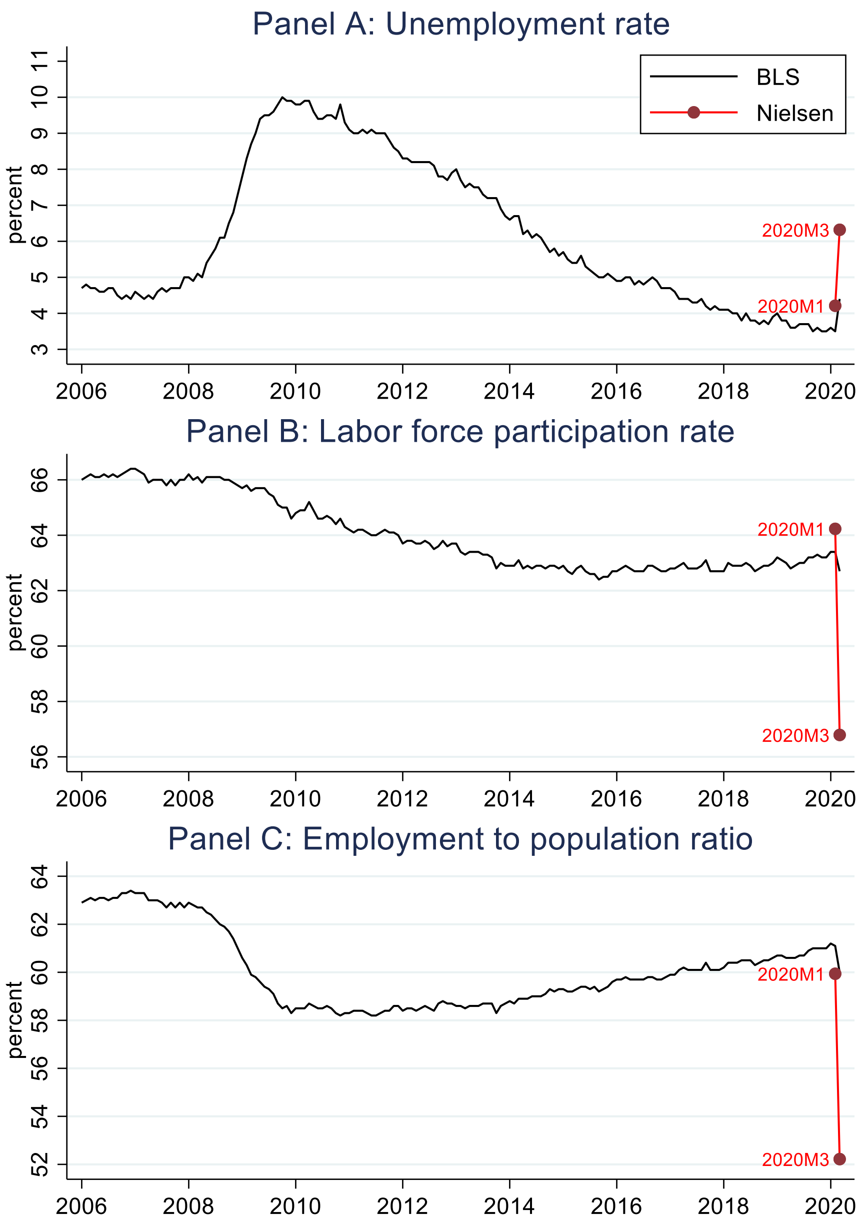

Notes: Each panel plots monthly time series of an employment statistic. The black, solid line shows data from the Bureau of Labor Statistics (BLS). The red, solid line with circles shows the corresponding values from the Nielsen survey. 2020M2 are the values from the Nielsen pre-crisis survey. 2020M3 are the values from the Nielsen crisis wave.

Covid Economics 21, 22 May 2020: 40-58

Second, we find a much smaller increase in the unemployment rate. As documented in Table 1, the adjusted unemployment rate rose from 4.2% to 6.3%. Panel A of Figure 1 plots this rise relative to previous changes in unemployment over the last 15 years. While this increase is the single biggest discrete jump in unemployment over the time period, this change in unemployment corresponds only to about one-third of the increase observed during the Great Recession. For comparison with the employment to population ratio, if all twenty million newly unemployed people as measured by the decline in the employment to population ratio were counted in the unemployment rate, we would have found an increase in the unemployment rate from 4.2% to 16.4%, the highest level since 1939.

The reason for the discrepancy between the two is that many of the newly non-employed people are reporting that they are not actively looking for work, so they do not count as unemployed but rather as exiting the labor force (recall that in contrast to the BLS, we do not automatically count laid-off workers as unemployed). Consistent with this, we find an extraordinary decline in the adjusted labor force participation rate, from 64.2% to 56.8%. For comparison, Panel B of Figure 1 plots the historical evolution of the labor force participation rate over the last 15 years which includes a historically large decline in participation between 2008 and 2016 of 3 percentage points. Even this cumulatively large decline in participation over an eight year period is dwarfed by the historic decline in participation that we document.[3]

How unusual are these patterns? We have already seen that the size of the changes in each variable is exceptional, at least for employment and changes in labor force participation. What about their simultaneous changes? Historically, unemployment and employment are very strongly negatively related. Within short periods, movements in one are reflected almost perfectly in the other as workers move from employment to unemployment and back (Appendix Figure 3). Slow-moving demographics cause the relationship to gradually change over time, as can be seen by decadal shifts in the curve, but short-run movements are close to linear. The change that we document since the covid-19 crisis jumps out: we see an enormous change in the employment-to-population ratio with a much smaller change in unemployment that would have typically been expected. This pattern is therefore qualitatively different from the historical experience of U.S. labor markets, even after taking into account the size of the changes. On the other hand, we find a less unusual pattern relative to historical experience when looking at the change in labor force participation and employment to population ratios (Appendix Figure 3). The two tend to commove positively and closely on average: periods when employment growth is strong are also periods during which more people are entering the labor force. In that sense, the simultaneous decline in employment and labor

[3] We find similar results at the state level (Appendix Figure 2) which plots state-level changes in unemployment, labor force participation and employment to population ratios around the time of the covid-19 crisis. The forty-five degree lines indicate no change. When looking at labor force participation rates and employment-to-population ratios, we see that almost all states fall below the 45 degree line, indicating the declines in each variable are widespread throughout the country and relatively homogenous in size. The change in unemployment across states is noisier, due in part to higher measurement error in measuring unemployment rates at local levels, but also indicates geographically dispersed increases in unemployment. Still, the pattern that comes out from state-level variation is one of broad-based declines in both employment and labor force participation across the country.

CEPR PRESS

force participation is mostly unusual because of the size of the changes. Still, the drop in labor force participation appears large relative to the historical experience given the size of the decline in employment, which is consistent with the smaller than normal increase in unemployment.

Why do so many unemployed choose not to look for work? Both surveys included a question asking those who said they were neither working nor looking for a job to select among possible answers why this was the case. The results for both waves are presented in Table 2. Prior to the crisis, most respondents out of the labor force claimed that it was because they were retired, disabled, homemakers, raising children, students, or did not need to work. Only 1.6% of those out of the labor force were claiming that they could not find a job as one of their reasons for not searching. At the height of the covid-19 crisis with a much larger number of people now out of the labor force, we see corresponding declines in the share of homemakers, those raising children and the disabled. However, we see a large increase in those who claim to be retired, going from 53% to 60%.[4] This makes early retirement a major force in accounting for the decline in the labor-force participation. Given that the age distribution of the two surveys is comparable, this suggests that the onset of the covid-19 crisis led to a wave of earlier than planned retirements. With the high sensitivity of seniors to the covid-19 virus, this may reflect in part a decision to either leave employment earlier than planned due to higher risks of working or a choice to not look for new employment and retire after losing their work in the crisis.

To see this more clearly, we exploit the panel dimension of the survey and identify respondents who were out of the labor force in April 2020 but in the labor force in January 2020. Column (3) of Table 2 presents reasons for being out of the labor force reported by those who were *employed* in January. Of those, 28% report that they are now out of the labor force because of retirement. Similarly, of those who were *unemployed* in January and out of the labor force in April (column 4), 21% report that it is because they retired. And while 9% of those going from employment to out of the labor force between January and April report that they are on break from working (as might be the case for some affected by temporary work closures due to covid-19), the equivalent proportion is 8% for those going from unemployment to out of the labor force, suggesting that these breaks are voluntary ones, not ones due to temporary work closings. In short, these results point to an unusual rise in the share of retirements accounting for the exceptional decline in labor force participation during this time period.

IV Distribution and Drivers of Retirement

To better understand which parts of the age distribution might drive the increase of retirees in our survey and whether economic incentives at least partially play a role, we plot in Panel A of Figure 2 the fraction of those claiming being retired (left scale) both in the pre-crisis wave (black, dashed) and in the crisis wave (red, solid) together with the difference between the two (blue, right scale). The crisis has shifted the whole distribution up, that is, for each part of the age distribution a larger fraction of the survey population now claims being retired.

[4] Coile and Levine (2007) document that early retirement increased in previous recessions but the magnitudes are much larger in the current crisis.

Covid Economics 21, 22 May 2020: 40-58

Table 2. Reasons for not looking for a job (for those who do not have a job).

Reason	Share of people choosing a reason			
	All people		Employed pre-crisis, out of labor force in crisis	Unemployed pre-crisis, out of labor force in crisis
	Pre-crisis	Crisis		
	(1)	(2)	(3)	(4)
Homemaker	0.204	0.155	0.158	0.244
Raising children	0.125	0.082	0.115	0.222
Student	0.023	0.018	0.050	0.019
Retiree	0.527	0.595	0.277	0.210
Disabled, health issues	0.297	0.261	0.142	0.144
Couldn't find a job	0.016	0.018	0.061	0.133
On break	0.012	0.011	0.092	0.080
No financial need	0.049	0.037	0.039	0.029
None of the above	0.026	0.034	0.352	0.184

Notes: the Nielsen survey question is "Here are a number of possible reasons why people who are not working choose not to look for work. **Please select all that apply to you**." This question is asked for people who do not have a job and are not looking for a job.

Covid Economics 21, 22 May 2020: 40-58

Figure 2. Non-employment status by age
Panel A: Retirement by age (3-year moving average).

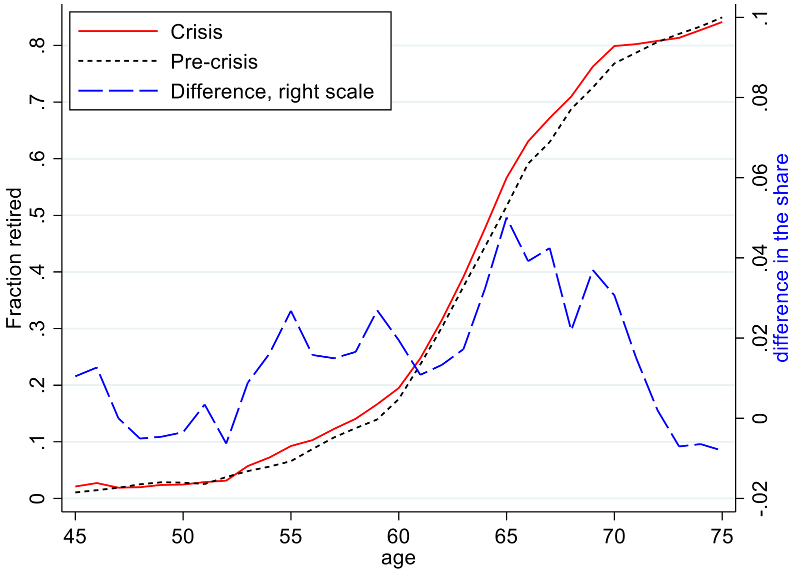

Panel B: Disability by age (3-year moving average).

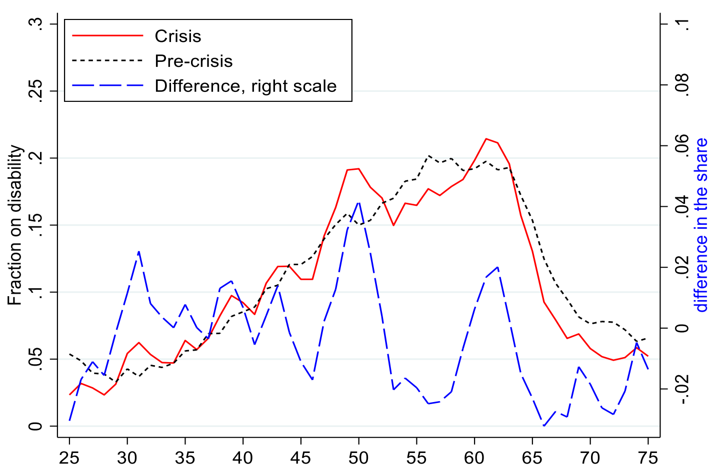

Notes: each panel plots fractions of population claiming retirement (Panel A) or disability (Panel B) by age as a reason of non-employment in the Nielsen survey. Red, solid line shows the distribution in the pre-crisis wave of the survey. Black, dashed line shows the distribution in the crisis (April) wave of the survey. The blue, long-dash line shows the difference in the distributions (right scale)

Covid Economics 21, 22 May 2020: 40-58

Covid Economics 21, 22 May 2020: 40-58

Table 3. Retirement by demographic characteristics

Demographic characteristic	Demographic group	Pre-crisis	Crisis	Difference
All	All	0.25	0.29	0.04
		(0.01)	(0.01)	(0.01)
gender	Male	0.29	0.31	0.02
		(0.01)	(0.02)	(0.02)
	Female	0.23	0.28	0.05
		(0.01)	(0.01)	(0.01)
race	White	0.25	0.29	0.04
		(0.01)	(0.01)	(0.01)
	Black	0.24	0.32	0.08
		(0.02)	(0.03)	(0.04)
	Asian	0.35	0.26	-0.09
		(0.06)	(0.06)	(0.08)
	Other	0.26	0.26	-0.00
		(0.03)	(0.04)	(0.06)
Housing situation	Own our house/apartment without a mortgage	0.34	0.34	0.01
		(0.01)	(0.01)	(0.02)
	Own our house/apt. and have a fixed-rate mortgage	0.22	0.26	0.04
		(0.01)	(0.01)	(0.02)
	Own our house/apt. and have a variable-rate mortgage	0.15	0.23	0.08
		(0.04)	(0.06)	(0.07)
	Rent our house/apartment	0.20	0.24	0.04
		(0.02)	(0.02)	(0.02)
	Other	0.14	0.28	0.13
		(0.03)	(0.05)	(0.06)
Household income	Under $5000	0.24	0.26	0.02
		(0.06)	(0.08)	(0.10)
	$5000-$7999	0.24	0.31	0.07
		(0.08)	(0.09)	(0.12)
	$8000-$9999	0.24	0.21	-0.03
		(0.06)	(0.06)	(0.09)
	$10,000-$11,999	0.24	0.26	0.02
		(0.05)	(0.06)	(0.08)
	$12,000-$14,999	0.27	0.29	0.02
		(0.04)	(0.05)	(0.07)
	$15,000-$19,999	0.31	0.35	0.04
		(0.04)	(0.04)	(0.06)
	$20,000-$24,999	0.31	0.31	0.00
		(0.03)	(0.04)	(0.05)
	$25,000-$29,999	0.33	0.36	0.03
		(0.04)	(0.04)	(0.05)
	$30,000-$34,999	0.28	0.32	0.03
		(0.03)	(0.03)	(0.05)
	$35,000-$39,999	0.30	0.30	-0.01
		(0.03)	(0.04)	(0.05)
	$40,000-$44,999	0.26	0.29	0.02
		(0.04)	(0.04)	(0.06)
	$45,000-$49,999	0.25	0.26	0.01
		(0.03)	(0.04)	(0.05)
	$50,000-$59,999	0.30	0.35	0.05
		(0.03)	(0.03)	(0.04)
	$60,000-$69,999	0.21	0.24	0.04
		(0.02)	(0.03)	(0.04)
	$70,000-$99,999	0.25	0.29	0.04
		(0.02)	(0.02)	(0.03)
	$100,000 +	0.22	0.25	0.03
		(0.01)	(0.02)	(0.02)

(continued on the next page)

Demographic characteristic	Demographic group	Pre-crisis	Crisis	Difference
Male head education	No head	0.22	0.27	0.05
		(0.01)	(0.02)	(0.02)
	Grade School	0.14	0.13	-0.02
		(0.08)	(0.12)	(0.15)
	Some High School	0.19	0.22	0.03
		(0.04)	(0.05)	(0.06)
	Graduated High School	0.27	0.28	0.01
		(0.02)	(0.02)	(0.03)
	Some College	0.30	0.30	-0.00
		(0.02)	(0.02)	(0.03)
	Graduated College	0.27	0.34	0.07
		(0.02)	(0.02)	(0.03)
	Post College Grad	0.24	0.29	0.05
		(0.03)	(0.03)	(0.04)
Female head education	No head	0.28	0.28	0.00
		(0.02)	(0.02)	(0.03)
	Grade School	0.29	0.29	-0.01
		(0.14)	(0.18)	(0.23)
	Some High School	0.26	0.34	0.09
		(0.05)	(0.08)	(0.10)
	Graduated High School	0.27	0.28	0.01
		(0.01)	(0.02)	(0.02)
	Some College	0.25	0.29	0.04
		(0.01)	(0.02)	(0.02)
	Graduated College	0.22	0.29	0.06
		(0.02)	(0.02)	(0.02)
	Post College Grad	0.21	0.33	0.12
		(0.02)	(0.03)	(0.04)

Notes: The table reports levels of retirement by survey wave (crisis and pre-crisis) as well as the difference between the waves. The econometric specification is $Retired_{it} = b_0 * 1\{precrisis\ wave\} + b_1 * 1\{crisis\ wave\} + error$. The sample is restricted to ages [55,66].

Hence, even for those that are well before retirement age, we see a large increase in early retirement. Moreover, a notable jump in the difference occurs at age 66 which is the first year people can claim retirement benefits without penalty from the social security administration (SSA).

Panel B follows the same structure but for the fraction of survey participants answering that they are not looking for a job because of disability. Here, the pre- and post-crisis lines do not show a general pattern with both lines crossing several times but we also see large spikes in the fraction being disabled during the crisis wave at ages 50 and 60 which are cut-off ages at which it becomes easier for individuals to claim disability benefits. Taken together, these results suggest that economic benefits could drive the decision to leave the labor force by retiring or claiming disability benefits. These results could also indicate that unemployment benefit extensions or increases in other welfare programs could decrease labor search intensities which would negatively affect a later recovery.

In Table 3, we report levels and changes in retirement status for ages 55 to 66 in various demographic groups. Unconditionally, we see a 4 percentage point increase in the fraction being retired which is predominantly driven by women in our survey who see a 5 percentage points increase. These results indicate a possible negative effect of the crisis on female labor force participation and possibly gender inequality. In addition, we see a large

Covid Economics 21, 22 May 2020: 40-58

increase in the fraction of blacks entering retirement. Homeowners with mortgages and renters are also more likely to retire relative to homeowners without mortgage. As for education, we typically find that the most educated enter retirement, both for women and men. Finally, we find general increases in the propensity of being retired for the whole income distribution with notable increases at the bottom and top of the income distribution.

Taken together, these results indicate that two different parts of the population drive the increase in retirement. On the one hand, the more vulnerable parts of the population, women and low-income individuals leave the labor force and enter retirement. On the other hand, highly educated and high income survey participants are also more likely to have moved into retirement during the crisis wave relative to the pre-crisis wave.

V Summary

It is still very early on in the covid-19 crisis, but preliminary indicators point toward catastrophic declines in employment. Our surveys provide additional evidence on this decline in employment, pointing to a 20 million decline in the number of employed workers. Most strikingly, we find a much less than proportional increase in unemployment, indicating that most of these newly unemployed workers are not looking for new work. Hopefully this reflects a transitory characteristic as these individuals face shelters-at-home and few work opportunities. But the wave of early retirements that we document suggests that more permanent changes may already be taking place.

References

Adams-Prassl, Abi, Teodora Boneva, Marta Golin, and Christopher Rauh, 2020. "Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys," Working Paper.

Allcott, Hunt, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Yang, 2020. "Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic," Working Paper.

Alvarez, Fernando, David Argente, and Francesco Lippi (2020) "A Simple Planning Problem for COVID-19 Lockdown", *Covid Economics – Real-Time and Vetted Papers* (14):1-32.

Atkeson, Andrew (2020) "What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios" NBER Working Paper w26867.

Baek, ChaeWon, Peter McCrory, Todd Messer and Preston Mui, 2020. "Unemployment Effects of Stay-at-Home Orders: Evidence from High Frequency Claims Data," Working Paper.

Barrios, John and Yael Hochberg, 2020. "Risk Perception Through the Lens of Politics in the Time of the COVID-19 Pandemic," Working Paper.

Barro, Robert J., José F. Ursúa, and Joanna Wang (2020) "The Coronavirus and the Great Influenza Pandemic: Lessons from the 'Spanish Flu' for the Coronavirus's Potential Effects on Mortality and Economic Activity" NBER Working Paper w26866.

Covid Economics 21, 22 May 2020: 40-58

Beland, Louis-Philippe, Abel Brodeur, and Taylor Wright, 2020. "The Short-Term Economic Consequences of COVID-19: Exposure to Disease, Remote Work and Government Response," Working Paper.

Bick, Alexander and Adam Blandin, 2020. "Real Time Labor Market Estimates During the 2020 Coronavirus Outbreak," Working Paper.

Binder, Carola, 2020. "Coronavirus Fears and Macroeconomic Expectations" Working Paper.

Coile, Courtney C. , and Phillip B. Levine, 2007. "Labor market shocks and retirement: Do government programs matter?" *Journal of Public Economics* 91(10): 1902-1919.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko and Michael Weber, 2020a. "Forward Guidance and Household Expectations," NBER Working Paper w26778.

Coibion, Olivier, Yuriy Gorodnichenko and Michael Weber, 2019. "Monetary Policy Communications and their Effects on Household Inflation Expectations," NBER Working Paper w25482.

D'Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber, 2020a. "Exposure to Daily Price Changes and Inflation Expectations," NBER Working Paper w26237.

D'Acunto, Francesco, Ulrike Malmendier, and Michael Weber, 2020b. "Gender Roles and the Gender Expectations Gap," NBER Working Paper w26837.

Dietrich, Alexander M., Keith Kuester, Gernot J. Müller, and Raphael S. Schoenle (2020) "News and Uncertainty about COVID-19: Survey Evidence and Short-Run Economic Impact," Working Paper.

Foote, Christopher, Alan Gerber, William Nordhaus, and Douglas Rivers, 2020. "Work in the Time of COVID: Rapid Estimates of the US Employment Situation Using the CovidCPS Survey," Working Paper.

von Gaudecker, Hans-Martin, Radost Holler, Lena Janys, Bettina Siflinger, and Christian Zimpelmann, 2020. "Labour supply in the early stages of the CoViD-19 Pandemic: Empirical Evidence on hours, home office, and expectations," Working Paper.

Mongey, Simon and Alex Weinberg, 2020. "Characteristics of Workers in Low Work-From-Home and High Personal-Proximity Occupations," Becker Friedman Institute for Economic White Paper.

Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt (2020) "The Macroeconomics of Epidemics" NBER Working Paper w26882.

Fetzer, Thiemo, Lukas Hensel, Johannes Hermle, and Christopher Roth, 2020. "Perceptions of Contagiousness and Mortality of COVID-19 Shape Economic Sentiment," Working Paper.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Ivan Werning, 2020. "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" NBER Working Paper w26918.

Hanspal, Tobin, Annika Weber, and Johannes Wohlfart, 2020. "Income and Wealth Shocks and Expectations during the Covid-19 Pandemic" Working Paper.

Covid Economics 21, 22 May 2020: 40-58

Online Appendix

(not for publication)

Covid Economics 21, 22 May 2020: 40-58

Appendix Figure 1: Pre-crisis employment statistics at the state level.

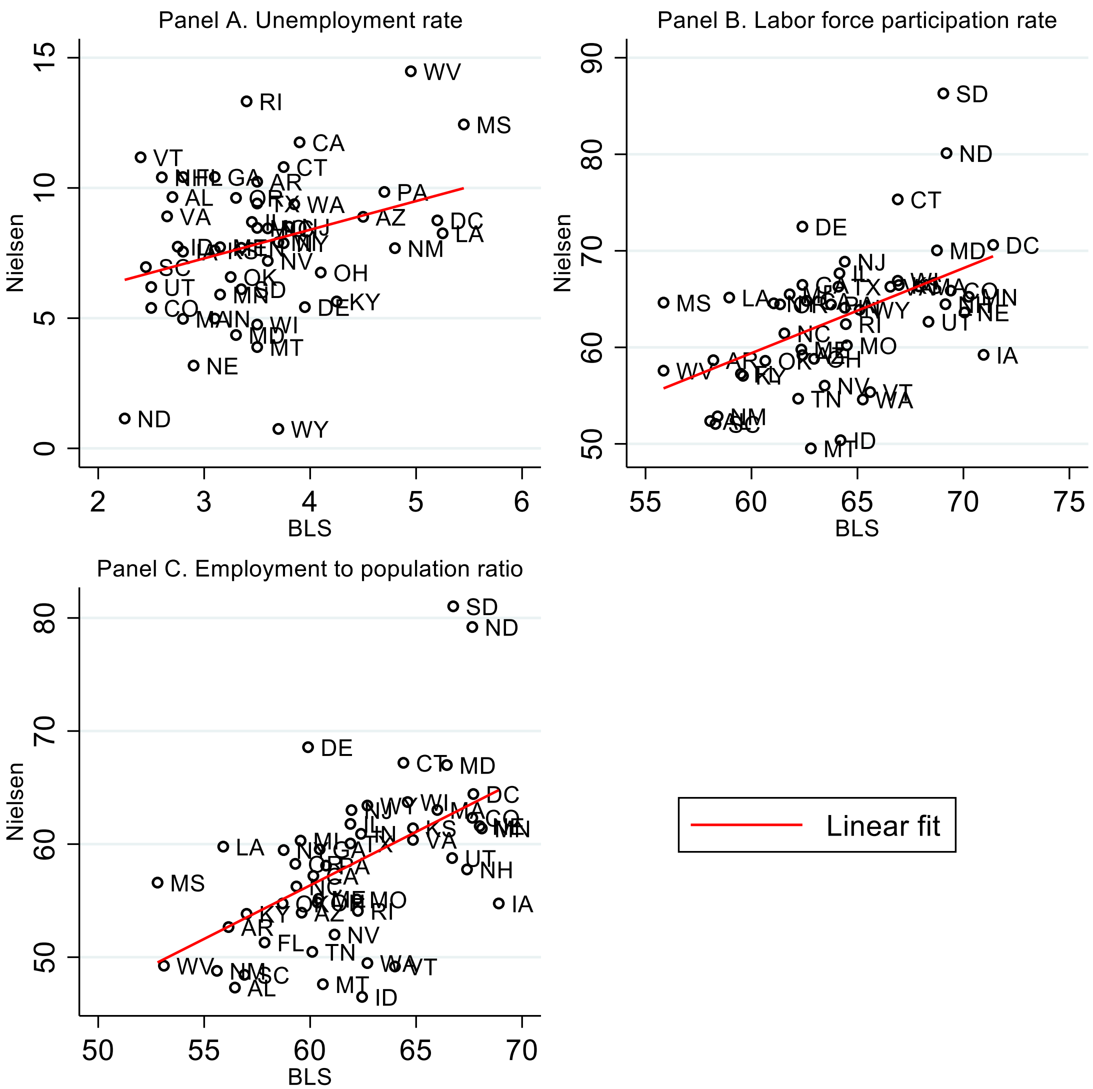

Notes: Bureau of Labor Statistics (BLS) data refer to the January-February 2020 period. Nielsen data are for the pre-crisis wave of the survey.

Appendix Figure 2. Adjusted employment statistics by state, pre-crisis vs. crisis levels, Nielsen survey.

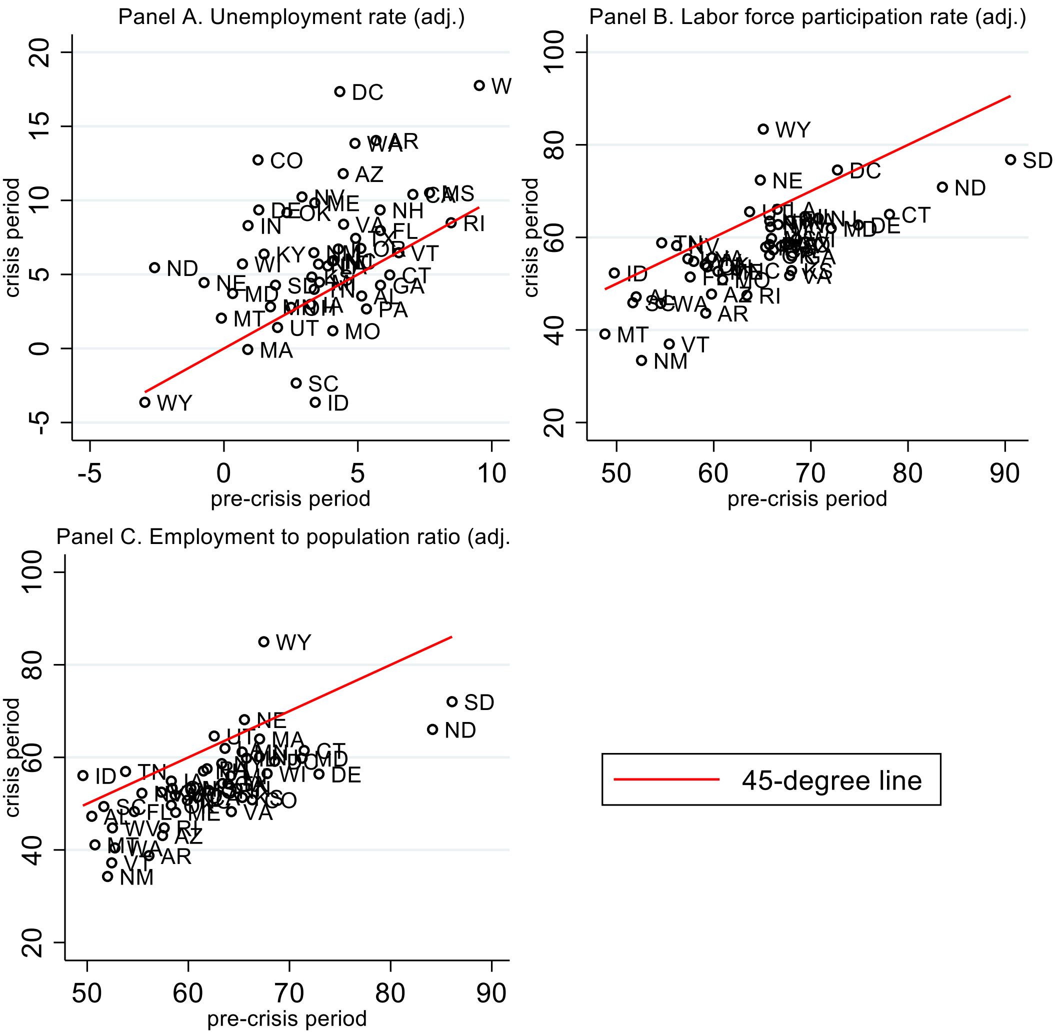

Notes: Nielsen employment statistics are adjusted to match average pre-crisis levels observed in the official data compiled by the Bureau of Labor Statistics (BLS). The adjustment is described in equations (1)-(2).

Covid Economics 21, 22 May 2020: 40-58

Appendix Figure 3. Historical comovement of key employment statistics.
Panel A: Employment-to-population ratio vs Unemployment rate.

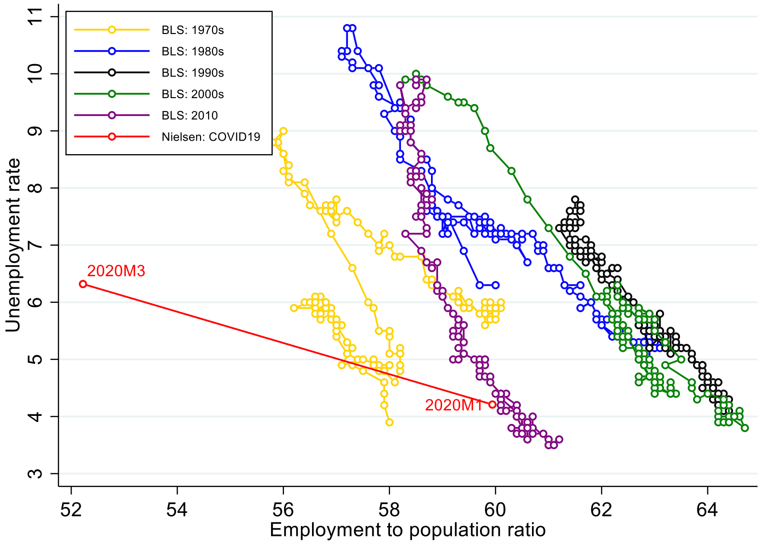

Panel B: Employment-to-population ratio vs Labor force participation rate.

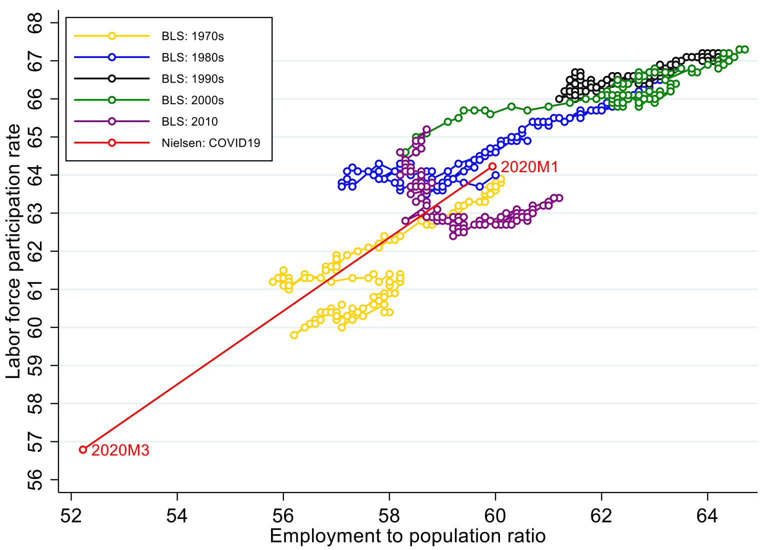

Notes: Each panel shows comovement (by decade) of key official employment statistics compiled by the Bureau of Labor Statistics (BLS) as well as employment statistics based on Nielsen surveys (red circles with dates). 2020M2 are the values from the Nielsen pre-crisis survey. 2020M3 are the values from the Nielsen crisis wave.

CEPR PRESS

Covid Economics **Issue 21, 22 May 2020**

International trade of essential goods during a pandemic[1]

Fernando Leibovici[2] and Ana Maria Santacreu[3]

Date submitted: 14 May 2020; **Date accepted:** 19 May 2020

This paper studies the role of international trade of essential goods during a pandemic. We consider a multi-country, multi-sector model with essential and non-essential goods. Essential goods provide utility relative to a reference consumption level, and a pandemic consists of an increase in this reference level. Each country produces domestic varieties of both types of goods using capital and labor subject to sectoral adjustment costs, and all varieties are traded internationally subject to trade barriers. We study the role of international trade of essential goods in mitigating or amplifying the impact of a pandemic. We find that the effects depend crucially on the countries' trade imbalances in essential goods. Net importers of these goods are relatively worse off during a pandemic than net exporters. The welfare losses of net importers are lower in a world with high trade barriers, while the reverse is the case for net exporters. Yet, once a pandemic arrives, net exporters of essential goods benefit from an increase in trade barriers, while net importers benefit from a decrease in them. These findings are consistent with preliminary evidence on changes in trade barriers across countries during the COVID-19 pandemic.

Covid Economics 21, 22 May 2020: 59-99

1 We thank Matthew Famiglietti and Makenzie Peake for excellent research assistance. The views expressed in this paper are those of the individual 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. First version: May 5, 2020. This version: May 20, 2020. Download the latest version of the paper from our websites: www.fernandoleibovici.com and www.anamariasantacreu.com.

2 Economist, Federal Reserve Bank of St. Louis.
3 Senior Economist, Federal Reserve Bank of St. Louis.

Covid Economics 21, 22 May 2020: 59-99

1 Introduction

The ongoing COVID-19 pandemic has led to a massive increase in the demand for essential medical equipment to combat it. For instance, goods such as masks, gowns, gloves, and respirators, among others, are playing a key role in allowing healthcare workers to address the ongoing pandemic. While supply has been gradually increasing to try to satisfy the high demand for these goods, countries are starting to face supply shortages and being increasingly forced to ration these goods.

Amidst the growing fear of supply shortages on these essential goods, countries have been resorting to trade policy. According to Evenett and Winters (2020), "as they scramble to find medical supplies to tackle COVID-19, some countries are eliminating their restrictions in imports while others are curtailing their exports."[1] That is, while some countries are lowering their import trade barriers to ease access to these goods, others are making it harder for their domestic firms to sell these goods internationally. Indeed, we document that the heterogeneous response of trade policy across countries is systematically related to the extent to which countries are net importers or net exporters of essential medical goods. According to data collected by Global Trade Alert as of April 24, 2020, 86% of the countries with a trade surplus in these goods have recently imposed restrictive export policies, while only 46% of the countries with a trade deficit have done so.

These findings suggest that international trade plays a fundamental role in the cross-country impact of a pandemic. In this paper, we investigate the role of international trade of essential goods in mitigating or amplifying this impact. In particular, to what extent is the impact of a pandemic heterogeneous across countries depending on whether they are net exporters or net importers of these essential goods? And what is the impact of alternative trade barriers across these countries in the short vs. long run? We address these questions using a dynamic quantitative general equilibrium model with multiple countries and multiple sectors.

While the paper is certainly motivated by and applied to the ongoing COVID-19 pandemic, the implications of our analysis extend well beyond the specific case of trade in essential medical goods during a pandemic. Our approach allows us to investigate the role of trade in any type of essential good that might be subject to shocks. For instance, to the

[1] See Baldwin and Evenett (2020) for a detailed analysis of trade policy changes during COVID-19.

extent that agricultural goods are essential for the survival of a country's population, our analysis can be extended to examine the impact of a pest that might destroy a country's agricultural production.

We thus begin the paper by investigating the patterns of trade imbalances and trade policy across countries in a range of essential goods broader than those directly necessary to address the COVID-19 pandemic. We consider essential goods as ones that are hard to substitute intertemporally. In particular, we classify food, defense, and medical goods as essential and all other traded goods as non-essential. We then ask: To what extent do trade patterns of and trade policy on these goods vary systematically relative to those for non-essential goods? We document that, in the average country, trade imbalances in food and defense are smaller in absolute value than those among non-essential goods. Moreover, both tariff and non-tariff barriers are systematically higher, on average, among food and defense goods relative to those on non-essential goods. In contrast, the average country runs a systematically larger trade deficit and imposes considerably lower tariffs on medical goods than on non-essential goods.

This evidence shows that international trade patterns and barriers among essential goods are systematically different from those among non-essential goods. Moreover, trade patterns and barriers among goods traditionally considered to be essential such as food and defense are also different from those among medical goods. The latter are proving to be essential during a pandemic but might have not been previously perceived as key from the perspective of international trade policy.

Motivated by the observed cross-sectoral heterogeneity, we set up a model of international trade that allows us to investigate the role of trade of essential goods in the economic impact of a shock to either the demand or supply of these goods. Each country in our model produces domestic varieties of both essential and non-essential goods using capital and labor. All these goods are traded internationally, since both countries consume domestic and imported varieties of essential and non-essential goods. Capital and labor can be reallocated across sectors in response to shocks, but this adjustment is subject to costs. Moreover, we assume that firms are myopic and do not internalize the impact of their production decisions on the welfare of households; in our application, this assumption allows us to capture that firms might not adjust their decisions for the possibility of a pandemic.

We assume that the fundamental difference between essential and non-essential goods is

Covid Economics 21, 22 May 2020: 59-99

Covid Economics 21, 22 May 2020: 59-99

accounted for by differences in household preferences for these goods. While non-essential goods are valued according to a standard logarithmic utility function, we assume that essential goods are valued according to a novel utility function in the spirit of Stone (1954) and Geary (1950). In particular, we assume that utility from essential goods is derived from the consumption of these goods relative to a time-varying *reference level*.[2] While analogous to the subsistence level featured by Stone-Geary preferences, we allow households to consume below the reference level but at an exponentially increasing utility cost. Moreover, the utility obtained from the consumption of essential goods is bounded above, allowing us to capture the limited potential to increase utility from ever increasing consumption of these goods.

We use this model as a laboratory to investigate the role of international trade on the cross-country impact of a global pandemic. Thus, throughout our analysis we interpret essential goods as consisting solely of essential medical goods. We then model a global pandemic as a substantial increase in the reference level relative to the level at which households in both countries derive utility from their consumption of essential goods. Our analysis therefore abstracts from most multifaceted effects of a pandemic (e.g., social distancing policies, infection and death rates, increased unemployment), solely restricting attention to the increased demand for essential goods.

In this preliminary version of the paper, we conduct a simple numerical exercise to illustrate the various mechanisms at play during a pandemic. In particular, we assume that both countries are symmetric throughout except for one dimension: we assume that one country is relatively more productive in essential goods, while the other country is relatively more productive in non-essential goods, leading the former to be a net exporter and the latter to be a net importer of these goods. Our goal in future versions of the paper is to quantify the role of these mechanisms by disciplining the model to match salient features of cross-country data.

We find that a global pandemic in which the reference consumption level of essential goods increases identically in both countries has significantly heterogeneous effects across them. There are two main forces at play. First, given that the production of essential goods is subject to capital and labor adjustment costs, producers of these goods cannot sufficiently reallocate production inputs across sectors to boost production, leading to a substantial

[2]One interpretation of this reference level in the context of essential medical goods is the level of healthcare consumption recommended by healthcare professionals. A pandemic such as COVID-19 can then be interpreted as an increase in this reference level.

increase in their relative price. Then, while the price of essential goods increases to the same extent in both economies, the effect on welfare depends on the countries' sectoral trade imbalances. Both countries earn relatively more for their sales of essential goods, and both countries have to spend relatively more for their purchases of these goods. However, given that the net importer is less productive in essential goods, this country sells relatively less than it purchases, while the reverse is the case for net exporters. As a result, net importers of these goods are relatively worse off during a pandemic than net exporters.[3]

We then conduct an extensive analysis to isolate the forces underlying these findings. We first find that capital and labor adjustment costs play a key role in generating a shortage of essential goods that raises their price, yielding a heterogeneous impact across countries. We then show that our preference specification for essential goods is key to generating a substantial increase in the demand for these goods, hence increasing their price substantially. We also document the importance of ex-ante trade imbalances in essential goods; in a world economy with balanced trade of essential goods, the welfare implications are significantly more muted. We conclude this analysis by showing that the assumption that firms are myopic does not play a key role in accounting for our findings in the current parametrization of the model.

We then investigate the role of international trade policy in mitigating or exacerbating the impact of a pandemic. We find that the level of international trade barriers at the onset of a pandemic can significantly alter the economic implications. Keeping all other parameters unchanged, higher initial trade costs imply that the net importer runs a smaller trade deficit of essential goods in the steady state, while the net exporter runs a smaller trade surplus in these goods. Therefore, while a pandemic continues to lead to a substantial increase in the relative price of essential goods, the relative impact on the home and foreign countries is mitigated. We thus conclude that while higher trade barriers on essential goods may reduce the amount of these goods consumed in the steady state, they mitigate the potential vulnerability of net importers of these goods when a global pandemic hits.

These findings have important implications for the design of international trade policy. Even if countries may benefit on average from having cheaper access to goods, the reliance on international trade might put these countries in a vulnerable position if these goods are

[3]The role of sectoral trade imbalances as an amplification mechanism of shocks to relative prices is akin to their role played in Kohn, Leibovici, and Tretvoll (forthcoming) to account for business cycle volatility differences between developed and emerging economies.

Covid Economics 21, 22 May 2020: 59-99

essential and either the supply or demand of these goods is subject to global shocks.

While the previous question is informative about the ex-ante design of international trade policy, it is silent about the impact of trade policy changes implemented once a pandemic hits. We thus conclude our analysis by asking: What is the impact of raising trade barriers in response to a pandemic? In contrast to the previous findings, we now find that raising trade barriers when a pandemic hits can significantly exacerbate the welfare costs for net importers of essential goods while making net exporters of these goods relatively better off. As a net importer of essential goods, the home country's production structure is such that it relies considerably on the foreign country for its consumption of essential goods. Thus, an increase in trade barriers on essential goods exacerbates net importers' already difficult task of meeting the demand for these goods and also makes the relative price of these essential goods even higher.

These results suggest that the design of international trade policy for essential goods may suffer from a time-consistency problem. Net exporters of essential goods may ex-ante prefer to live in a world with low trade costs; but when a global pandemic hits, they might be tempted to renege on their commitments and increase trade barriers. The reverse is the case for net importers of essential goods: they may ex-ante prefer to live in a world with high trade costs on essential goods; but when a global pandemic hits, they might be tempted to renege on their commitments and decrease trade barriers.

Our findings also raise the question of whether countries would benefit from protecting essential sectors such as those that produce defense, food, and medical goods. While protectionism may decrease welfare on average, it might be particularly beneficial to mitigate the negative consequences of relying on other countries for the supply of these essential goods. That is, the argument for the protection of essential sectors might be that free trade reduces the ability of countries to produce those goods in the presence of a national emergency such as a war, a natural disaster or, as is the case today, a global health shock. In general, countries tend to specialize in those goods in which they have comparative advantage. During a global disaster in which a country becomes isolated from its main suppliers, its only option is to start producing goods it would have otherwise imported. It may not be feasible, however, to do so if it does not have some installed capacity already in place to increase domestic production of those goods.

This question has been indeed a recurrent topic of debate in international trade policy

Covid Economics 21, 22 May 2020: 59-99

Covid Economics 21, 22 May 2020: 59-99

and, more specifically, in World Trade Organization negotiations. For instance, the importance of agriculture in global trade led to specific agreements through which governments are allowed to support domestic production via domestic or export subsidies. Moreover, goods such as sugar and steel have been subject to protectionist measures. Most recently, the United States has imposed tariffs on steel and aluminum, alluding to national security reasons.

Our paper raises questions related to a very recent and growing literature on the role of international trade of essential medical equipment during a pandemic. In particular, the shortage of essential medical equipment in countries that depend on imports of these goods during the ongoing COVID-19 pandemic has led to a surge of papers proposing international coordination mechanisms to avoid future shortages. Many of these very recent papers are contributions to the recent VoxEU eBook on COVID-19 and trade policy (Baldwin and Evenett 2020). For instance, Stellinger, Berglund, and Isakson (2020) and others advocate for the role of international trade as an insurance device for those countries that rely on imports of essential goods, and hence are specialized in the production of other types of goods.

Several authors have also emphasized the need for greater international coordination. In particular, Evenett (2020) has proposed an international agreement with export incentives to the main suppliers of essential goods and low import tariffs by the main importers. These arguments were also used in a less recent literature on the role of alliances. A very influential paper on this issue is Olson and Zeckhauser (1966). They develop a model of international cooperation on the basis that countries want to participate in these organizations to provide a public good. The model is applied to the case of NATO, hence focusing on defense as a global public good.

Our modeling framework builds on a large literature that studies international business cycles. In particular, we combine the multi-country framework of Backus, Kehoe, and Kydland (1992) with the multi-sector setups of Mendoza (1995), Schmitt-Grohé and Uribe (2018), and Kohn, Leibovici, and Tretvoll (forthcoming). As in the latter, our model features sectoral adjustment costs in both labor and capital. We extend these setups to model essential goods as different from non-essential goods via differences in their contribution to household utility.

The paper is also related to a recent trade literature studying the short-run and long-run effects of international trade on welfare. In particular, Ravikumar, Santacreu, and Sposi (2019) develop a multi-country two-sector model with endogenous capital accumulation and trade imbalances to evaluate dynamic gains from trade. In contrast to their work, we focus on the implications of a transitory unanticipated shock.

2 Evidence on international trade of essential goods

In this section we investigate the patterns of trade imbalances and trade policy on essential goods across countries. We first examine the extent to which these cross-sectional patterns differ systematically between essential and non-essential goods. Then, we examine the relationship between trade imbalances and trade policy changes on essential medical goods in response to the COVID-19 pandemic.

2.1 Data

International trade flows We collect product-level data on international trade flows at the HS 6-digit level of disaggregation from UN COMTRADE (United Nations Statistics Division 2020). We restrict attention to countries that have a population above one million, which leaves us with a total of 109 countries; these countries account for 96% of world trade and 97% of world GDP.[4] We collect data on both exports and imports of these countries vis-à-vis the rest of the world for the year 2018. Data are in current US Dollars.

We classify the 5,203 products available at this level of disaggregation into four product categories: three essential good categories (defense, food, and medical) and an aggregate of non-essential goods. Defense consists of 20 product codes that include vessels and warships, military weapons, and ammunition, among others. Food consists of 847 product codes that include agriculture and food processing. Medical consists of 71 product codes that include pharmaceuticals, antibiotics, and personal protective equipment. The non-essential sector contains the remaining 4,265 HS 6-digit product codes. Hence, 82% of the products traded are classified as non-essential, while 18% are classified as essential.[5]

[4]Population and GDP data are from the World Bank World Development Indicators (https://datacatalog.worldbank.org/dataset/world-development-indicators).

[5]A complete list of the countries in the dataset and the codes used to classify goods across categories is available for downloaded at `http://github.com/LeiboviciSantacreu/EssentialGoods`.

Covid Economics 21, 22 May 2020: 59-99

We focus on two key statistics to document the cross-country pattern of international trade in these goods. For each country i and product category j, we first compute the share of country i's aggregate exports accounted for by product category j, x_i^j as:

$$x_i^j = \frac{X_i^j}{X_i},$$

where X_i^j denotes the value of exports of product j by country i and X_i denotes country i's aggregate exports. We refer to this statistic as the "export share" of these goods. We compute the analogous statistic for imports and refer to it as the "import share" of these goods.

We then compute each country i's trade imbalance in product category j as the ratio between net exports (exports minus imports) and the total amount of trade (exports plus imports) in these goods:

$$\frac{\mathrm{NX}_i^j}{\text{Total trade}_i^j} = \frac{X_i^j - M_i^j}{X_i^j + M_i^j},$$

where M_i^j denotes imports of product j by country i. We refer to this statistic as the "trade imbalance" in these goods. Limitations on the availability of disaggregated output data across countries prevent us from computing trade imbalances relative to output rather than total trade.

Tariffs We obtain data on tariffs from UNCTAD's Trade Analysis Information System (United Nations Conference on Trade and Development 2020). Data are reported at the HS 6-digit product level. We contrast tariffs across countries and product categories, restricting attention to effectively applied tariffs in the year 2018. For each country, we compute tariffs applied by a given country on imported goods in a given product code as the average tariffs on these goods across all source countries. We obtain data on applied tariffs for 93 of the 109 countries we use to document the patterns of international trade flows.

Non-tariff barriers We obtain data on non-tariff barriers from the Non-Tariff Measure database also collected by UNCTAD's Trade Analysis Information System (United Nations Conference on Trade and Development 2020). This database contains information on a range of trade policy instruments, from traditional trade policy instruments, such as quotas and

Covid Economics 21, 22 May 2020: 59-99

Covid Economics 21, 22 May 2020: 59-99

price controls, to regulatory and technical measures related to health and environmental protection. In particular, non-tariff barriers are classified into 16 different categories, which refer to different types of measures (import licensing, quotas, price controls, export restrictions, etc.). For each reporter country, partner, and HS 6-digit product and category, the dataset reports the number of barriers imposed.[6]

The data spans the period from 2012 to 2018, but not all countries report in every year. We keep the latest year for which a country reports implementing a non-tariff measure to a partner. We find that 59 of the 109 countries we focus on report applying non-tariff measures in this period.

We follow UNCTAD and The World Bank (2018) in using these data to compute two summary statistics on the prevalence of non-tariff barriers across countries and product categories. First, we compute a *frequency index* for each country and product category: the share of traded product lines subject to at least one non-tariff barrier. This statistic allows us to capture the share of traded goods subject to non-tariff barriers regardless of their values. Then, we capture the share of traded value subject to non-tariff barriers by computing a *coverage ratio* for each country and product category: the share of a country's total trade (i.e., the sum of exports and imports) of these goods subject to at least one non-tariff barrier.

2.2 Trade imbalances and trade policy on essential goods

Patterns of trade in essential goods We begin our empirical analysis by contrasting the patterns of international trade flows between essential and non-essential goods. To do so, Table 1 reports the average export and import shares as well as the average trade imbalances across countries for each of the four product categories described above.

We observe that trade of non-essential goods accounts for the vast majority of both exports and imports in the average country. In particular, the average shares of exports and imports of non-essential goods are 78% and 83%, respectively. Among essential goods, food is the most traded, accounting for 16% of exports and 10% if imports, respectively. It is followed by trade of medical goods with 3% and 4%, respectively, and then defense.

We document that while the size of trade imbalances is also heterogeneous across product categories, it appears to vary systematically between essential and non-essential goods. In

[6]One salient limitation of this data is that it only provides information on the number of barriers per country and product code, but not on the intensity of these barriers.

Table 1: Pattern of trade in essential goods

	Export share (%)	Import share (%)	Trade imbalance (%)
Defense	0.07	0.10	-13.88
Food	16.31	10.09	-5.96
Medical	3.24	3.95	-33.22
Non-essential	77.54	82.84	-14.93

Note: This table reports averages across 109 countries in the year 2018. Sectoral imbalances are computed as net exports relative to total trade. Average trade imbalances are computed trimming the bottom 4% of observations.

Covid Economics 21, 22 May 2020: 59-99

particular, we find that trade is relatively more balanced on average in traditional essential goods (food and defense) than in non-essential goods. That is, while all imbalances are negative on average, non-essential goods feature larger trade deficits on average than essential goods.[7]

An exception is observed for trade of medical goods, which features a larger trade imbalance than both non-essential and traditional essential goods. That is, the average country relies heavily on imports of medical goods, with only 17 countries in our sample featuring a surplus in these goods. We thus observe that while countries tend to rely less on other countries to fulfill their demand for traditional essential goods relative to that for non-essential goods, they are very dependent on other countries to satisfy their demand for essential medical goods needed to combat a global pandemic such as COVID-19.

Tariffs We now investigate the extent to which essential and non-essential goods feature systematically different international trade barriers. To do so, Table 2 contrasts summary statistics on the levels of tariffs across product categories.

While average tariffs are low on average across all product categories, we do observe significant heterogeneity both across countries within product categories as well as across product categories. Despite the low average tariff levels, there appears to be significant dispersion across countries, as evidenced by standard deviations across tariffs that are approximately identical to the average values of the tariffs.

Moreover, Table 2 suggests that there are systematic differences across sectors. In par-

[7] Average trade imbalances are computed trimming the bottom 4% of observations. This eliminates countries that report having zero or very low exports in a sector, yielding a trade imbalance (% of total trade) above 95%.

Covid Economics 21, 22 May 2020: 59-99

Table 2: Tariffs (%)

	Average	Std. Dev.	Min	Max
Defense	7.2	7.4	0.0	25.0
Food	7.9	7.2	0.0	34.5
Medical	1.7	1.9	0.0	8.3
Non-essential	5.4	4.3	0.0	14.2

Note: This table reports averages across 93 countries in year 2018. All values are expressed as percentages.

ticular, we observe that traditional essential goods (defense and food) feature higher average tariffs than non-essential goods: 7.2% vs. 5.4%. In contrast, we find that tariffs on imports of medical goods are low both in absolute and relative terms. The average country applies a 1.7% tariff, on average, on medical goods, with a maximum of 8.3%. These values are substantially lower than those applied to the other product categories.

These patterns resemble the patterns of trade documented above. It may be the case, however, that countries impose higher tariffs on imports of traditional essential goods, thus leading to lower dependence of the average country on the rest of the world for the consumption of these goods. Or it may be the case that imports of medical products are subject to lower tariffs than the other product categories, which may account for the higher dependence of the average country on international trade to fulfill the demand for these goods.

Non-tariff barriers While tariffs are an important and popular trade policy instrument that are particularly easy to quantitatively compare across countries and goods, countries rely on a variety of other instruments to impose barriers to international trade. Thus, we now compare the prevalence of non-tariff barriers across the different product categories under analysis. To do so, Table 3 reports the average frequency and coverage indexes for each product category.

We find that, according to both indicators, non-tariff barriers are significantly more prevalent on essential than non-essential goods. In particular, more than 90% of the traditional essential goods traded are subject to some type of non-tariff barriers, while this is only the case for 50% of non-essential goods. Similarly, while more than 87% of the import value of traditional essential goods are subject to non-tariff barriers, this is only the case for 64% of

CEPR PRESS

Table 3: Non-tariff barriers

	Frequency index (%)	Coverage index (%)
Defense	90.58	87.73
Food	92.59	92.96
Medical	74.51	86.00
Non-essential	50.21	64.12

the import value of non-essential goods.

Note that medical goods are the least protected type of essential good, with 75% of these products and 86% of these imports subject to non-tariff barriers. Although medical goods are more protected than non-essential goods, this finding suggests that prior to the COVID-19 pandemic, countries might not have considered medical goods as essential in the same way as food or defense.

2.3 Trade of essential medical goods during COVID-19

We conclude this section by investigating the response of international trade policy across countries during the ongoing COVID-19 pandemic. On the one hand, some countries have been imposing export restrictions on medical supplies that are essential to fighting the pandemic. On the other hand, some countries have been implementing trade liberalization policies, mostly in the form of import tariff reductions, to make it easier to import these essential medical goods. We ask: To what extent are these heterogeneous changes in international trade policy systematically related to the countries' patterns of international trade imbalances in these goods prior to the pandemic?

To answer this question, we combine data on import liberalizations and export restrictions from Global Trade Alert with the international trade flow data used above. We find that 29 countries have implemented some type of liberalization on medical products necessary to fight COVID-19, whereas 59 countries have imposed restrictions. Table 4 decomposes these changes across countries based on whether they ran a trade surplus or deficit in essential medical goods in the year 2018.

We find that trade policy changes on essential medical goods in the aftermath of COVID-19 appear to be systematically related with the extent to which countries are net suppliers or

Covid Economics 21, 22 May 2020: 59-99

Table 4: Trade of essential medical goods during COVID-19

	Number of countries	Number of countries		Share of countries (%)	
		Import liberalization	Export curbs	Import liberalization	Export curbs
Surplus	22	4	19	18.18	86.36
Deficit	87	25	40	28.74	45.98

Note: Surplus (deficit) refers to countries with positive (negative) net exports in essential medical goods in 2018.

net purchasers of these goods vis-a-vis the rest of the world. In particular, about 86% of the countries that have a surplus in medical equipment have implemented export restrictions, whereas only 46% of those with a deficit of these goods have done so. Similarly, about 18% of the countries with a surplus in these goods have reduced import barriers on these goods, whereas only about 30% of those with a deficit in them them did so.

3 Model

We study a world economy populated by two countries, home and foreign, and two sectors: a sector that produces essential goods and one that produces non-essential goods. Each country produces a domestic variety in each sector. Thus, there are four goods in the world economy: a home and a foreign variety of essential goods, and a home and a foreign variety of non-essential goods. All of these goods are traded internationally.

Each country is populated by five types of representative agents: a household, a producer of domestic essential goods, a producer of domestic non-essential goods, a producer of an essential good composite, and a producer of a non-essential good composite.

While the structures of the two countries are identical, we allow some parameters to be country specific. Thus, throughout the rest of this section we describe each of these agents focusing on the home country, and referring to variables *chosen* by the foreign country with an asterisk ("*"). We refer to variables corresponding to the essential and non-essential goods using subscripts e and c, respectively.

Covid Economics 21, 22 May 2020: 59-99

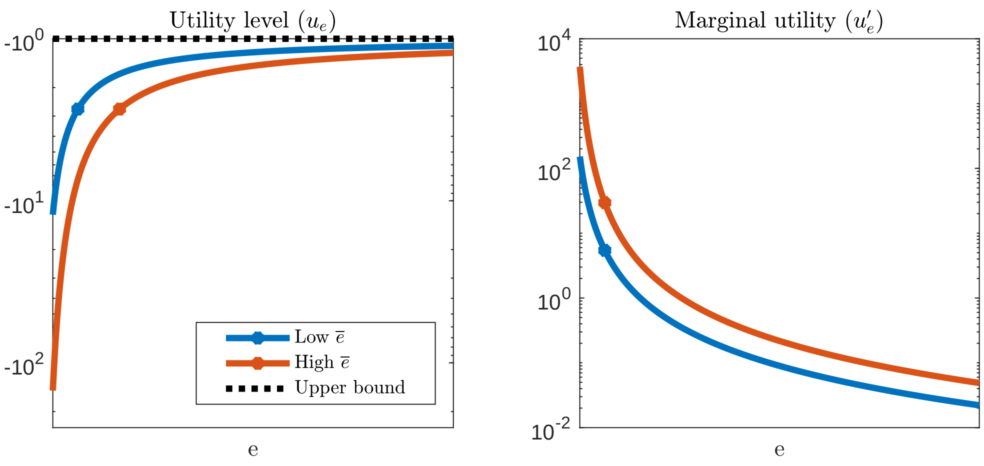

Figure 1: Utility function

3.1 Household

Each country is populated by a representative household who is infinitely lived and discounts the future at rate $\beta < 1$. The household is endowed with one unit of labor that is supplied inelastically at wage rate w_t, and the household also owns domestic producers of essential and non-essential goods. Thus, every period the household earns labor income w_t as well as the profits or losses $\pi_{e,t}$ and $\pi_{c,t}$ incurred by the respective domestic producers. We assume that the household does not have access to international financial markets. The household's budget constraint in a given period is then given by:

$$p_{c,t}c_t + p_{e,t}e_t = w_t + \pi_{c,t} + \pi_{e,t},$$

where $p_{c,t}$ and $p_{e,t}$ denote the price of non-essential and essential goods, respectively, and c_t and e_t denote the consumption of non-essential and essential goods, respectively.

We assume that the household's period utility function is given by:

$$u(c_t, e_t) = \ln c_t - \gamma \exp\left(\frac{\overline{e}_t}{e_t}\right).$$

The parameter γ controls the relative importance of the two goods for the household's utility, and we refer to $\overline{e}_t$ as the "reference level" of essential goods relative to which household

Covid Economics 21, 22 May 2020: 59-99

Covid Economics 21, 22 May 2020: 59-99

consumption of these goods is evaluated. We model this reference level as exogenous and time varying following a stochastic process that we describe below. Notice that period utility is separable in the consumption of essential and non-essential goods. Then, while the level of utility derived from the consumption of non-essential goods is given by standard logarithmic utility, the utility derived from the consumption of essential goods is non-standard.

The goal of this non-standard utility specification is to capture some dimensions along which essential goods might be different from non-essential goods. First, the utility derived from the consumption of essential goods is a function of the ratio between the consumption level e_t and a reference level $\overline{e}_t$ of essential goods. Thus, a given level of essential goods providing high utility levels depends on whether e_t is sufficiently higher than $\overline{e}_t$ rather than on the absolute level of consumption. This captures the idea that households have reference levels for the consumption of essential goods such as food or health services, and a given level of consumption is high or low depending on the comparison with the respective reference level. While akin to the subsistence level featured by Stone-Geary preferences, our specification allows $e_t < \overline{e}_t$. Later in the paper we investigate the impact of changes in this reference level $\overline{e}_t$ to capture an increase in the required amount of essential goods during a pandemic.

Second, while utility derived from the consumption of essential goods is strictly increasing for every $e_t > 0$, the utility level attained features an asymptote at $-\gamma$ for $e_t \to \infty$. This captures the idea that there is a satiation level (in the limit) for essential goods such as food or medical services: increasing levels of food or health services might increase utility marginally, but there is an upper bound to how much they can do so.

Figure 1 plots the utility levels (left panel) and the marginal utilities (right panel) corresponding to a high and low value of $\overline{e}_t$. Notice that essential goods consumed below the reference level are penalized with exponentially decreasing levels of utility. Moreover, higher levels of the reference level $\overline{e}_t$ lead to uniformly lower levels of utility. Thus, given a level of consumption of essential goods e_t, an increase in $\overline{e}_t$ can lead to a substantial decline in utility as well as to a substantial increase in the marginal utility at such consumption level.

The household's problem can be written as:

$$\max_{\{c_t,e_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\ln c_t - \gamma \exp\left(\frac{\bar{e}_t}{e_t}\right)\right]$$

subject to

$$p_{c,t}c_t + p_{e,t}e_t = w_t + \pi_{c,t} + \pi_{e,t} \quad \forall t = 0, ..., \infty,$$

where the expectations operator is taken conditional on the information set at time period 0. While the household's decisions are static, we set up the household's infinite horizon problem to ease the computation of welfare effects later in the paper.

3.2 Domestic producers of good $j \in \{c, e\}$

In each sector $j \in \{c, e\}$, a representative firm produces a domestic variety of either non-essential (if $j = c$) or essential (if $j = e$) goods using capital ($k_{j,t}$) and labor ($n_{j,t}$) at a given level of productivity A_j. The amount produced y_j is given by $y_{j,t} = A_j n_{j,t}^{\alpha} k_{j,t}^{1-\alpha}$, where α denotes the share of labor in production; and notice that we assume that sectoral productivity is time invariant and that the technology has constant returns to scale.

Every period firms choose the amount of labor to use in that period and the amount of investment to alter the capital stock in the following period. We assume that investment and the capital stock consist of non-essential goods so that increasing the amount of capital by one unit in the following period requires investing $i_{j,t}$ units of non-essential goods today. Capital depreciates at rate δ, which implies that next period's capital stock $k_{j,t+1}$ is given by $(1-\delta)k_{j,t} + i_{j,t}$.

Given our focus on investigating the adjustment of the world economy for increased demand for essential goods, we introduce capital and labor adjustment costs to help us discipline the degree to which countries can reallocate production across sectors. We assume that capital and labor adjustment costs are quadratic, consist of non-essential goods, and are given by:

$$\phi_k(k_{j,t+1}, k_{j,t}) = \frac{\Omega_k}{2}\left(\frac{k_{j,t+1}}{k_{j,t}} - 1\right)^2$$

$$\phi_n(n_{j,t}, n_{j,t-1}) = \frac{\Omega_n}{2}\left(\frac{n_{j,t}}{n_{j,t}-1} - 1\right)^2,$$

Covid Economics 21, 22 May 2020: 59-99

where Ω_k and Ω_n are positive constants that control the degree to which adjusting production inputs is costly.

The firm's problem then consists of choosing the amount of labor and investment in each period to maximize lifetime discounted profits, where future returns are discounted at rate m_{t+1}. The firm's problem can be expressed as:

$$
\begin{aligned}
\max_{\{n_{j,t},i_{j,t},k_{j,t+1},y_{j,t}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} m_t \left[q_{j,t}y_{j,t} - w_t n_{j,t} - p_{c,t}i_{j,t} - p_{c,t}\phi_k(k_{j,t+1},k_{j,t}) - p_{c,t}\phi_n(n_{j,t},n_{j,t-1})\right] \\
&\text{subject to} \\
&k'_{j,t} = (1-\delta)k_{j,t} + i_{j,t} \quad \forall t = 0,...,\infty \\
&y_{j,t} = A_j n_{j,t}^{\alpha} k_{j,t}^{1-\alpha} \quad \forall t = 0,...,\infty.
\end{aligned}
$$

Notice that the price of the domestic variety of good j is given by $q_{j,t}$ and that adjustment costs are expressed in units of the non-essential good.

Given our interest in understanding the optimality of production decisions in the context of essential goods, we assume in our baseline model that firms are myopic and do not internalize the impact of their production decisions on the household's utility. Thus, in our baseline model we assume that the rate at which firms discount the future is given by $m_t = \beta^t$. In the quantitative analysis, we investigate the importance of this externality relative to an economy in which firms internalize the impact of production and investment decisions on the utility derived from the consumption of essential goods.

3.3 Producers of composite good $j \in \{c, e\}$

A representative firm produces a composite good $z_{j,t}$ by combining varieties of the good produced in both the home ($z_{j,h,t}$) and foreign ($z_{j,f,t}$) countries. To do so, the firm operates a constant elasticity of substitution technology with elasticity $\sigma > 0$ given by:

$$
z_{j,t} = \left[\omega_j z_{j,h,t}^{\frac{\sigma-1}{\sigma}} + (1-\omega_j) z_{j,f,t}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}},
$$

where $\omega_j \in (0,1)$ denotes the relative weight of home vs. foreign goods in the production of the composite good.

Covid Economics 21, 22 May 2020: 59-99

The problem of the firm consists of choosing the amounts of inputs $z_{j,h,t}$ and $z_{j,f,t}$ to maximize profits. While the price paid in the home country for the variety of good j produced in that country is given by $q_{j,t}$, the price paid in the home country for the variety produced in the foreign country is given by $\tau_j q^*_{j,t}$, the product of the foreign price and trade costs τ_j. We assume that these trade costs are ad valorem and such that $\tau_j \geq 1$.

The firm's problem in period t is then given by:

$$\max_{z_{j,t}, z_{j,h,t}, z_{j,f,t}} p_{j,t} z_{j,t} - q_{j,t} z_{j,h,t} - \tau_j q^*_{j,t} z_{j,f,t}$$

subject to

$$z_{j,t} = \left[\omega_j z_{j,h,t}^{\frac{\sigma-1}{\sigma}} + (1-\omega_j) z_{j,f,t}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}.$$

3.4 Reference level of essential goods

The process for the time-varying reference level of essential goods $\bar{e}$ is given by:

$$\log \bar{e}_{t+1} = (1-\rho) \log \bar{e} + \rho \log \bar{e}_t + \varepsilon_t$$

where ρ denotes the persistence of the reference level, $\bar{e}$ denotes the steady-state reference level, and $\varepsilon_t \sim N(0, \sigma_e^2)$.

3.5 Market clearing conditions

We let the price of the domestic variety of non-essential goods in the home country $q_{c,t}$ be our numeraire. Then, a *competitive equilibrium of the world economy* consists of:

- prices $\left\{w_t, w^*_t, p_{c,t}, p_{e,t}, p^*_{e,t}, p^*_{c,t}, q_{e,t}, q^*_{e,t}, q^*_{c,t}\right\}_{t=0}^{\infty}$,
- home country allocations $\{c_t, e_t, \pi_{c,t}, \pi_{e,t}, n_{c,t}, n_{e,t}, k_{c,t}, k_{e,t}, i_{c,t}, i_{e,t}, y_{c,t}, y_{e,t}, z_{c,t}, z_{e,t}\}_{t=0}^{\infty}$,
- foreign country allocations $\left\{c^*_t, e^*_t, \pi^*_{c,t}, \pi^*_{e,t}, n^*_{c,t}, n^*_{e,t}, k^*_{c,t}, k^*_{e,t}, i^*_{c,t}, i^*_{e,t}, y^*_{c,t}, y^*_{e,t}, z^*_{c,t}, z^*_{e,t}\right\}_{t=0}^{\infty}$

such that the following conditions hold:

- Home country:

Covid Economics 21, 22 May 2020: 59-99

1. Given prices, allocations solve the household's problem
2. Given prices, allocations solve problem of domestic producers
3. Given prices, allocations solve problem of composite good producers
4. Labor market clears: $n_{c,t} + n_{e,t} = 1\ \forall t$
5. Home essential goods market clearing: $e_{h,t} + \tau_e^* e_{h,t}^* = y_{e,t}\ \forall t$
6. Home non-essential goods market clearing: $c_{h,t} + \tau_c^* c_{h,t}^* = y_{c,t}\ \forall t$
7. Essential composite good market clearing: $e_t = z_{e,t}\ \forall t$
8. Non-essential composite good market clearing:

$$c_t + \sum_{j\in\{c,e\}} \left[i_{j,t} + \frac{\Omega_k}{2}\left(\frac{k_{j,t+1}}{k_{j,t}} - 1\right)^2 + \frac{\Omega_n}{2}\left(\frac{n_{j,t}}{n_{j,t-1}} - 1\right)^2 \right] = z_{c,t}\ \forall t$$

- Foreign country:

1. Given prices, allocations solve household's problem
2. Given prices, allocations solve problem of domestic producers
3. Given prices, allocations solve problem of composite good producers
4. Labor market clearing: $n_{c,t}^* + n_{e,t}^* = 1\ \forall t$
5. Foreign essential goods market clearing: $\tau_e e_{f,t} + e_{f,t}^* = y_{e,t}^*\ \forall t$
6. Foreign non-essential goods market clearing: $\tau_c c_{f,t} + c_{f,t}^* = y_{c,t}^*\ \forall t$
7. Essential composite good market clearing: $e_t^* = z_{e,t}^*\ \forall t$
8. Non-essential composite good market clearing:

$$c_t^* + \sum_{j\in\{c,e\}} \left[i_{j,t}^* + \frac{\Omega_k}{2}\left(\frac{k_{j,t+1}^*}{k_{j,t}^*} - 1\right)^2 + \frac{\Omega_n}{2}\left(\frac{n_{j,t}^*}{n_{j,t-1}^*} - 1\right)^2 \right] = z_{c,t}^*\ \forall t.$$

4 Quantitative impact of a pandemic

We use the model presented in the previous section to investigate the impact of a pandemic on a wide range of economic outcomes as well as on welfare. We then evaluate the role played by key ingredients of the model and the model's parametrization on our findings. In

Covid Economics 21, 22 May 2020: 59-99

the next section, we examine the extent to which international trade policy in the short and long run can impact our findings.

We begin by parametrizing the model presented in the previous section. Our goal is to conduct a numerical exercise in a hypothetical world in which the two countries are fully symmetric except for their productivity in the production of domestic goods. In particular, we assume that the home country is relatively more productive in the production of non-essential goods, while the foreign country is relatively more productive in the production of essential goods.

Thus, in this preliminary analysis we do not quantify the effect of a specific pandemic such as COVID-19. Instead, we use the model to investigate the effect of a generic pandemic in a world in which countries are heterogeneous in their dependence on other countries for their consumption of essential goods.

4.1 Parameterization

To parametrize the model, we partition the parameter space into two groups: a set of predetermined parameters and a set of parameters that we estimate to match moments chosen to ensure that the world economy we study resembles actual economies along some key dimensions.

Predetermined parameters The set of predetermined parameters consists of the discount factor β, the elasticity of substitution between domestic and foreign varieties of the essential and non-essential goods σ, the labor share α, the capital depreciation rate δ, the steady-state reference level $\overline{e}$, capital adjustment costs Ω_k, labor adjustment costs Ω_n, weights ω_e and ω_c, the productivity of the home country in essential goods, and the productivity of the foreign country in non-essential goods. All of these parameters except for the productivities are assumed to be identical across both countries. Table 5 reports the parameter values used throughout.

We set the value of β, σ, α, and δ to standard values from the literature. In particular, we interpret periods in the model as quarters and thus set β to 0.99, which implies a quarterly real interest rate of 1%. Also, we set the elasticity of substitution to 4 following the work of Simonovska and Waugh (2014), implying that domestic and foreign varieties are relatively substitutable.

Covid Economics 21, 22 May 2020: 59-99

CEPR PRESS

Covid Economics 21, 22 May 2020: 59-99

Table 5: Predetermined parameters

Parameter	Value	Description
A. Parameters common across countries		
β	0.99	Discount factor
σ	4	Elasticity of substitution
α	0.66	Labor share
δ	0.06	Capital depreciation rate
$\bar{e}$	0.40	Reference level of essential goods
Ω_k	50	Capital adjustment costs
Ω_n	50	Labor adjustment costs
$\omega_e = \omega_c$	0.50	Weight on home goods
ρ	0.85	Persistence of shocks to $\bar{e}$
σ_e	0.10	Std. dev. of shocks to $\bar{e}$
B. Country-specific parameters		
$A_e = A_c^*$	1	Sectoral productivities

The remaining parameters are set to illustrate the mechanisms at play. For instance, we set capital and labor adjustment costs Ω_k and Ω_n to 50, allowing us to study an economy in which the adjustment of capital and labor following a shock is neither instantaneous nor zero.[8] Similarly, we set the steady-state reference level $\bar{e}$ to equal 0.40, which given the rest of the parameters yields that shocks to the reference level of essential goods have a significant impact on allocations. We assume the process followed by $\bar{e}_t$ has a persistence ρ equal to 0.85 and a standard deviation σ_e equal to 0.10.

Finally, we normalize to 1 the productivity of the home country in the production of essential goods A_e as well as the productivity of the foreign country in non-essential goods A_c^*. We also set the weights on home goods to $\omega_e = \omega_c = 0.50$.

Estimated parameters The set of estimated parameters consists of the home country's productivity in the production of non-essential goods A_c; the foreign country's productivity in the production of essential goods A_e^*; the weight γ of essential goods in the household's

[8]This parametrization of sectoral adjustment costs is in the range used by Kohn, Leibovici, and Tretvoll (forthcoming) in a similar economic environment.

period utility function; and the international trade costs $\tau_e = \tau_e^*$ and $\tau_c = \tau_c^*$ in essential and non-essential goods, respectively. We assume that $A_c = A_e^*$ such that the two countries are symmetric in all parameters except that they have reverse patterns of productivity across sectors. Thus, there are four parameter values that need to be pinned down.

We choose these four parameter values to ensure that the following moments hold in the home country's steady state: (i) the net exports-to-output ratio in essential goods is equal to -20%, (ii) the share of essential goods in aggregate GDP is equal to 10%, (iii) the import share in essential goods is 25%, and (iv) the import share in non-essential goods is 25%. While these moments are arbitrary, we believe they capture reasonable features of actual economies: (i) some countries run sectoral trade imbalances in essential goods, (ii) the production of essential goods constitutes a small share of aggregate GDP, and $(iii) - (iv)$ the shares of imports in the consumption of essential and non-essential goods are not too high.

The estimated parameters and the model counterpart of the target moments are reported in Table 6. We find that these four parameter values can be chosen to match the four target moments exactly. To do so, the model requires that the home country is more productive in the production of non-essential goods, while the foreign country is more productive in the production of essential goods. The model requires a very low utility weight on essential goods in order to match the low share of essential goods in aggregate GDP. And, finally, to rationalize a 25% import share in each of the goods, the model requires that trade costs on essential goods are considerably higher than those on non-essential goods.

Finally, note that the different specification of the period utility function that correspond to each of the goods combined with the assumption that one country is relative more productive than the other in the production of essential goods (and vice-versa for non-essential goods) implies that the steady-state allocations across countries are asymmetric. The set of estimated and predetermined parameters imply that the foreign country features a trade surplus in essential goods, with these goods accounting for 14% of aggregate GDP. In addition, the foreign country imports a small fraction of essential goods but a very large fraction of non-essential goods.

Covid Economics 21, 22 May 2020: 59-99

CEPR PRESS

Table 6: Estimated parameters

Parameter	Value	Description
$A_c = A_e^*$	1.13	Sectoral productivities
γ	3.20×10^{-4}	Utility weight on essential goods
$\tau_e = \tau_e^*$	1.75	Trade costs on essential goods
$\tau_c = \tau_c^*$	1.36	Trade costs on non-essential goods

	Targeted		*Untargeted*
	Home country S.S.		Foreign country S.S.
Moment	Target value	Model	Model
NX_e/GDP_e	−0.20	−0.20	0.16
GDP_e/GDP	0.10	0.10	0.14
$\frac{M_e}{p_e e}$	0.25	0.25	0.09
$\frac{M_c}{p_c c}$	0.25	0.25	0.31

4.2 Dynamics following a global pandemic

In the rest of the paper, we use the model as parametrized above to investigate the macroeconomic dynamics following a global pandemic. We model a global pandemic as a shock to the reference level of essential goods $\overline{e}_t$ of the same extent in the two countries. This modeling approach is motivated by the COVID-19 pandemic, where the amount of medical goods required to keep individuals safe and healthy increased suddenly, e.g., protective medical equipment such as sterile gloves, medical protective clothing, protective goggles, and masks. Our analysis therefore abstracts from most multifaceted effects of a pandemic (e.g., social distancing policies, infection and death rates, increased unemployment), solely restricting attention to the increased demand for essential goods.

To examine the effect of an increase in $\overline{e}_t$ in the two countries, we compute impulse response functions following a shock that increases the value of $\overline{e}_t$ by 50% (in logs).[9] Figure 2 plots the impulse response functions of key variables of the model in response to a global pandemic. Each panel plots two lines: a solid line that plots the dynamics of the variable

[9]Throughout the paper we restrict attention to the pruned third-order approximation of the model around its deterministic steady state. All impulse response functions are generalized impulse response functions computed at the ergodic mean.

Covid Economics 21, 22 May 2020: 59-99

CEPR PRESS

in the home country, and a dashed line that plots its foreign counterpart. Unless otherwise specified, each panel presents the dynamics of a variable as a log-deviation from its average value.[10] Given the moderate persistence of the shock, as parametrized above, we plot the dynamics over the first 20 periods, when most variables are close to their average values.

Our key finding is that a global pandemic in which the reference level of essential goods increases identically in both countries nevertheless has significantly heterogeneous effects across them. In each country, the increase in the reference level of essential goods leads to an increase in the demand for these goods; as Figure 1 shows, at any given level of consumption of essential goods, a higher reference level increases the marginal utility of consumption of essential goods.

Given that the production of essential goods in each of the countries is subject to capital and labor adjustment costs, producers of these goods cannot costlessly reallocate production inputs across sectors to boost production of essential goods in sufficient amounts to satiate the higher demand due to the pandemic. Therefore, the excess demand for essential goods leads to a substantial increase in p_e/p_c, the price of these goods relative to non-essential goods.

Now, while p_e/p_c increases to the same extent in both economies, it leads to a persistent decline in the period utility function of households in the home country but to a persistent increase of that of households in the foreign country. Recall that the two countries are symmetric except for their sectoral productivities: the home country is relatively more productive in non-essential goods, while the foreign country is relatively more productive in essential goods. This difference implies that the foreign country is a net exporter of essential goods, while the home country is a net importer of these goods. Thus, a substantial increase in the relative price of essential goods has a negative welfare effect on net importers of these goods but a positive one on net exporters of these goods.

Both countries earn relatively more for their sales of essential goods, and both countries have to spend relatively more for their purchases of essential goods. However, given that the foreign country is the more productive producer of essential goods, this country sells relatively more essential goods than it purchases, thus experiencing a net benefit from the pandemic.

[10]The only exceptions are the four panels depicting the dynamics of sectoral net exports which are expressed in level deviations around their average value.

Covid Economics 21, 22 May 2020: 59-99

Covid Economics 21, 22 May 2020: 59-99

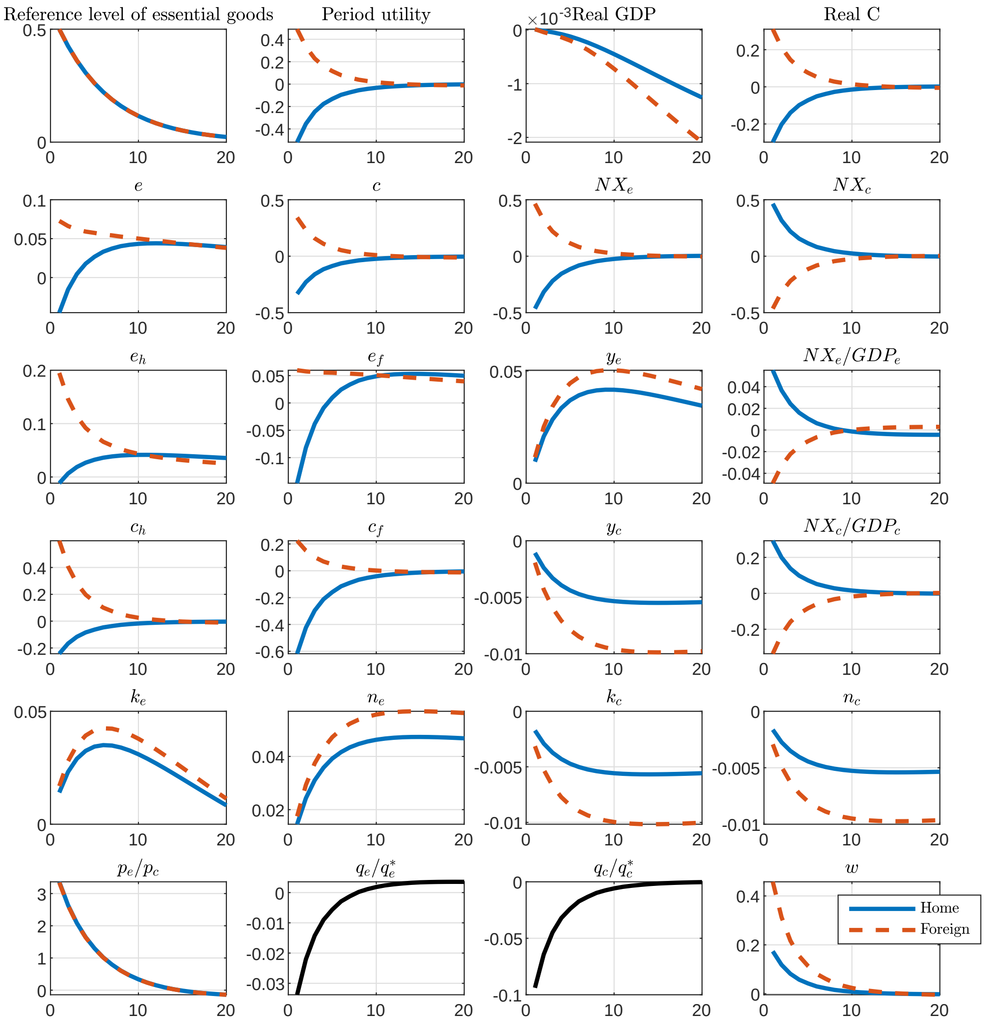

Note: x-axes denote time periods. x-axes are expressed as log deviations from long-run averages except for the panels featuring net exports or net export- to-GDP ratios, which are expressed as deviations from long-run average levels.

Figure 2: Dynamics following a global pandemic

But how can this country *benefit* from a global pandemic? The substantial increase in the price of essential goods allows foreign households to increase their consumption of domestic and imported essential goods while the country remains a net exporter of these goods. On the other hand, the increased value of exported essential goods allows foreign households to

Covid Economics 21, 22 May 2020: 59-99

also increase their consumption of non-essential goods, both domestic and imported.

The home country experiences welfare losses from a global pandemic through exactly the reverse forces. As a net importer of essential goods, the increase in the price of essential goods relative to the price of the good produced domestically implies that this country is forced to cut back its consumption of essential goods despite its higher need for them. As aggregate consumption decreases, households in the home country cut their consumption of non-essential goods to reallocate resources to prevent the consumption of essential goods from declining even further.

4.3 Welfare implications of a global pandemic

The discussion and impulse response functions above show that the effects of a global pandemic can be substantially heterogeneous across countries depending on whether a country is a net exporter or net importer of essential goods. Indeed, the analysis above shows that the period utility function increases for net exporters but decreases for net importers of these goods. We now further investigate the welfare implications of a global pandemic by quantifying the lifetime impact of a pandemic in terms of consumption-equivalent units. To do so, for each country we ask: What fraction of the consumption of non-essential goods would a household living forever in the deterministic steady state of the model be willing to give up to avoid experiencing the global pandemic examined in the impulse response functions above?

We present the welfare impact of a global pandemic in the home and foreign countries in the first row of Table 7. Consistent with our findings above, we find that a global pandemic leads to a substantial welfare loss in the home country and to a sizable welfare gain in the foreign country. In particular, an agent living in the deterministic steady state of the home country would be willing to sacrifice 1.77% of consumption of non-essential goods every period to avoid being hit by a global pandemic. In contrast, an agent living in the deterministic steady state of the foreign country would demand her consumption to increase 0.93% every period to deter her from preferring to experience a global pandemic.

A word of caution is in order regarding the interpretation of our quantitative findings. Recall that the quantitative exercise in the current version of the paper is designed to illustrate the mechanisms of the model rather than to obtain a quantification that might be informative about the welfare implications of the COVID-19 or other pandemics. Our goal

CEPR PRESS

Covid Economics 21, 22 May 2020: 59-99

Table 7: Welfare implications of a global pandemic

Model	Welfare gain for...	
	Home country	Foreign country
Baseline	-1.77%	0.93%
Low adjustment costs	-1.46%	-0.12%
High adjustment costs	-2.35%	2.44%
Stone-Geary	-3.16%	0.68%
Cobb-Douglas weight	-5.20%	-4.30%
Rational firms	-1.68%	0.86%
No sectoral imbalances	-0.38%	-0.38%
Higher trade barriers	-0.87%	-0.0004%
Raise trade barriers when pandemic hits	-2.16%	1.03%

is to eventually quantify these effects by estimating the model to match salient features of the data; future versions of the paper will attempt to bridge this gap and provide welfare estimates that might be more directly applicable to pandemics in actual economies.

4.4 Understanding the mechanism

The results presented in the previous sections show that a global pandemic has substantially heterogeneous effects across countries depending on the extent to which countries are net exporters or net importers of essential goods. We now investigate the role played by various ingredients of the model in accounting for these findings.

4.4.1 Role of adjustment costs

We begin by investigating the role played by the adjustment costs that need to be incurred by producers of essential and non-essential goods, to adjust the amount of capital and labor used in production. To do so, we compute impulse response functions for the global pandemic examined in the previous section under alternative degrees of capital and labor adjustment costs. On one end, we consider an economy with lower adjustment costs on both capital and labor; in particular, we set $\Omega_n = \Omega_k = 10$. On the other end, we consider an economy with

Covid Economics 21, 22 May 2020: 59-99

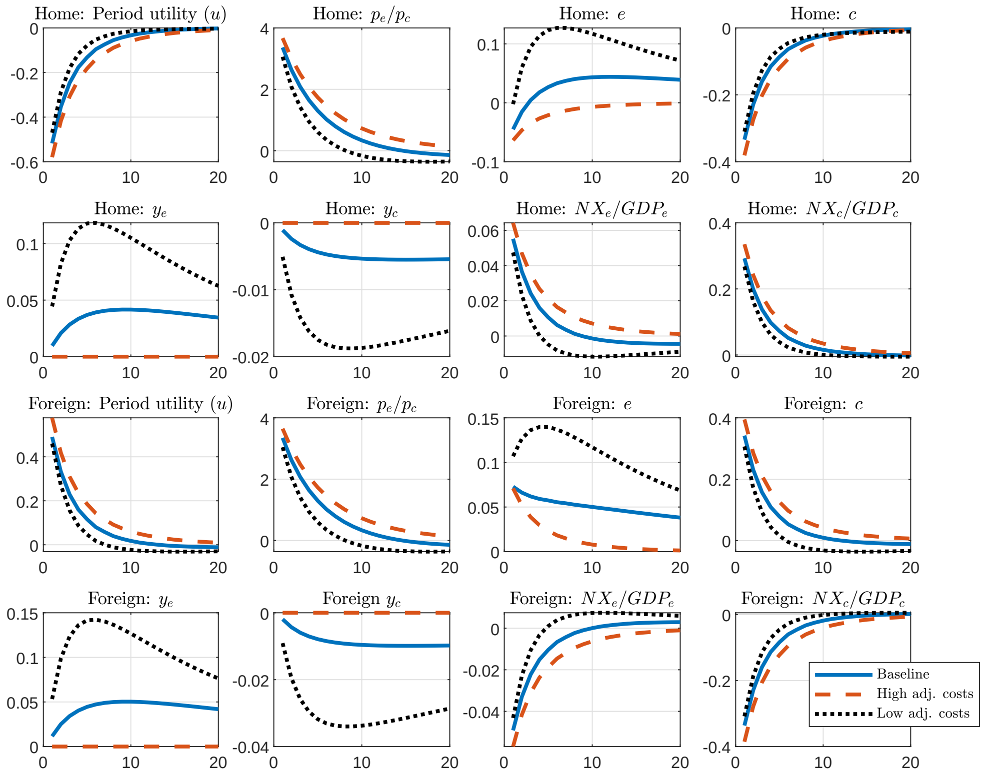

Note: x-axes denote time periods. y-axes are expressed as log deviations from long-run averages except for the panels featuring net exports or net export-to-GDP ratios, which are expressed as deviations from long-run average levels.

Figure 3: Role of adjustment costs

much higher adjustment costs on both capital and labor; in particular, we set $\Omega_n = \Omega_k = 10^6$.

Figure 3 plots the impulse response functions for the baseline model and for the economies with alternative degrees of factor adjustment costs. In contrast to the impulse response functions presented above, each panel now presents the dynamics of a given variable in a given country, and each line in a panel corresponds to a different adjustment cost specification: a solid corresponds to the baseline model, a dashed line corresponds to the economy with high adjustment costs, and a dotted line corresponds to the economy with low adjustment costs. To ease the presentation, we restrict attention to a subset of key variables that allow us to illustrate the main effects.

Covid Economics 21, 22 May 2020: 59-99

We find that adjustment costs play a crucial role in accounting for the implications of our baseline model. Insofar as producers of essential and non-essential goods are not able to easily increase the production of essential goods when a pandemic hits, the relative price of essential to non-essential goods increases persistently, leading to the effects described above in our baseline model. Higher adjustment costs then simply reinforce these effects, leading to a higher increase in the relative price of essential goods and, thus, to a larger increase in utility in the foreign country and to a larger decline in utility in the home country. Notice that production of essential and non-essential goods remains approximately unchanged under high adjustment costs. Thus, in this case, the total amount produced of each good remains almost unchanged; so the adjustment only occurs via the reallocation of consumption across countries.

In contrast, in an economy with lower adjustment costs, producers of essential and non-essential goods are able to easily adjust their production in response to changing economic conditions. The relative price of essential goods increases on impact given that capital cannot be reallocated within the period, as investment takes one period to materialize. Yet, after the first period, the relative price of essential goods decreases rapidly as both countries are able to rapidly reallocate their factors of production towards this sector. These effects have very significant implications for the dynamics of utility. While utility increases on impact in the foreign country, the gains rapidly evaporate, declining even below their long-run average values a period after the initial shock; in every period the foreign country is worse off than in the baseline model with higher adjustment costs. In contrast, the reverse is the case for the home country: with low adjustment costs, utility declines relatively less than in the baseline model.

The welfare implications in terms of consumption-equivalent units reported in the second and third rows of Table 7 are consistent with these findings. In particular, households living in the deterministic steady state in the foreign country need to be compensated with 0.93% and 2.44% of consumption of non-essential goods every period in the baseline and high-adjustment-cost economies, respectively, to be prevented from preferring to experience a global pandemic. In contrast, with low adjustment costs, they are willing to give up 0.12% of consumption of non-essential goods every period to avoid experiencing a global pandemic. The effects are reversed for households in the home country, for whom the welfare cost of a pandemic is lower (-1.46%) in an economy with low adjustment costs and considerably higher in an economy with high adjustment costs (-2.35%).

4.4.2 Role of essential good preferences

We now investigate the role played by our specification of the utility derived from the consumption of essential goods relative to a reference level. To do so, we contrast the impulse response functions of our baseline model with those arising from two alternative ways of modeling preferences for essential goods.

First, we consider an economy with Stone-Geary preferences, where the household's utility function in each country is given by $u(c_t, e_t) = \ln c_t + \gamma \ln(e_t - \bar{e}_t)$. This more standard way of modeling preferences for essential goods features a subsistence level, so we begin by comparing the implications of our model with this alternative specification. However, these preferences are problematic for our purposes, since they are only well defined for $e > \bar{e}_t$, preventing us from considering large shocks to the subsistence level that could lead to levels of consumption of essential goods below subsistence.

Second, we consider an economy that abstracts altogether from featuring a reference or subsistence level. Instead, we assume that preferences for essential goods are symmetric to preferences for non-essential goods: $u(c_t, e_t) = \ln c_t + \gamma \ln e_t$. These preferences are thus Cobb-Douglas. We engineer a pandemic in this economy by shocking the weight γ on essential goods given that this alternative economy has no reference level to shock.

Figure 4 plots the impulse response functions for the baseline model and for the economies with the alternative preference specifications. To keep the magnitudes comparable to those implied by the baseline model, in the model with Stone-Geary preferences, we consider a 10% shock (in logs) to the subsistence level, and we calibrate this level such that $\frac{e-\bar{e}}{e} = 0.20$ in the steady state. For the economy with Cobb-Douglas preferences, we consider a 50% increase (in logs) in γ.

Our findings are twofold. On the one hand, we find that the implications of our baseline model are very similar to those of the analogous model with Stone-Geary preferences. This is reassuring, as our preference specification for that model attempts to capture the role for subsistence featured by Stone-Geary and avoid the sharp kink faced when consumption reaches the subsistence level.

On the other hand, we find that an alternative specification that abstracts from shocks to a reference or subsistence level and, instead, shocks the Cobb-Douglas weight on essential

Covid Economics 21, 22 May 2020: 59-99

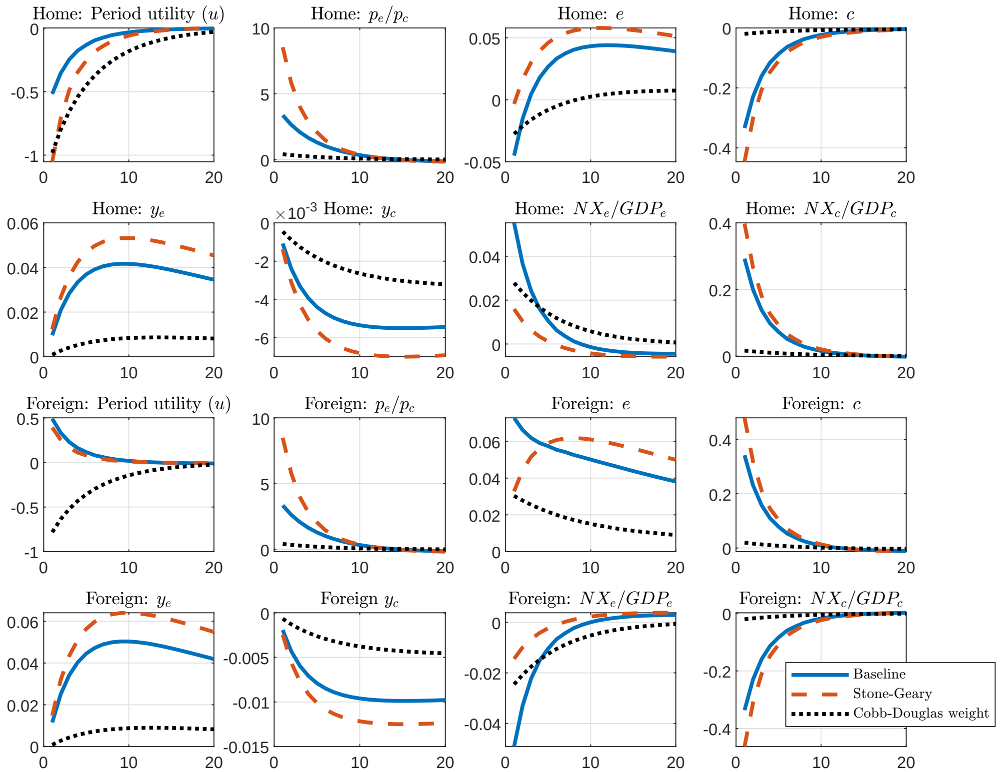

Note: x-axes denote time periods. y-axes are expressed as log deviations from long-run averages except for the panels featuring net exports or net export-to-GDP ratios, which are expressed as deviations from long-run average levels.

Figure 4: Role of essential good preferences

goods generates considerably milder and qualitatively different effects. The key difference is that even a substantial increase in the weight on essential goods appears unable to generate a considerable increase in the relative price of essential goods despite increasing the demand for these goods. Without significant effects on prices as featured in our baseline model, we then observe that both the home and foreign countries experience a decline in utility when the pandemic arrives. Cursory experimentation with larger shocks to γ appears unable to generate significant effects on prices of the magnitude featured by the baseline model. The welfare implications in terms of consumption-equivalent units reported in the fourth and fifth rows of Table 7 are consistent with these findings. We thus conclude that our preferences

Covid Economics 21, 22 May 2020: 59-99

Covid Economics 21, 22 May 2020: 59-99

that include a reference level for essential goods play a critical role in accounting for the implications of the baseline model.

4.4.3 Role of myopic firms

Another dimension that we investigate is the role played by our assumption that producers of both essential and non-essential goods are myopic in their discounting of the future when making hiring and investment decisions. In particular, while households discount the future using a stochastic discount factor that is a function of the relative marginal utilities between consumption today vs. next period, we assume that firms simply discount the future at discount rate β. The goal of this assumption is to introduce a mismatch between the desires of households to ensure a smooth supply of essential goods vis-a-vis the profit maximization motive of producers of such goods. Thus, the model features an externality such that producers of essential and non-essential goods do not internalize the impact of their decisions on the welfare of households.

To investigate the role of this assumption, Figure 5 contrasts the baseline impulse response functions with those arising from a model in which firms discount the future using the household's stochastic discount factor (SDF); we refer to this alternative model as one with "rational" firms. Specifically, we let the SDF in the home country m_t be given by $m_t = \beta^t \lambda_t$, where λ_t is the Lagrange multiplier on the household's budget constraint in period t; the SDF of firms in the foreign country is defined analogously.

We find that the assumption that firms are myopic does not appear to play a fundamental role in accounting for our findings. Most variables respond very similarly both qualitatively and quantitatively in the two models under consideration. Yet, we interestingly find that a pandemic in an economy with rational firms leads to less of a decline in utility in the home country and less of an increase in utility in the foreign county. This is indeed what the welfare implications presented in Table 7 also show.

What accounts for these differential effects? In the economy with myopic firms, producers of essential goods in the home country do not internalize that there are some states of the world (e.g., an increase in the reference level of essential goods) in which producing the essential good would be extremely valuable to the household, even if during normal times producing the essential good in the home country would not be very profitable. Thus, in the economy with myopic firms, producers of essential goods undersupply these goods, and over

Covid Economics 21, 22 May 2020: 59-99

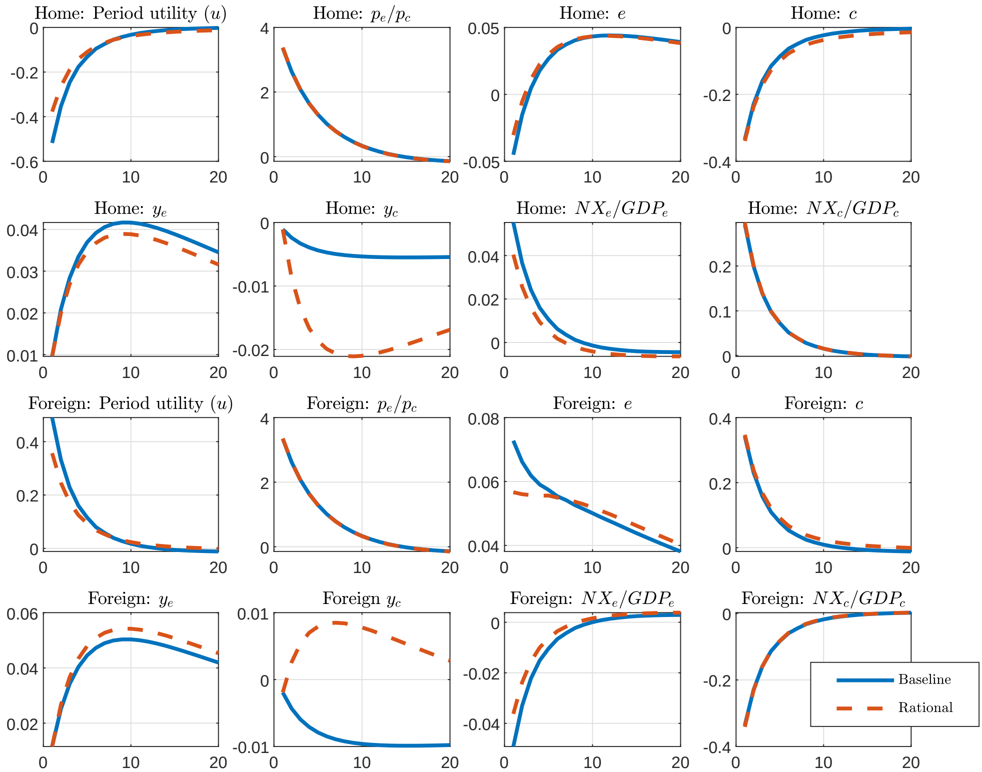

Note: x-axes denote time periods. y-axes are expressed as log deviations from long-run averages except for the panels featuring net exports or net-export-to GDP ratios, which are expressed as deviations from long-run average levels.

Figure 5: Role of myopic firms

In contrast, in an economy with rational firms, producers of the essential good in the home country do internalize that even if they are not very productive, it might nevertheless be worthwhile to produce this good since there are states of the world in which the household would really benefit from it. These firms then produce more essential goods on average than in the model with myopic firms, and the home country runs smaller trade deficits of essential goods on average when firms are rational. Thus, a pandemic in the home country induces lower welfare losses in the home country under rational firms than under myopic ones; the reverse is the case for the foreign country.

These findings show that whether or not firms internalize the effect of their decisions on the utility of households may affect the welfare implications of the model. More work remains to be done to identify conditions that may induce larger differences in the impulse response functions across these two models.

4.4.4 Role of sectoral trade imbalances

We finally investigate the role played by sectoral trade imbalances on the impact of a pandemic across countries. Our discussion above suggests this is a key channel through which changes in the relative price of essential goods triggers a very heterogeneous response across net exporters and net importers of essential goods. In Figure 6 we contrast our baseline impulse response functions with those of an alternative calibration of our model. We keep the estimation approach as in our baseline except for the targeted sectoral imbalance in essential goods of the home country: we now recalibrate the model such that the home country's trade of essential goods is balanced rather than featuring a deficit as in our baseline calibration.

We find that, while the relative price of essential goods increases even more than in our baseline model, the welfare effect in both countries is significantly more muted. Both countries are now able to increase their consumption of essential goods at the expense of producing and consuming fewer non-essential goods. The home country is thus relatively better off than in our baseline model given it no longer runs sectoral trade deficits in essential goods, while the foreign country is relatively worse off since it no longer benefits from selling its excess production of essential goods internationally. The lifetime welfare effects presented in Table 7 are consistent with these findings.

We thus conclude that sectoral trade imbalances in essential goods play a fundamental role in accounting for the cross-country effects of a pandemic implied by our baseline model.

5 Trade policy and the impact of a pandemic

The previous section shows that international trade plays a fundamental role in accounting for the impact of a pandemic across countries. In a world economy in which producers face costs to adjust capital and labor in the short run, net importers of essential goods might experience severe welfare losses, while net exporters of these good might experience welfare gains. We now investigate the role of international trade policy in mitigating or exacerbating

Covid Economics 21, 22 May 2020: 59-99

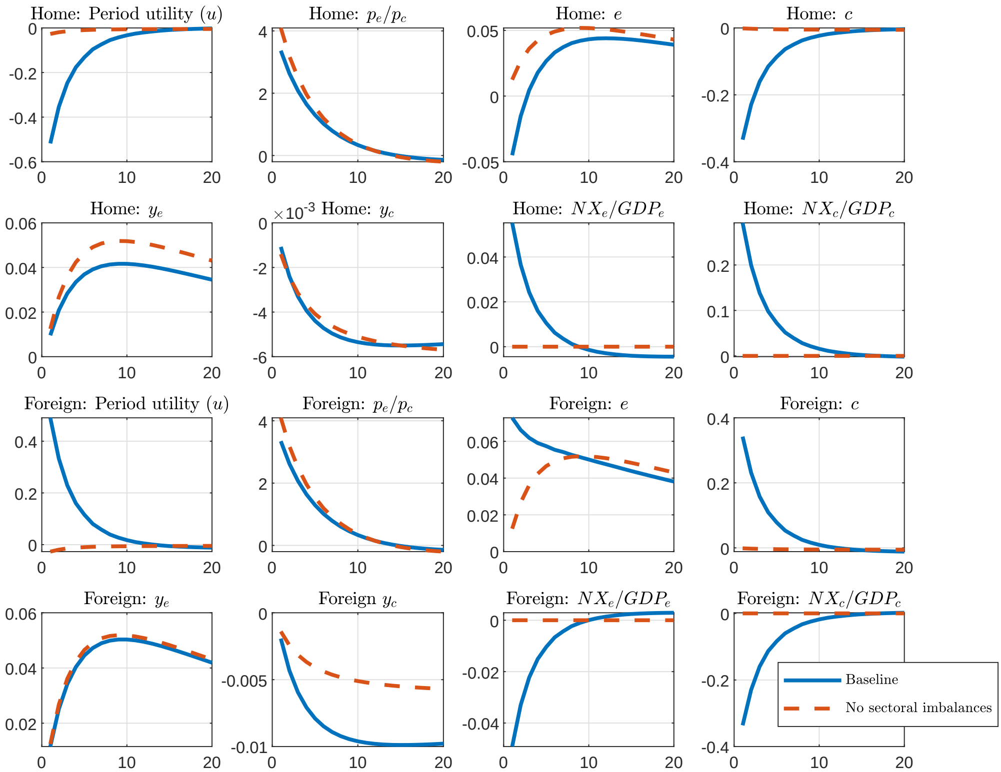

Note: x-axes denote time periods. y-axes are expressed as log deviations from long-run averages except for the panels featuring net exports or net export-to-GDP ratios, which are expressed as deviations from long-run average levels.

Figure 6: Role of sectoral trade imbalances

the impact of a pandemic.

5.1 A pandemic in a world with low vs. high trade barriers

We begin by examining the extent to which the long-run level of international trade barriers at the onset of a pandemic affects its impact on economic outcomes. To do so, Figure 7 contrasts the baseline impulse response functions examined in the previous section with their counterparts from a model featuring higher trade barriers on essential goods. In particular, we consider a world economy in which $\tau_e = \tau_e^* = 3$ (vs. $\tau_e = \tau_e^* = 1.75$ in the baseline), keeping all other parameters unchanged from their baseline values.

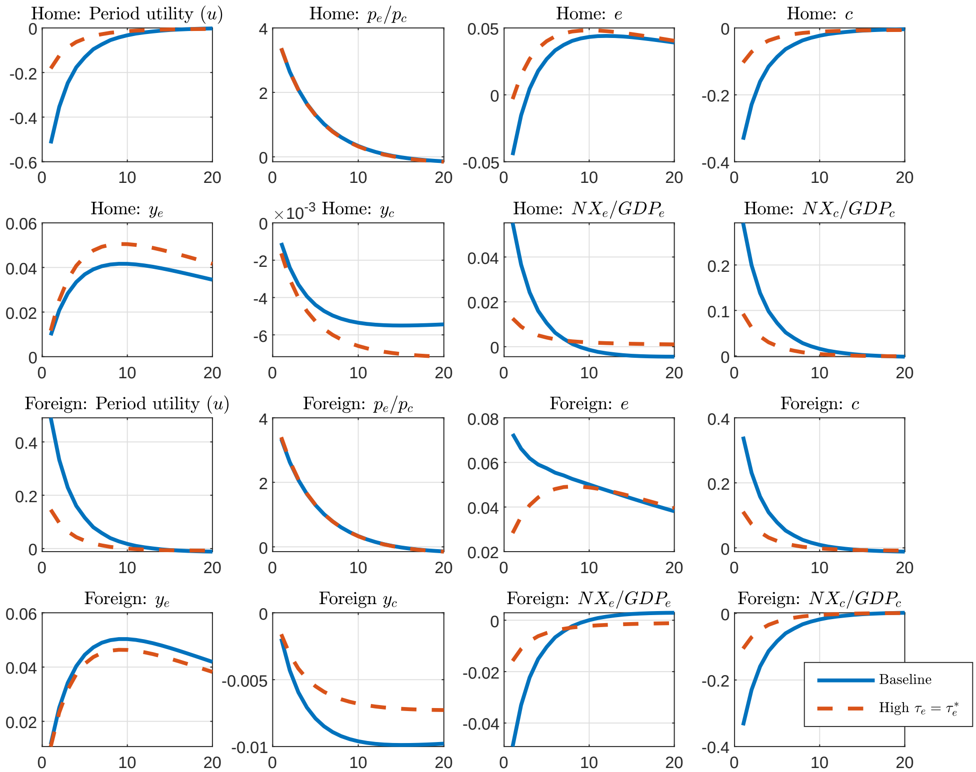

Note: x-axes denote time periods. y-axes are expressed as log deviations from long-run averages except for the panels featuring net exports or net export-to-GDP ratios, which are expressed as deviations from long-run average levels.

Figure 7: A pandemic in a world with low vs. high trade barriers

We find that differences in international trade barriers at the onset of a pandemic can significantly alter the economic implications of a pandemic. Keeping all other parameters unchanged, higher trade costs imply that the home country runs a smaller trade deficit of essential goods in the steady state, while the foreign country runs a smaller trade surplus in these goods. Therefore, while a pandemic continues to lead to a substantial increase in the relative price of essential goods, the relative impact on the home and foreign countries is mitigated by higher trade barriers. As the home country runs a smaller trade deficit in essential goods, the increase in the price of essential goods leads to a smaller decline in the utility of households in this country. Conversely, the smaller trade surplus in these goods

Covid Economics 21, 22 May 2020: 59-99

in the foreign country implies that this country benefits relatively less from a pandemic than in the baseline model. These effects indeed translate to substantially mitigated welfare implications of a pandemic, as shown in Table 7.

We thus conclude that while higher trade barriers on essential goods may reduce the amount of these goods consumed in the steady state, they mitigate the potential vulnerability of net importers of these goods when a global pandemic hits. These findings have important implications for the design of international trade policy. Even if countries may benefit on average from having cheaper access to goods, the reliance on international trade might put these countries in a vulnerable position if these goods are essential and either the supply or demand of these goods is subject to shocks. Thus, these findings raise the following question: To what extent should countries protect, if at all, goods that are essential and whose demand or supply may be subject to shocks? We will address this and other related questions in future versions of the paper.

5.2 A pandemic in a world that raises trade barriers

While the previous question is informative about the design of international trade policy from an ex-ante perspective, it is silent about what countries should do once a pandemic hits. We thus ask: Given some level of trade barriers prevalent in the world when a pandemic hits, what is the impact of raising these barriers in responses to the pandemic? Our answer to this question might help us interpret the policy response of several countries over the first few weeks of the COVID-19 pandemic, when they raised the trade barriers on essential medical equipment.

To answer this question, Figure 8 contrasts our baseline impulse response functions with those implied by our baseline model when a pandemic is accompanied by a simultaneous increase in international trade barriers on essential goods in both countries of 0.50 log points relative to their steady-state values.[11] Our findings contrast sharply with the results presented in the previous subsection. We now find that raising trade barriers when a pandemic hits can significantly amplify the welfare implications in both countries; these effects can be observed both directly via the period utility plotted in the figure and via the welfare implications provided in Table 7. As a net importer of essential goods, the home country's

[11] We assume the increase is transitory and returns back to the steady state at the same rate as the reference level of essential goods $\overline{e}_t$.

Covid Economics 21, 22 May 2020: 59-99

Covid Economics 21, 22 May 2020: 59-99

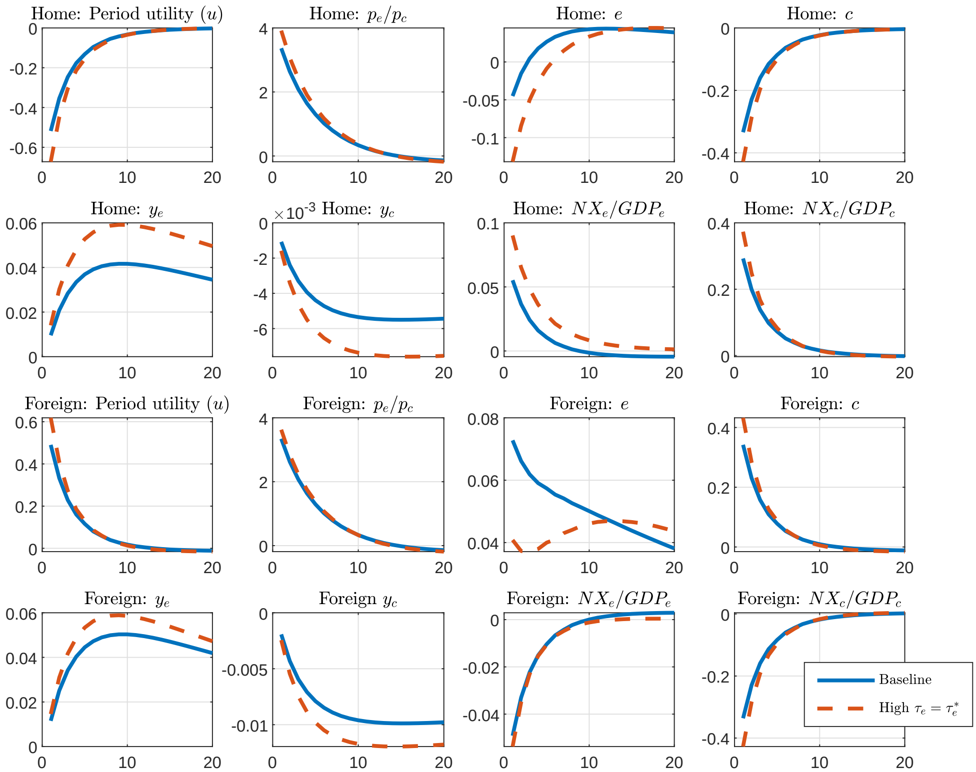

Note: x-axes denote time periods. y-axes are expressed as log-deviations from long run averages except for the panels featuring net exports or net export-to-GDP ratios, which are expressed as deviations from long-run average levels.

Figure 8: A pandemic in a world that responds by raising trade barriers

production structure is such that it relies considerably on the foreign country for its consumption of essential goods. Thus, if trade barriers on essential goods are raised when the pandemic hits, they exacerbate the shortages of these goods the home country already faces and raises the relative prices of these goods even higher.

The exacerbated welfare losses of the home country when trade barriers are increased in the short run during a pandemic contrast sharply with the mitigated welfare loss of the country in a world with persistently higher trade barriers, as examined in the previous subsection. The reason is simply that persistently higher trade barriers make the home country less reliant on international trade for the consumption of essential goods, while a short-run increase in international trade barriers in a world with capital and labor adjustment costs exacerbates the impact of a higher price of essential goods on the country because it cannot rapidly adjust its production structure to the changes in demand.

These findings suggest that the design of international trade policy for essential goods may suffer from a time-inconsistency problem. Net exporters of essential goods may ex-ante prefer to live in a world with low trade costs; but when a pandemic hits, they might be tempted to renege on their commitments and increase trade barriers. The reverse is the case for net importers of essential goods; they may ex-ante prefer to live in a world with high trade costs on essential goods; but when a pandemic hits, they might be tempted to renege on their commitments and decrease trade barriers. We will investigate these issues further in future versions of the paper.

6 Concluding remarks

This paper studies the role of international trade of essential goods during a pandemic. In particular, we investigate the extent to which trade of essential goods mitigates or amplifies the impact of a pandemic and the degree to which these effects depend on international trade policy both in the short and long run.

Our findings so far suggest that the effects of a pandemic across countries depend crucially on countries' trade imbalances in essential goods. Net exporters of essential goods can experience welfare gains during a pandemic, while net importers can experience significant welfare losses. The welfare losses of net importers are lower in a world with high trade barriers, while the reverse is the case for net exporters. Yet, once a pandemic arrives, net exporters of essential goods benefit from an increase in trade barriers, while net importers benefit from a decrease in them. These findings are consistent with preliminary evidence on changes in trade barriers across countries during the COVID-19 pandemic.

Covid Economics 21, 22 May 2020: 59-99

Covid Economics 21, 22 May 2020: 59-99

References

Backus, D. K., P. J. Kehoe, and F. E. Kydland (1992): "International real business cycles," *Journal of political Economy*, 100(4), 745–775.

Baldwin, R. E., and S. J. Evenett (eds.) (2020): *COVID-19 and Trade Policy: Why Turning Inward Won't Work*. CEPR Press.

Evenett, S. (2020): "Tackling COVID-19 together: The trade policy dimension," *Global Trade Alert*, 23.

Evenett, S. J., and L. A. Winters (2020): "Preparing for a Second Wave of COVID-19: A Trade Bargain to Secure Supplies of Medical Goods," Discussion paper.

Geary, R. C. (1950): "A note on A constant-utility index of the cost of living," *The Review of Economic Studies*, 18(1), 65–66.

Kohn, D., F. Leibovici, and H. Tretvoll (forthcoming): "Trade in Commodities and Business Cycle Volatility," *American Economic Journal: Macroeconomics*.

Mendoza, E. G. (1995): "The terms of trade, the real exchange rate, and economic fluctuations," *International Economic Review*, pp. 101–137.

Olson, M., and R. Zeckhauser (1966): "An economic theory of alliances," *The review of economics and statistics*, pp. 266–279.

Ravikumar, B., A. M. Santacreu, and M. Sposi (2019): "Capital accumulation and dynamic gains from trade," *Journal of International Economics*, 119, 93–110.

Schmitt-Grohé, S., and M. Uribe (2018): "How Important are Terms-Of-Trade Shocks?," *International Economic Review*, 59(1), 85–111.

Simonovska, I., and M. E. Waugh (2014): "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 92(1), 34–50.

Stellinger, A., I. Berglund, and H. Isakson (2020): "How trade can fight the pandemic and contribute to global health," *COVID-19 and Trade Policy: Why Turning Inward Won't Work*.

Stone, R. (1954): "Linear expenditure systems and demand analysis: An application to the pattern of British demand," *The Economic Journal*, 64(255), 511–527.

UNCTAD, and The World Bank (2018): "The Unseen Impact of Non-Tariff Measures: Insights from a New Database," Discussion paper.

United Nations Conference on Trade and Development (2020): "Trade Analysis Information System (TRAINS)," data accessed in April 2020.

United Nations Statistics Division (2020): "UN COMTRADE," data accessed in April 2020.

CEPR PRESS

Covid Economics Issue 21, 22 May 2020

Learning from friends in a pandemic: Social networks and the macroeconomic response of consumption[1]

Christos A. Makridis[2] and Tao Wang[3]

Date submitted: 18 May 2020; **Date accepted:** 19 May 2020

How do individuals adjust their consumption in response to information disseminated through peers and the social network? Using United States data on consumption, coupled with geographic friendship ties to measure social connectivity, this paper quantifies the role of social networks as a propagation mechanism for understanding aggregate fluctuations in consumption. Using the COVID-19 pandemic as a source of variation, we find that a 10% rise in cases and deaths in counties connected through the social network is associated with a 0.64% and 0.33% decline in consumption expenditures--roughly three to seven times as large as the direct effects of local cases or deaths. Counties more socially connected to epicenter countries of the pandemic also saw a bigger drop in consumption. These effects are concentrated among consumer goods and services that rely more on social contact, suggesting that individuals incorporate the experiences from their social network to inform their own consumption choices. We are working on incorporating this microeconomic evidence into a heterogeneous agent model and social interaction to study the aggregate demand implications.

1 We would like to thank Safegraph, Facteus and Facebook (especially Mike Bailey) for providing the access to the data used in this paper. We also thank Nick Bloom, Chris Carroll, and Laura Veldkamp for comments. These views are our own and do not reflect those of any affiliated institutions.

2 Arizona State University, MIT Sloan.
3 : John Hopkins University

Covid Economics 21, 22 May 2020: 100-133

1 Introduction

The ongoing COVID-19 pandemic represents the largest world-wide shock in at least a century, leading to substantial declines in employment (Bartik et al., 2020; Cajner et al., 2020), consumption (Baker et al., 2020b; Coibion et al., 2020), and output (Makridis and Hartley, 2020; Guerrieri et al., 2020). The vast majority of empirical contributions thus far have focused on the direct effects of the first moment shocks associated with the virus and the resulting national quarantine.[1]

The primary purpose of this paper is to explore the macroeconomic effects of the pandemic on consumption mediated through the presence of social networks. Since social networks are now a primary vehicle for obtaining information in the average household (Westerman et al., 2014), individuals may adjust their consumption in response to information communicated through friends in connected regions even if their own county has fairly low exposure to the virus. Quantifying how individuals make consumption and savings decisions in response to shocks to not only their fundamentals, but also those of their connected friends is important for understanding the sources of aggregate fluctuations, particularly during episodes of uncertainty and panic.

This paper quantifies an elasticity of the individual consumption and the composition of consumption in response to fluctuations in infections in connected counties. Using a combination of US micro card transaction data from Facteus and the Social Connectedness Index (SCI) from Facebook, we exploit plausibly exogenous variation in individuals' exposure to different geographically remote counties based on ties in their social network that were formed prior to the pandemic. We find that a 10% increase in SCI-weighted cases and deaths is associated with 0.64% and 0.33% decline in consumption, respectively. We also find that these declines are greater among social-

[1] See Baker et al. (2020a) for an exception.

contact-based consumption categories and activities away from home. For instance, each 10% increase in socially-connected cases is associated with a 2% decrease in clothing/footwear/cosmetics, a 1.3% decrease in contract-based service, and a 1.1% decrease in travel. These are twice to three times as large as the drop in average spending. This finding provide direct evidence for the uneven impacts of the pandemic on different sectors, which, augmented with market incompleteness, could result in a permanent drop in aggregate demand (Guerrieri et al., 2020).

We investigate the mechanisms and show that our results are not driven by time-varying shocks that are also correlated with infections in connected counties. First, we control for state $\times$ day fixed effects, which isolates variation across counties in the same state. As a placebo, we also show that increases in infections among connected counties do not have systematically different effects on consumption in states after the adoption of stay-at-home orders, which we would expect them to have if the results were driven by state-specific policies that simply curtailed foot traffic. Second, we conduct a wide array of heterogeneity exercises, showing that the heterogeneous treatment effects align with theory (e.g., greater effect in younger counties since social networks are more prevalent with millennials). Finally, we exploit counties' heterogeneous exposure to day-to-day changes in infections across South Korea, Italy, France, and Spain. Restricting our sample to February 15th to March 15th before the United States' national emergency was launched in full force, we find similar results, suggesting that our results reflect an information-driven response.

Our paper directly contributes to a large literature on the household response of consumption to macroeconomic shocks. This literature largely focuses on the impact of income volatility and borrowing constraints (Zeldes, 1989; Pistaferri, 2001; Gourinchas and Parker, 2002), stimulus (Di Maggio et al., 2017; Fuster et al., 2018), and tax rebates (Souleles, 1999; Johnson et al., 2006;

Covid Economics 21, 22 May 2020: 100-133

Covid Economics 21, 22 May 2020: 100-133

Agarwal et al., 2007) on consumption.[2] Quantifying how shocks affect consumption is important for understanding the presence of partial insurance and the pass-through of shocks (Blundell et al., 2008; Kaplan and Violante, 2010, 2014; Heathcote et al., 2014).[3] Our results provide evidence that individuals may adjust their consumption in response to shocks to affect their "friends" even beyond any direct effects on themselves. While we are not the first to point out the presence of peer effects in economic and financial behavior (Moretti, 2011; Bursztyn et al., 2014), our results nonetheless build on emerging literature that points to the real economic consequences of social networks, as in the case of renting versus owning in the housing market (Bailey et al., 2018a).

Our paper also contributes to an older literature on social externalities, specifically the spread of disease (Diamond and Maskin, 1979; Kremer and Morcom, 1998).[4] Recent research has began investigating the effects of pandemics on economic activity, placing a central role on the optimizing behavior of households. Whereas Eichenbaum et al. (2020) and Garibaldi et al. (2020) build models featuring households that fail to internalize the risk of infecting others through contact, Krueger et al. (2020) allow for heterogeneity in infection probabilities across sectors. This heterogeneity allows Krueger et al. (2020) to explore how individuals may optimally choose to avoid certain activities because of the risk of infection, meaning that the economy can avoid contagion without as severe of government intervention. Because we show how individuals adjust their consumption in response to information about the severity of the pandemic before stay-at-home orders were even implemented, we provide additional evidence that individuals make adjustments to their behavior in response to information. Our results are also consistent with Farboodi et al. (2020)

[2] Closely related is a larger literature on precautionary saving (Carroll, 1992, 1997; Carroll and Samwick, 1998) and the relationship between economic sentiment and consumption (Carroll et al., 1994).

[3] See Jappelli and Pistaferri (2010) for a survey.

[4] There is also a related literature in applied psychology and computational social science that has identified evidence of contagion through the dissemination of information through social networks, i.e. see Kramer et al. (2014), as well as Fowler and Christakis (2008) for some of the early evidence outside of social media.

who use the Safegraph data to quantify the effect of the pandemic on foot traffic.

Finally, our paper is related with an emerging empirical literature on the role of personal experience in expectation formation. Studies have highlighted the role of personal experience in forming beliefs about future returns (Cogley and Sargent, 2008), inflation (Malmendier and Nagel, 2016; Coibion and Gorodnichenko, 2015), energy prices (Binder and Makridis, 2020), housing prices (Kuchler and Zafar, 2019), macroeconomic activity (Malmendier and Nagel, 2011; Makridis, 2020; Makridis and McGuire, 2020), asset prices (Malmendier et al., 2018), political preferences (Giuliano and Spilimbergo, 2014), and consumption (Malmendier and Shen, 2018). By showing how consumption activity is linked with shocks that are diffused throughout the social network, our paper builds closely on Carroll (2003) who finds that household expectations are informed by news reports and the views of professional forecasters. If social networks amplify negative shocks by "spreading the bad news," they can potentially help account for the potentially persistent effect that the pandemic will have on expectations (Kozlowski et al., 2020a).[5] Binder (2020) also finds that individuals worried about the pandemic also have greater inflation expectations, suggesting that bad news about the pandemic spills over into the broader expectation formations process.

The structure of the paper is as follows. Section 2 describes our data and measurement approach. Section 3 introduces our identification strategy. Section 4 presents our main results and investigates heterogeneity across consumption goods. Section 5 investigates the prospective social networks mechanism. Section 6 concludes. We are continuing to produce additional results, specifically focusing on taking these empirical patterns to an aggregate model of the economy.

[5] Kozlowski et al. (2020a) builds a model along the lines of Kozlowski et al. (2020b) where transient shocks can have persistent effects on beliefs. If expectations are trending in one direction, but agents experience a large shock, beliefs can take time to recover.

Covid Economics 21, 22 May 2020: 100-133

2 Data and Measurement

The transaction-level data is provided by Safegraph and Facteus based on an anonymized panel of roughly 5.18 million debit card users' daily spending records between January 1st, 2017 to April 17th, 2020. Transactions are collected from primarily four types of cards providers across the United States: (1) bank debit cards whose majority users are young people; (2) general-purpose debit cards that are primarily distributed by merchants and retailers. (3) payroll cards used between employers and employees. (4) government cards. Average nationwide daily spending of the whole sample is 194 million dollars from a total of 2.3 million transactions.

Three features make the data particularly useful to our analysis.[6] First, there is rich geographic heterogeneity. In particular, transactions are partitioned by the residential zipcode of the card user. We then aggregate zip-level transactions into county-level consumption observations of 3051 counties (out of 3141 in the United States as of 2019). For zip zones that are associated with multiple counties, we allocate total consumption to its multiple corresponding counties based on its population weights. To ensure the county-level consumption is not biased by abnormal individual users' records and extreme values, we restrict our sample to include only county-day observations with more than 30 card users. Daily average consumption expenditures per card user is roughly $40 based on 0.5 transactions.

Second, there is high-frequency variation. In particular, we exploit the daily variation in transactions to identify the response of consumption to news about the pandemic. Since the

[6] However, one limitation of our data is that the location of a transaction differs from the location of residence; we only observe the latter. While we suspect that exploiting county-level (rather than zipcode-level) variation mitigates this concern, since people consume locally most of the time, we view potential misclassification as a source of measurement error (Chen et al., 2011). This would bias us against finding a result. We nonetheless conduct robustness where we investigate potential heterogeneous treatment effects in areas that have high versus lower levels mobility in "normal times", i.e. college towns.

Covid Economics 21, 22 May 2020: 100-133

epidemic crisis has eclipsed nearly ever other national and international event with release of daily news on the number of infections and deaths, daily records provide much cleaner variation than the common alternative of monthly data to recover the effects of news and social media. Although the transaction data goes back to 2017, we restrict the sample to the period between February 15th and April 17th with a total of 60 days and roughly 250,000 county $\times$ day observations. This period spans from the early spreading stage of the COVID-19 in Asian and European continents to the peak of the crisis within the United States. Depending on if the focus of analysis is domestic or international, we split our sample with the cutoff date March 15—a widely acknowledged watershed in nationwide response to the crisis in the country.

Finally, spending transaction is recorded by the merchant's type identified by its merchant classification code (MCC), a commonly adopted classification scheme by major card providers such as Visa/Mastercard. This allows us to study the consumption responses by category. We group each one of the 982 MCCs into 17 broad categories based on its degree of exposure to the infection risks, as well as its demand elasticity.[7] For instance, eating/drinking/leisure outside the home, contact-based service such as barbershop, and travel are expected to be most severely hit by the infection risk. Grocery and food shopping, financial services, and housing utilities, in contrast, are expected to have mild responses to the pandemic news during this period.

Figure 1 plots the average daily spending of each month since February 2020 by consumption category. The bulk of the consumption is accounted for by goods and services that are generally most exposed to the pandemic, including eating and drinking, leisure outside of the home, contact-based services, travel and transportation, and clothing, footwear, and cosmetics. However, some goods and services, such as financial services, grocery shopping, and home leisure, have actually

[7] See Section 6 for examples of merchant types that fall into each category.

CEPR PRESS

increased in March, relative to the two months prior. One important difference in the data, however, is that grocery shopping and other necessary purchases account for a large share in total spending, reflecting the fact that the composition of consumers in the sample is lower income and younger than a more nationally representative sample.

[INSERT FIGURE 1 HERE]

While the data contains these three important advantages over the traditional sources, we nonetheless are concerned about whether the data is nationally representative enough to map elasticities identified in the micro-data to the aggregate economy. We explore several validation exercises. First, Figure 2 plots monthly total spending based on our transaction records and the monthly total retail sales provided by the Census Bureau. To ensure that the series are as comparable as possible, we exclude from the card spending both financial services (e.g., insurance premiums and wire transfers) and housing rent and utilities, leaving the purchase of durable and non-durable goods and services. The two series track each other fairly well given that they are not apples to apples comparisons—the correlation is 0.41 over the three-year period that overlaps. Some of these differences may emerge because the sample selects lower-income individuals and does not have complete coverage throughout the country.[8]

[INSERT FIGURE 2 HERE]

In addition, we also compare the transaction series with subcategory of consumption separately reported by the Census Bureau. The positive correlation remains between our series with categories

[8]These low income and younger groups are widely known in the literature to have a high Engel index, i.e. a large share of spending on necessities such as grocery/food. That means the composition of the spending recorded in the transaction is geared toward basic items. Moreover, both low-income and young people tend to have a high marginal propensity to consume (MPC) due to under insurance. This will undoubtedly induce more volatility in consumption spending across different periods.

Covid Economics 21, 22 May 2020: 100-133

CEPR PRESS

Covid Economics 21, 22 May 2020: 100-133

such as grocery shopping, food and beverage stores, general merchandise and eating/drinking places, for which the correlation coefficients are all above 0.4.

Another major dataset used in this analysis measures the social network connectedness between different pairs of counties and between a U.S. county to different foreign countries. We combine these data with the Social Connectedness Index (SCI) from Facebook. Introduced by Bailey et al. (2018a) to study the role of social networks in propagating housing price shocks across space, these data are beginning to be used more widely to understand how social ties are related with economic activity (Bailey et al., 2018b). The index is constructed off of aggregated and anonymized information between all Facebook users, counting the number of friendship ties between county c and every other county c' in the United States. We use the 2019 data extract. Each user is limited to a total of 5,000 friends on a profile. Friendship ties require that both sides agree. We also draw upon the number of COVID-19 infections and deaths at the county $\times$ day level from the Center for Systems Science and Engineering from Johns Hopkins.[9]

3 Identification Strategy

While several emerging papers now document a substantial drop in consumption and its composition over the course of the pandemic (Baker et al., 2020b; Coibion et al., 2020), these studies focus on the direct effects of the national quarantine on spending patterns. However, since individual financial behaviors are also a function of peer effects (Moretti, 2011; Bursztyn et al., 2014), we investigate whether there is evidence of a decline in consumption prior to the national quarantine mediated through social networks. In particular, we draw on the Social Connectedness Index (SCI) to produce an SCI-weighted index of COVID-19 cases and deaths:

CEPR PRESS

$$COVID_{ct}^{SCI} = \sum_{c'}(COVID_{ct}^{d} \times SCI_{c,c'})$$

where $COVID_{ct}^{SCI}$ denotes the logged SCI-weighted number of cases or deaths in connected counties, $COVID_{ct}^{d}$ denotes the (direct) logged number of cases or deaths in the county. Importantly, we constructed Equation 3 by excluding ties between county c and all other counties c' within the same state to avoid potential mechanical effects between state-level policies and infections. Using this SCI-weighted index of the number of cases and deaths, we consider regressions of the following form that also control for local infections:

$$Y_{ct} = \gamma COVID_{ct}^{SCI} + \phi COVID_{ct}^{d} + \zeta_c + \lambda_t + \epsilon_{ct}$$

where y_{ict}^{k} denotes logged consumption for county c on day t for category-k consumption good, and ϕ and λ denote fixed effects on county and day-of-the-year. We cluster standard errors at the county-level to allow for arbitrary degrees of autocorrelation over time (Bertrand et al., 2004).

Equation 3 exploits plausibly exogenous variation in the exposure of an individual to counties that have more versus less severe COVID-19 shocks over time. For instance, whereas Maricopa County in Arizona has a scaled SCI with King County (Seattle) in Washington of 3,626, San Francisco in California has a scaled SCI of 12,294 with King County. Then, because Seattle was one of the hardest hit cities at first, we would expect that individuals in San Francisco would experience a greater drop in their consumption, relative to Maricopa County. Moreover, since we are controlling for the direct effects of COVID-19 in county c, we are exploiting only the variation that arises from social networks. We also explore potentially heterogeneous treatment effects.

4 Empirical Results

Covid Economics 21, 22 May 2020: 100-133

CEPR PRESS

Covid Economics 21, 22 May 2020: 100-133

4.1 Main Results

Table 1 documents the results associated with Equation 3. Starting with columns 1 and 4, we find that a 10% rise in SCI-weighted cases and deaths are associated with a 0.64% and 0.33% decline in consumption expenditures. Throughout all specifications, our SCI-weighted index excludes connections among counties within the same state as we focus on the effect through social networks instead of physical connect. Compared to column (1) and (4), columns (2), (3) (5), and (6) control for local (county) cases and deaths. While our coefficients decline slightly in magnitude, they remain statistically significant at a 1% level.

Importantly, the gradients on our SCI-weighted index of cases and deaths are roughly three to seven times as large as the direct effects of cases and deaths. This suggests that information transmitted through social networks might be even more quantitatively significant at informing consumption decisions than local activity does. Finally, recognizing the presence of time-varying state-specific containment policies, which are associated with meaningful effects on different activities such as job postings (Ali et al., 2020), we control for state $\times$ day fixed effects, isolating variation in counties' heterogeneous social networks even in the same state. Table 2 replicates these results using consumption growth and SCI-weighted growth in cases and deaths, producing almost identical results.

[INSERT TABLES 1 AND 2 HERE]

How do these information shocks potentially heterogeneously affect spending across different types of consumption goods? Figure 5 documents these results by reporting the coefficients associated with major categories of goods, which we created based on merchant category codes (MCC)

Covid Economics 21, 22 May 2020: 100-133

in the transaction data. We report the coefficients associated with both the direct effect of infections and the indirect effect through propagation from social networks. Not surprisingly, we find that clothing, footwear, and comsetic products decline the most, followed by contact-based services, durables, travel, and eating or drinking outside the home. For example, a 10% rise in SCI-weighted infections is associated with nearly a 3% decline in contact-based service spending, which is roughly three-times as large as the effects obtained on grocery / food or home leisure spending. These results are also consistent with Coibion et al. (2020) who find a 31 log point drop in consumer spending concentrated with a decline in travel and clothing.

[INSERT FIGURE 5 HERE]

4.2 Heterogeneous Treatment Effects Across Space

We now turn towards evidence of heterogeneity in the treatment effects by county characteristics. We control for the direct effects of county infections and deaths, focusing on variation in the SCI-weighted infections. We focus on per capita income, the age distribution, population, the share of digitally-intensive employees as defined by Gallipoli and Makridis (2018), and the share of teleworking employees as defined by Dingel and Neiman (2020). We partition each variable based on the median value, allowing for heterogeneity above and below the median. Our results with the digital and telework shares are both estimated on a restricted sample because we obtain them from the American Community Survey micro-data, which does not cover every county.

Table 3 documents these results. While not all the differences across different types of counties are statistically distinguishable from one another, they are consistent with theory. For example, a 10% rise in the SCI-weighted infections is associated with a 0.54% decline in consumption among

Covid Economics 21, 22 May 2020: 100-133

the counties below the median in per capita income, but a 0.46% decline among the rest. This could be consistent with the fact that counties with higher per capita income also have higher social capital and hygiene (Makridis and Wu, 2020). Turning towards heterogeneity in the age distribution, we distinguish among those counties that rank above and below the median in terms of the share of individuals below age 35 and the share of individuals above age 65.

Interestingly, we see that the effects are concentrated in counties that rank above the median share of individuals below age 35 and below the median share of individuals above age 65. This is consistent with the fact that younger individuals are more likely to pay attention to information from social media (Smith and Anderson, 2018). We also find that lower population counties have a nearly two-times as large of an elasticity, which could be explained by the fact that individuals in urban areas learn faster through their own surroundings. Finally, we see that counties with lower shares of digitally-intensive and teleworking employees are more adversely affected. Since both of these measures from Gallipoli and Makridis (2018) and Dingel and Neiman (2020) are measuring occupational tasks that cushion against the national quarantine—since digital services and teleworking (unlike, for example, hotels) are not directly affected by the national quarantine—we would expect to see that the counties with fewer of such workers being the ones that are harder hit.

[INSERT TABLE 3 HERE]

5 Understanding the Mechanisms

We have shown that there is an economically and statistically meaningful decline in consumption associated with increases in the number of COVID-19 infections in socially connected counties even

Covid Economics 21, 22 May 2020: 100-133

after controlling for time invariant characteristics across space and time, as well as time-varying shocks to local health outcomes (e.g., infections and deaths). However, one concern is that these results are plagued by other time-varying omitted variables that jointly affect connected counties and local consumption outcomes. This section provides further evidence that the results reflect a genuine information effect, rather than potential omitted variables.

One of the primary examples of omitted variables bias is the introduction of state-specific policies. For example, one possibility is that the introduction of emergency orders within a state naturally lead to declines in consumption by significantly disrupting foot traffic and leading to closures of businesses. While we show that our results are robust to controlling for state $\times$ day fixed effects, we nonetheless explore this possibility further by exploiting variation in the staggered introduction of state-specific stay-at-home orders (SAHOs) using data from Ali et al. (2020). If, for example, the introduction of SAHOs and other state policies account for the decline in consumption (Coibion et al., 2020), then we should see that the effect of the SCI-weighted infections loads on the interaction between it and the SAHOs. However, when we estimate these fixed effect specifications, we find a statistically insignificant point estimate of -0.002. This placebo counters the possibility remains that there are other unobserved and time-varying county-specific policies that vary with both consumption and connected counties.

We further investigate the role of social networks by turning towards measures of international exposure for each county, leveraging the fact that some countries began experiencing the surge in COVID-19 cases much sooner and more severely than the United States. We focus on four countries—South Korea, Italy, Spain, and France—although our results hold on a broader set of countries exposed early on.[10] Each of these four countries successively experienced large number

[10]Although we would, of course, ideally include China, the Facebook data does not have representative coverage

Covid Economics 21, 22 May 2020: 100-133

of infections in different scale since late February preceding the United States.

We exploit variation along two dimensions. First, counties vary cross-sectionally in their exposure to these countries. For example, whereas Maricopa County in Arizona has an SCI of 142,771 with France, San Francisco has an SCI of 258,825. Second, countries vary in their intensity of COVID-19 shocks. For example, Figure 6 shows how Italy experienced a sharper and more severe surge in cases than France even though its population is roughly 6 million smaller. We now consider regressions of logged consumption on the product of the cross-sectional exposure to a country and its time series variation in infections, conditional on the usual county and day fixed effects. Importantly, we restrict our sample to the period between February, 15th to March, 15th, which covers the time leading up to the full-scale outbreak in the United States.[11] This allows us purge variation that is possibly correlated with time-varying shocks in the United States.

Table 4 documents these results. We find that there is a robust negative association between the SCI-weighted number of infections / deaths and consumption for each country. For example, a 10% rise in infections (deaths) in Italy for counties that are more closely connected to Italy is associated with a 0.07% (0.52%) decline in consumption. One reason for the potentially larger coefficient on deaths over infections stems from the way that media covers international deaths more intensively than the number of infections, although we cannot say conclusively. We see broadly similar treatment effects for each country, although they are smaller for France, perhaps because the United States had already witnessed the experience of Asian countries, like South Korea, and Spain and Italy earlier in the month of March.

of ties with China because their government prohibits the use of Facebook.

[11] We also conduct the same analysis for the period after March, 15th for a different consideration. Since the Federal government of the U.S. announced the travel ban from Europe in the same week, focusing on this later period potentially shuts down the channel via which socially connected cases posed a real risk of infection. The negative impacts of consumption by SCI weighted cases from each of this country, if any, becomes more significant.

CEPR PRESS

Covid Economics 21, 22 May 2020: 100-133

[INSERT TABLE 4 HERE]

Our finding that consumption in one county depends in part on the effects of infections among connected counties—even if they are geographically distant—builds directly on an emerging literature on the real effects of social connectedness (Bailey et al., 2018b,a). However, separately identifying the causal effect of shocks to a network from selection effects is challenging (Goldsmith-Pinkham and Imbens, 2013). Our diagnostics—the combination of domestic and international connectivity—suggest that we are detecting meaningful effects from social networks, rather than just selection effects, but this remains an area of ongoing research. Our paper is also related with recent evidence from Charoenwong et al. (2020) that finds some counties were more likely to adopt social distancing and restrictions measures based on their exposure to Italy and China, although the data on social connectivity to China is confounded by the fact that use of Facebook Is blocked within the country.

6 Conclusion

The COVID-19 pandemic has led to substantial declines in employment (Bartik et al., 2020; Cajner et al., 2020), consumption (Baker et al., 2020b; Coibion et al., 2020), and output (Makridis and Hartley, 2020; Guerrieri et al., 2020), largely a function of the national quarantine policy. While the emerging empirical literature on the pandemic has focused on the direct effects of specific policies and/or the spread of the virus, this paper focuses on the role that social networks play in potentially propagating the effects on consumption. Using real-time data on consumption expenditures based on 5.18 millions debit card users' transaction, coupled with data on social connectivity across geographies from Facebook, we quantify the response of consumption to changes in a county's

COVID-19 exposure based on its social network. Our results suggest that these effects from the social network are significantly larger than the direct effects of the virus on consumption.

Covid Economics 21, 22 May 2020: 100-133

Tables and Figures

Figure 1: Descriptive Statistics on Consumption Expenditures, by Category

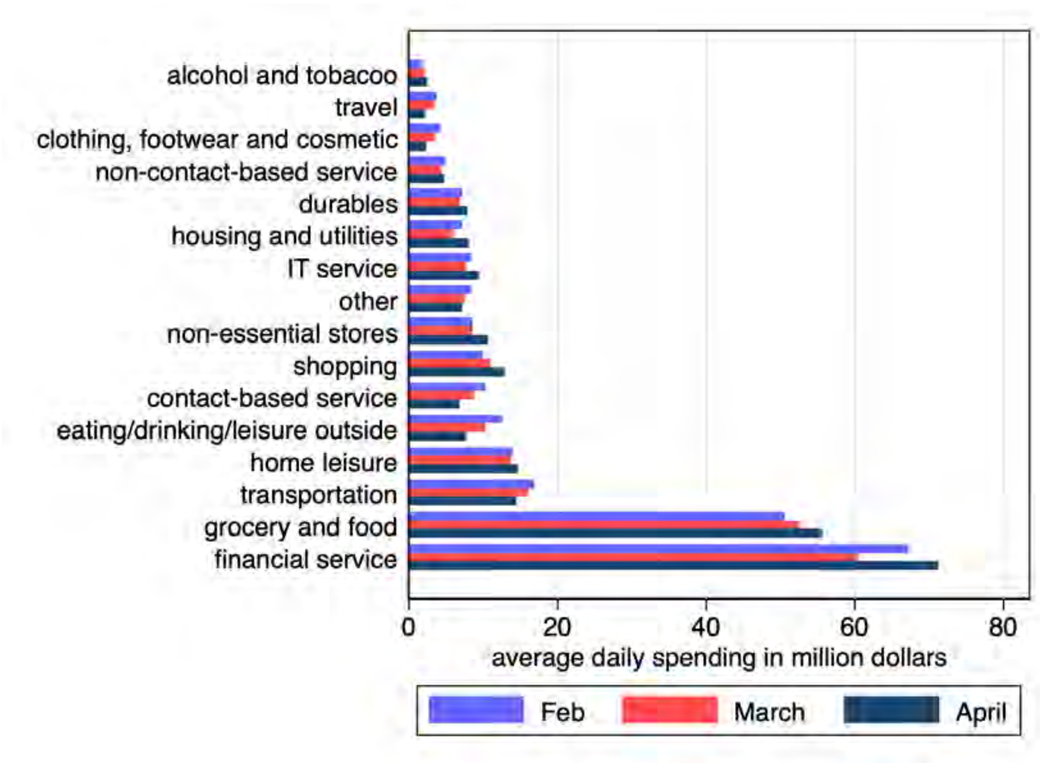

Notes.—Source: Facteus. Average daily consumption by category. Each bar plots the average spending in the specific category within each month. Data till the 17th is used for the average of April. See the Appendix for the examples of each consumption category.

Covid Economics 21, 22 May 2020: 100-133

Table 1: Consumption Responses to the COVID-19 Information on Facebook

Dep. var. =	log(spending)					
	(1)	(2)	(3)	(4)	(5)	(6)
log(SCI-weighted Cases)	-.064***	-.054***	-.067***			
	[.006]	[.006]	[.014]			
log(SCI-weighted Deaths)				-.033***	-.027***	-.054***
				[.003]	[.003]	[.008]
log(County Cases)		-.008***	-.006***		-.008***	-.006***
		[.001]	[.001]		[.001]	[.001]
log(County Deaths)		-.008***	-.007***		-.007***	-.006**
		[.002]	[.002]		[.002]	[.002]
R-squared	.97	.97	.98	.97	.97	.98
Sample Size	126106	126100	126086	126106	126100	126086
County FE	Yes	Yes	Yes	Yes	Yes	Yes
State/Time FE	No	No	Yes	No	No	Yes
Day FE	Yes	Yes	No	Yes	Yes	No

Notes.–Sources: Facebook, Facteus. The table reports the coefficients associated with regressions of logged consumption spending on logged SCI-weighted infections (excluding the counties in the same state) and logged county infections and deaths, conditional on county and time fixed effects. Standard errors are clustered at the county-level. The sample period is between March, 1st to April 17th, 2020

Table 2: Growth in Consumption and the COVID-19 News from Facebook

Dep. var. =	log spending growth					
	(1)	(2)	(3)	(4)	(5)	(6)
log(SCI-weighted Cases) growth	-.056***	-.053***	-.025			
	[.012]	[.012]	[.025]			
log(SCI-weighted Deaths) growth				-.014**	-.013**	-.039**
				[.006]	[.006]	[.016]
log(County Cases)		.005***	.001		.005***	.001
		[.001]	[.002]		[.001]	[.001]
log(County Deaths)		-.006***	-.005*		-.006**	-.005**
		[.002]	[.003]		[.002]	[.002]
R-squared	.33	.33	.42	.33	.33	.42
Sample Size	105606	105596	105589	105606	105596	105589
County FE	Yes	Yes	Yes	Yes	Yes	Yes
State/Time FE	No	No	Yes	No	No	Yes
Day FE	Yes	Yes	No	Yes	Yes	No

Notes.–Sources: Facebook, Facteus. The table reports the coefficients associated with regressions of logged consumption spending on week-to-week growth in SCI-weighted infections (excluding the counties in the same state) and logged county infections and deaths, conditional on county and time fixed effects. Standard errors are clustered at the county-level.The sample period is between March, 1st to April 17th, 2020

Covid Economics 21, 22 May 2020: 100-133

Figure 2: Benchmarking Consumption Expenditures with Retail Sales Over Time

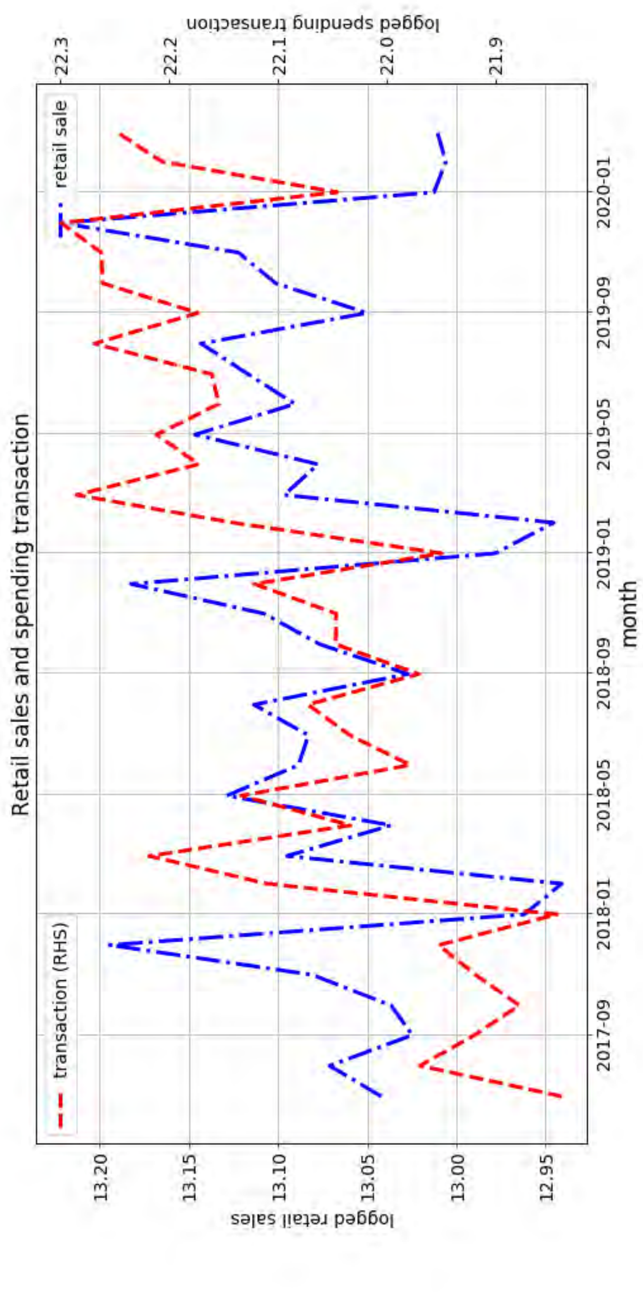

Notes.—Source: retail sales from the Census Bureau and transaction data from Facteus. Both are without seasonal adjustment and deflated by PCE price index. The total card transaction excludes financial-service-related such as insurance premium and wire transfer, as well as housing and utilities expenses. The retail sale uses the series "retail and food services, (total)", directly provided by the Census Bureau. The correlation coefficient of the two series is 0.41.

Figure 3: Actual and Socially-connected COVID-19 Case Infections

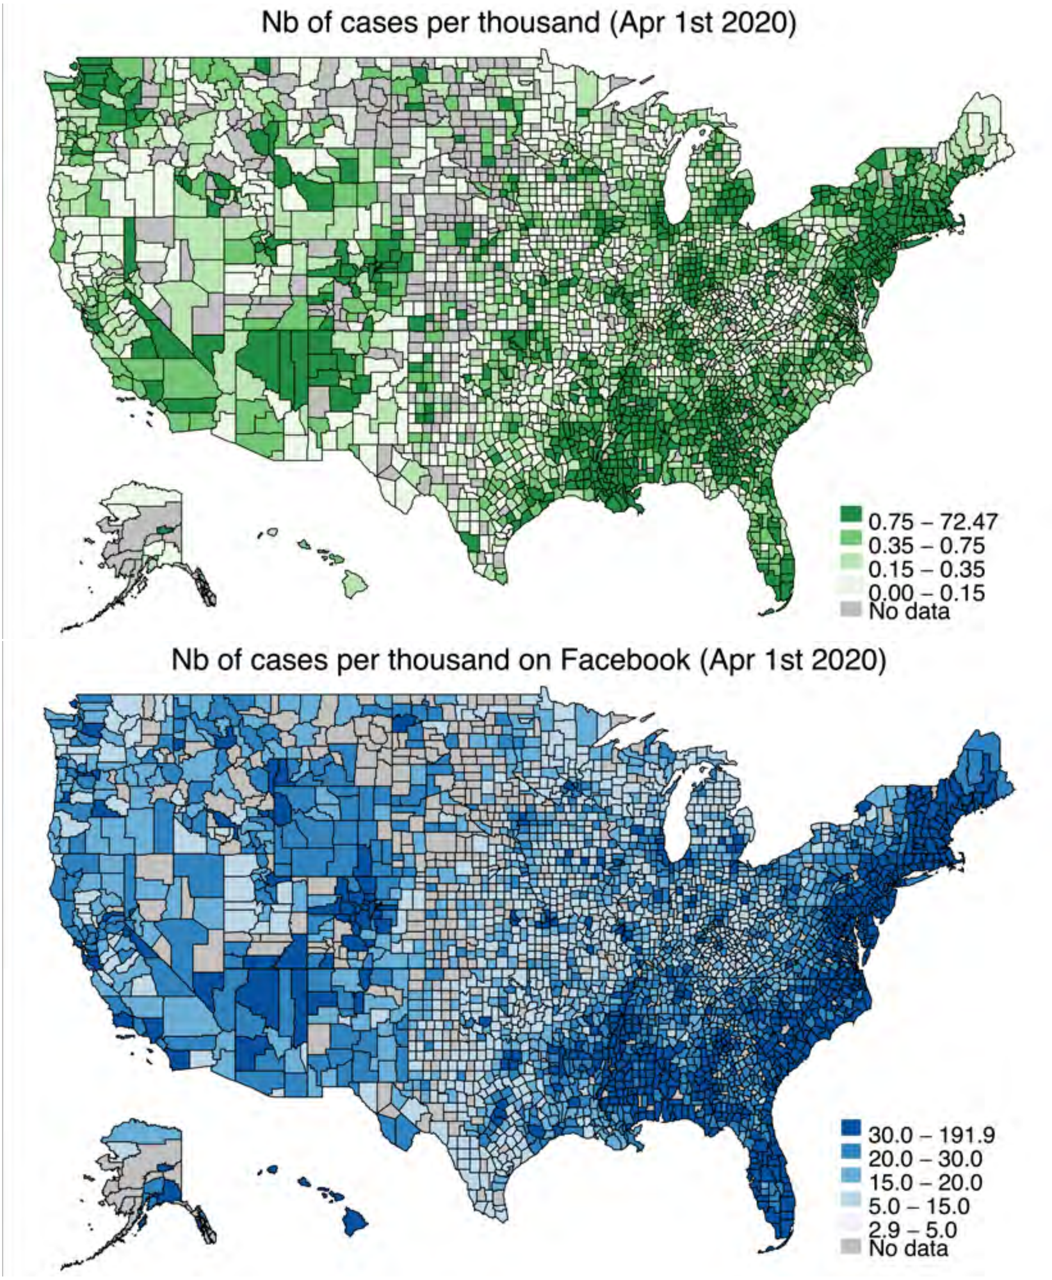

Notes.—Source: Facebook 2019 Social Connectedness Index (SCI). Panel A plots the number of COVID-19 infections per 1,000 individuals within each county as of April 1st, 2020. Panel B plots the SCI-weighted number of infections per 1,000 individuals, obtained by taking the population-weighted average across the product of infections in county c' and the SCI between county c and c'.

Covid Economics 21, 22 May 2020: 100-133

Figure 4: Actual and Socially-connected COVID-19 Deaths

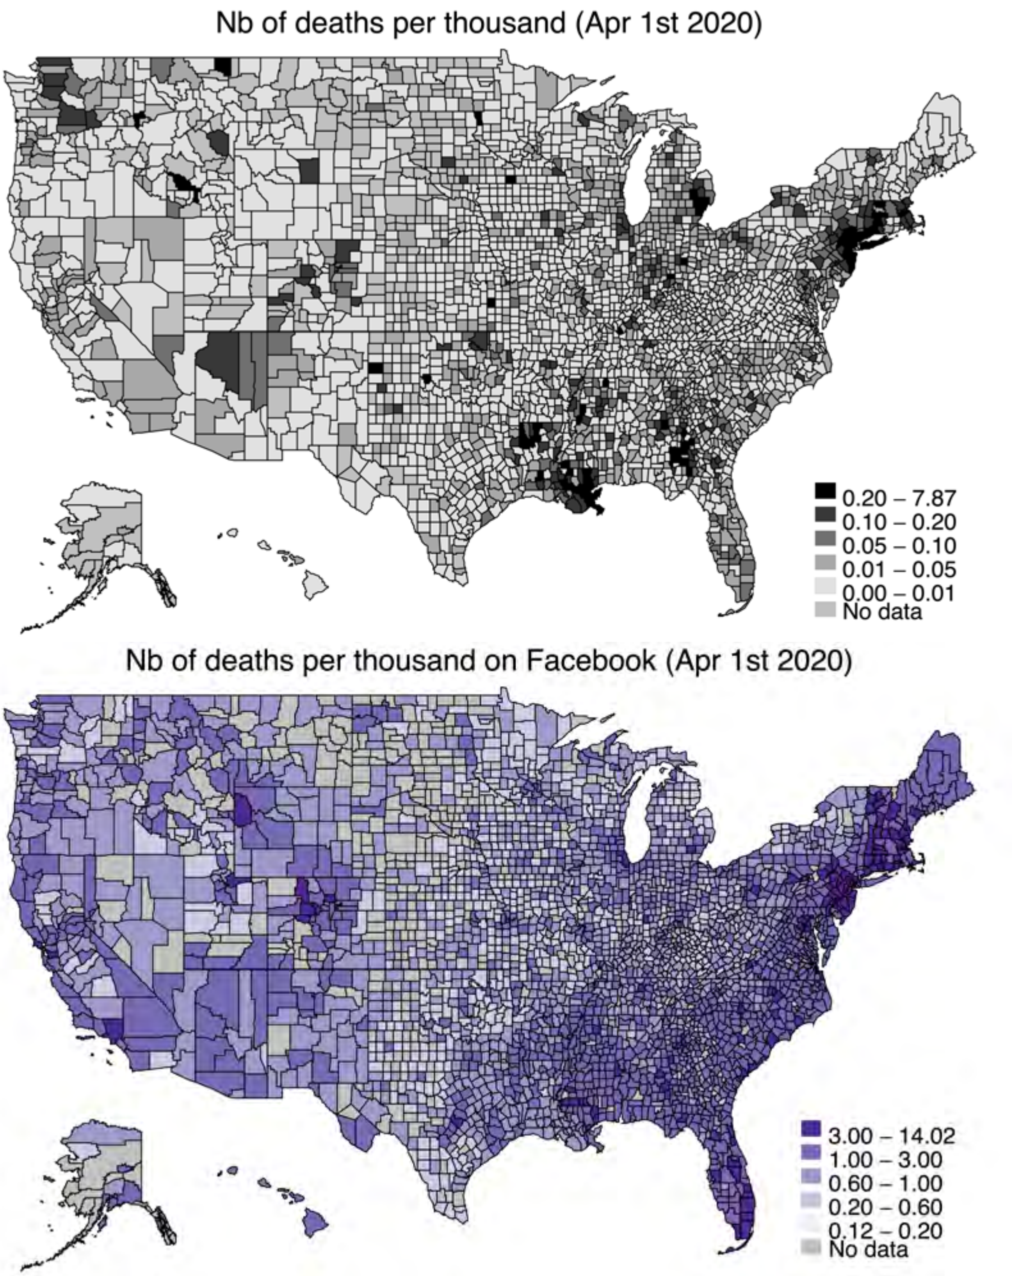

Notes.—Source: Facebook 2019 Social Connectedness Index (SCI). Panel A plots the number of COVID-19 deaths per 1,000 individuals within each county as of April 1st, 2020. Panel B plots the SCI-weighted number of deaths per 1,000 individuals, obtained by taking the population-weighted average across the product of deaths in county c' and the SCI between county c and c'.

CEPR PRESS

Covid Economics 21, 22 May 2020: 100-133

Figure 5: Consumption Response to COVID-19 Shocks, by Consumption Category

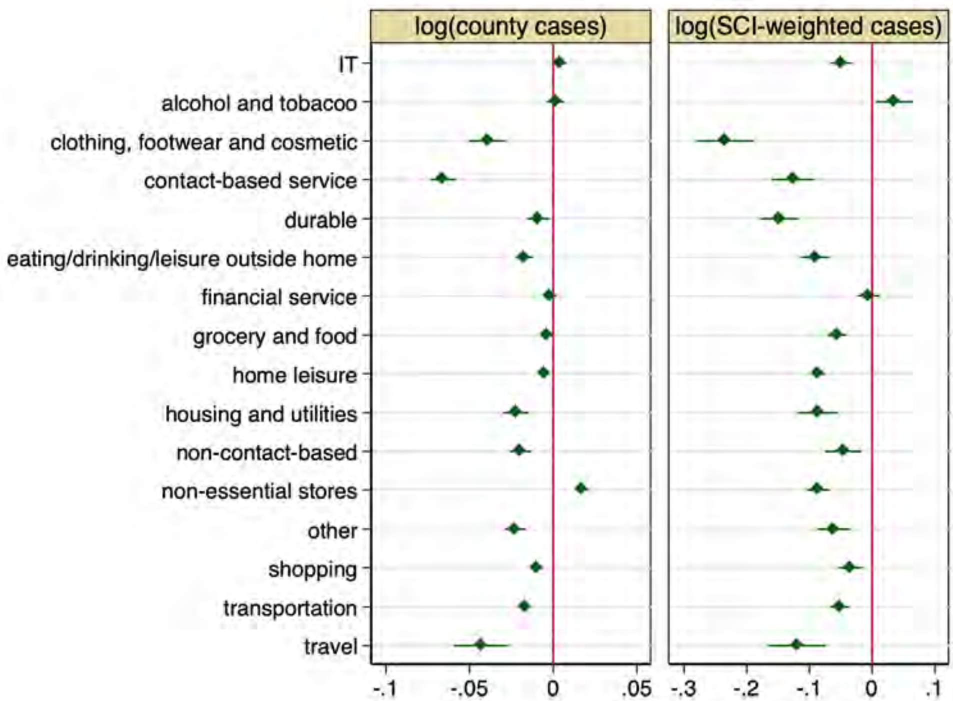

Notes.—Source: Facebook 2019 Social Connectedness Index (SCI) and Facteus. The figure reports the coefficients associated with regressions of logged consumption in a county on the logged number of COVID-19 cases (Panel A) and the logged number of SCI-weighted cases (Panel B) by category of consumption. Each transaction is classified as one of the following category based on its merchant category code (MCC).The sample period is between March, 1st to April 17th, 2020

Table 3: Heterogeneous Effects of the COVID-19 Information Shock on Consumption, by County Characteristics

RHS Variable Partition =	Per Capita Income		Share Under Age 35		Share Over Age 65		Population		Digital Intensity		Teleworking Intensity	
	High	Low	High	Low	High	Low	High	Low	High	Low	High	Low
log(SCI-weighted Cases)	-.046***	-.054***	-.070***	-.046***	-.046***	-.064***	-.038***	-.074***	-.019**	-.032***	-.021***	-.031***
	[.007]	[.009]	[.008]	[.009]	[.009]	[.007]	[.005]	[.012]	[.008]	[.010]	[.008]	[.010]
log(County Cases)	-.009***	-.009***	-.006***	-.006**	-.007***	-.005***	-.006***	-.013***	-.007**	-.003	-.006**	-.004
	[.002]	[.002]	[.002]	[.002]	[.003]	[.002]	[.001]	[.004]	[.003]	[.003]	[.003]	[.003]
log(County Deaths)	-.008***	-.004	-.005**	-.014***	-.005	-.008***	-.009***	-.005	-.007**	-.005	-.005	-.006
	[.002]	[.005]	[.002]	[.004]	[.006]	[.002]	[.002]	[.013]	[.003]	[.004]	[.003]	[.005]
R-squared	.98	.96	.98	.96	.95	.98	.98	.90	.98	.97	.98	.98
Sample Size	62235	63865	64855	61245	59862	66238	70597	55503	10741	9483	10360	9864
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes.–Sources: Facebook, Facteus, Census Bureau. The table reports the coefficients associated with regressions of logged consumption spending on logged SCI-weighted infections (excluding the counties in the same state) on logged county infections, conditional on county and time fixed effects. Standard errors are clustered at the county-level.

Figure 6: Time Series Patterns in COVID-19 Infections: Italy and France

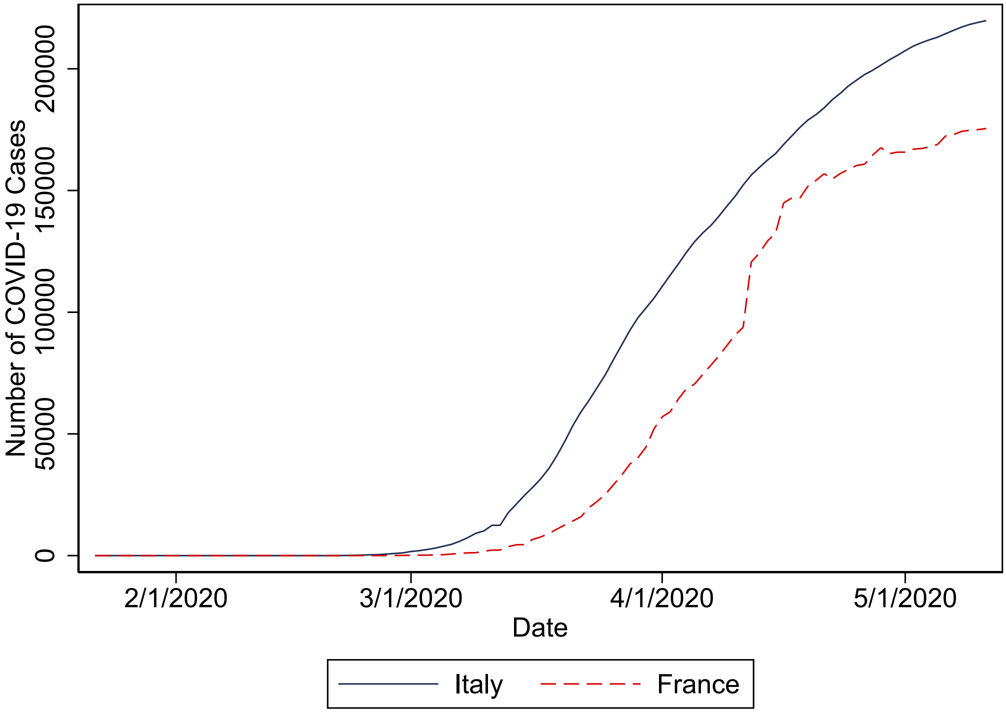

Notes.—Source: Johns Hopkins. The figure plots the number of COVID-19 infections for Italy and France over time.

Covid Economics 21, 22 May 2020: 100-133

Table 4: Consumption Responses to COVID-19 Information from Other Countries

Dep. var. =	log(spending)							
	ITA	ITA	SPA	SPA	FRA	FRA	SK	SK
log(SCI-weighted cases of the country)	-.007***		-.008***		-.011***		-.011***	
	[.001]		[.001]		[.001]		[.001]	
log(SCI-weighted deaths of the country)		-.052***		-.072***		-.014***		-.081***
		[.001]		[.001]		[.001]		[.002]
log(County Cases)	-.005	.015***	-.005	.003	-.005	-.005	-.005	.012***
	[.003]	[.004]	[.003]	[.004]	[.003]	[.003]	[.003]	[.004]
log(County Deaths)	-.004	-.025	-.004	-.019	-.004	-.004	-.004	-.025
	[.016]	[.018]	[.016]	[.018]	[.016]	[.016]	[.016]	[.018]
R-squared	.97	.98	.97	.98	.97	.97	.97	.98
Sample Size	78550	62925	78550	34148	78550	78550	78550	65552
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day FE	No	No	No	No	No	No	No	No

Notes.–Sources: Facebook, Facteus. The table reports the coefficients associated with regressions of logged consumption spending on logged SCI-weighted infections or deaths of a given foreign country, conditional on county and time fixed effects. These SCI-weighted infections / deaths are obtained by taking the time-varying number of infections in country i and multiplying it by the exposure of county c to country i, producing a Bartik-like measure. The four countries are Italy(ITA), Spain (SPA), France (FRA) and South Korea (SK). The sample period is between February 15th and March 15th, 2020. Standard errors are clustered at the county-level.

Covid Economics 21, 22 May 2020: 100-133

Online Appendix

Grocery and food. 1. grocery stores and super markets; 2. convenience stores; 3. drug stores and pharmacies; 4. miscellaneous retail stores; 5. meat provisions; 6. bakery, etc.

Transportation. 1. bus lines; 2. railway stations 3. car rentals; 4. toll and bridge fees, etc.

Home leisure. 1. TV cable fees; 2. digital goods, i.e. games, etc.

Housing and utilities. 1. housing rent payment; 2. home utilities, etc.

Shopping. 1. department stores; 2. discount stores; 3. variety stores; 4. general merchandise; 5. wholesale clubs, etc.

Eating, drinking, and leisure outside the home. 1. restaurants; 2. bars/taverns/clubs; 3. different kinds of parks; 4. outdoor sport and sports events; 5. orchestra and theaters, etc.

Information technology services. 1. computer network; 2. telegraph; 3. telecommunication, etc.

Contact-based services. 1. barber and beauty shops; 2. child care; 3. home cleaning; 4. repair stores; 5. veterinary services; 6. home furnishing; 7. laundry; 8. auto repair, etc.

Durables. 1.vehicles/motorcycle /auto parts; 2. furniture; 3. home appliances; 4. electronics and equipment; 5. home supplies; 6. music instruments, etc.

Non-contact-based services. 1. accounting/auditing; 2. business services; 3. programming; 4. consultations; 5. horticultural/ landscaping, etc.

Clothing, footware, and cosmetics. 1. clothing stores of different kinds; 2. cosmetic stores; 3. footwear and shoe stores, etc.

Alcohol and tobacco. 1. package stores selling wine, beer and other liquor; 2. cigar and tobacco stores, etc.

Travel. 1. airlines; 2. lodging and hotels; 3. duty-free stores; 4. airports; 5. travel agencies, etc.

Financial services. 1. insurance; 2. money orders; 3. wire transfers, etc.

Non-essential stores. 1.antique stores; 2. book stores; 3. art dealers, etc.

Other. 1. public organizations; 2. government fees; 3. educations; 4. medical spending such as a dental clinic, etc.

Covid Economics 21, 22 May 2020: 100-133

References

Agarwal, S., Liu, C., and Souleles, N. S. (2007). The reaction of consumer spending and debt to tax rebates: Evidence from consumer credit data. *Journal of Political Economy*, 115(6):986–1019.

Ali, U., Herbst, C. M., and Makridis, C. A. (2020). No More Home Alone: The Effects of Stay-at-Home Orders on the U.S. Child Care Market. *Working paper*.

Bailey, M., Cao, R., Kuchler, T., and Stroebel, J. (2018a). The economic effects of social networks: Evidence from the housing market. *Journal of Political Economy*, 126(6):2224–2276.

Bailey, M., Cao, R., Kuchler, T., Stroebel, J., and Wong, A. (2018b). Social connectedness: Measurement, determinants, and effects. *Journal of Economic Perspectives*, 32(3):259–280.

Baker, S. R., Bloom, N., Davis, S. J., and Terry, S. J. (2020a). COVID-induced economic uncertainty. *NBER working paper*.

Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., and Yannelis, C. (2020b). How Does Household Spending Respond to an Epidemic? Consumption During the 2020 COVID-19 Pandemic. *Working paper*.

Bartik, A. W., Betrand, M., Lin, F., Rothstein, J., and Unrath, M. (2020). Labor market impacts of COVID-19 on hourly workers in small- and medium- size businesses: Four facts from HomeBase data. *Chicago Booth Rustandy Center, Working Paper.*

Bertrand, M., Duflo, E., and Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*, 119(1):249–275.

Binder, C. (2020). Coronavirus Fears and Macroeconomic Expectations. *Review of Economics and Statistics*, forthcoming.

Binder, C. and Makridis, C. A. (2020). Stuck in the Seventies: Gas Prices and Macroeconomic Expectation. *Review of Economics & Statistics, R&R.*

Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921.

Bursztyn, L., Ederer, F., Ferman, B., and Yucghtman, N. (2014). Understanding mechanisms underlying peer effects: Evidence from a field experiment. *Econometrica*, 82(4):1273–1301.

Cajner, T., Crane, L., Decker, R. A., Grigsby, J., Hamins-Puertolas, A., Hurst, E., Kurz, C., and Yildirmaz, A. (2020). The U.S. labor market during the beginning of the pandemic recession. *BFI working paper.*

Carroll, C. (1992). The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings Papers on Economic Activity*, 23(2):61–156.

Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *Quarterly Journal of Economics*, 112(1):1–55.

Covid Economics 21, 22 May 2020: 100-133

Carroll, C. D. (2003). Macroeconomic Expectations of Households and Professional Forecasters. *Quarterly Journal of Economics*, 118(1):269–298.

Carroll, C. D., Fuhrer, J. C., and Wilcox, D. W. (1994). Does consumer sentiment forecast household spending? If so, why? *American Economic Review*, 84(5):1397–1408.

Carroll, C. D. and Samwick, A. A. (1998). How important is precautionary saving? *Review of Economics and Statistics*, 80(3):410–419.

Charoenwong, B., Kwan, A., and Pursiainen, V. (2020). Social connections with COVID-19 affected areas increase compliance with mobility restrictions. *Working paper*.

Chen, X., Hong, H., and Nekipelov, D. (2011). Nonlinear models of measurement errors. *Journal of Economic Literature*, 49(4):901–937.

Cogley, T. and Sargent, T. J. (2008). The market price of risk and the equity premium: A legacy of the Great Depression. *Journal of Monetary Economics*, 55(3):454–476.

Coibion, O. and Gorodnichenko, Y. (2015). Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. *American Economic Review*, 105(8):2644–2678.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). The cost of the COVID-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. *NBER working paper*.

Di Maggio, M., Kermani, A., Keys, B. J., Piskorski, T., Ramcharan, R., Seru, A., and Yao, V. (2017). Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review*, 107(11):3550–3588.

Covid Economics 21, 22 May 2020: 100-133

Diamond, P. A. and Maskin, E. (1979). An equilibrium analysis of search and breach of contract, I: Steady states. *Bell Journal of Economics*, 10(1):282–316.

Dingel, J. I. and Neiman, B. (2020). How many jobs can be done at home. *Covid Economics: Vetted and Real-Time Papers*, 1:16–24.

Eichenbaum, M., Rebelo, S., and Trabandt, M. (2020). The macroeconomics of epidemics. *NBER working paper*.

Farboodi, M., Jarosch, G., and Shimer, R. (2020). Internal and external effects of social distancing in a pandemic. *BFI working paper*.

Fowler, J. H. and Christakis, N. A. (2008). Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the Framingham Heart Study. *BMJ*, 337.

Fuster, A., Kaplan, G., and Zafar, B. (2018). Wht would you do with $500? Spending responses to gains, losses, news and loans. *NBER working paper*, Review of Economic Studies, R&R.

Gallipoli, G. and Makridis, C. (2018). Structural Transformation and the Rise of Information Technology. *Journal of Monetary Economics*, 97:91–110.

Garibaldi, P., Moen, E. R., and Pissarides, C. A. (2020). Modeling contacts and transitions in the SIR epidemics model. *Covid Economics, CEPR Press*, 5(16).

Giuliano, P. and Spilimbergo, A. (2014). Growing up in a Recession. *Review of Economic Studies*, 81(2):787–817.

Goldsmith-Pinkham, P. and Imbens, G. W. (2013). Social Networks and the Identification of Peer Effects. *Journal of Business & Economic Statistics*, 31(3):253–264.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):47–89.

Guerrieri, V., Lorenzoni, G., Straub, L., and Werning, I. (2020). Macroeconomic implications of COVID-19: Can negative supply shocks cause demand shortages. *Working paper*.

Heathcote, J., Storesletten, K., and Violante, G. L. (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review*, 104(7):1–52.

Jappelli, T. and Pistaferri, L. (2010). The consumption response to income changes. *Annual Review of Economics*, 2:479–506.

Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610.

Kaplan, G. and Violante, G. L. (2010). How much consumption insurance beyond self-insurance. *American Economic Journal: Macroeconomics*, 2(4):53–87.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Kozlowski, J., Veldkamp, L., and Venkateswaran, V. (2020a). Scarring body and mind: The long-term belief-scarring effects of COVID-19. *Covid Economics:Vetted and Real-Time Papers*, 8:1–26.

Kozlowski, J., Veldkamp, L., and Venkateswaran, V. (2020b). The tail that wags the economy: Beliefs and persistence stagnation. *Journal of Political Economy*, forthcoming.

Covid Economics 21, 22 May 2020: 100-133

Kramer, A. D. I., Guillory, J. E., and Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790.

Kremer, M. and Morcom, C. (1998). The effect of changing sexual activity on HIV prevalence. *Mathematical Biosciences*, 151(1):99–122.

Krueger, D., Uhlig, H., and Xie, T. (2020). Macroeconomic dynamics and reallocation in an epidemic. *NBER working paper.*

Kuchler, T. and Zafar, B. (2019). Personal experiences and expectations about aggregate outcomes. *Journal of Finance*, 74(5):2491–2542.

Makridis, C. (2020). Sentimental Business Cycles and the Protracted Great Recession. *Working paper.*

Makridis, C. A. and Hartley, J. (2020). The cost of COVID-19: A rough estimate of the 2020 GDP impact. *Mercatus Center, Policy Brief Special Edition.*

Makridis, C. A. and McGuire, E. (2020). Refined by Fire: The Great Depression and Entrepreneurship. *Working paper.*

Makridis, C. A. and Wu, C. (2020). Ties that Bind (and Social Distance): How Social Capital Helps Communities Weather the COVID-19 Pandemic. *SSRN working paper.*

Malmendier, U. and Nagel, S. (2011). Depression babies: Do macroeconomic experienes affect risk taking? *Quarterly Journal of Economics*, 126(1):373–416.

Covid Economics 21, 22 May 2020: 100-133

Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences. *Quarterly Journal of Economics*, 131(1):53–87.

Malmendier, U., Pouzo, D., and Vanasco, V. (2018). Investor experiences and financial market dynamics. *NBER Working Paper 24697*.

Malmendier, U. and Shen, L. S. (2018). Scarred consumption. *NBER working paper*.

Moretti, E. (2011). Social learning and peer effects in consumption: Evidence from movie sales. *Review of Economic Studies*, 78(1):356–393.

Pistaferri, L. (2001). Superior information, income shocks, and the permanent income hypothesis. *Review of Economics and Statistics*, 83(3):465–476.

Smith, A. and Anderson, M. (2018). Social media use in 2018. *Pew Center*.

Souleles, N. S. (1999). The response of household consumption to income tax refunds. *American Economic Review*, 89(4):947–958.

Westerman, D., Spence, P. R., and Van Der Heide, B. (2014). Social media as information source: Recency of updates and credibility of information. *Journal of Computer-Mediated Communication*, 19(2):171–183.

Zeldes, S. P. (1989). Consumption and liquidity constraints: An empirical investigation. *Journal of Political Economy*, 97(2):305–346.

Covid Economics 21, 22 May 2020: 100-133

Banks as lenders of first resort: Evidence from the Covid-19 crisis[1]

Lei Li,[2] Philip E. Strahan[3] and Song Zhang[4]

Date submitted: 20 May 2020; **Date accepted:** 20 May 2020

In March of 2020, banks faced the largest increase in liquidity demands ever observed. Firms drew funds on a massive scale from pre-existing credit lines and loan commitments in anticipation of cash flow disruptions from the economic shutdown designed to contain the Covid-19 crisis. The increase in liquidity demands was concentrated at the largest banks, who serve the largest firms. Pre-crisis financial condition did not limit banks' liquidity supply. Coincident inflows of funds to banks from both the Federal Reserve's liquidity injection programs and from depositors, along with strong pre-shock bank capital, explain why banks were able to accommodate these liquidity demands.

Covid Economics 21, 22 May 2020: 134-167

1 The results and conclusions of this paper represent those of the authors and do not represent the views of the Federal Reserve Board of Governors or any other institutions or persons affiliated with the Federal Reserve. Some data in this article were obtained through a confidential survey of depository institutions that requires confidential treatment of institution-level data and any information that identifies the individual institutions that reported the data.
2 Senior Economist, Federal Reserve Board of Governors.
3 John L. Collins Professor of Finance, Boston College.
4 Ph.D. candidate, Boston College.

Covid Economics 21, 22 May 2020: 134-167

I. INTRODUCTION

Firms go *first* to their bank(s) during a crisis. In the last three weeks of March 2020, anticipating disruptions to cash flow and external funding conditions, non-financial businesses drew funds from bank credit lines on an unprecedented scale. The shock occurred as asset prices had fallen across the capital markets. As a result, commercial and industrial (C&I) loans on bank balance sheets exploded, increasing by $482 billion between March 11 and April 1. For context, the weekly growth of bank C&I loans over the past 45 years has averaged 0.12% (with a standard deviation of 0.47%). Firms drew heavily on bank credit lines following the Lehman bankruptcy too, with lending increasing by about 6% during last three weeks of September 2008 and the first two weeks of October, or about 1.2% per week (about 10 times the average). In the last three weeks of March, however, lending grew more than 6% per week (or about 50 times the average). The growth in lending during these three weeks exceeded *every other* weekly growth rate going all the way back to 1973, when the Federal Reserve's H.8 releases (*Assets and Liabilities of Commercial Banks in the U.S.*) began (see Figure 1).[1]

[1] Calculations are the authors', based on the H.8 data for all commercial banks, not seasonally adjusted. See https://www.federalreserve.gov/releases/h8/current/h8.pdf

Figure 1: Weekly Growth in C&I Loans on Bank Balance Sheets: 1973-April, 2020

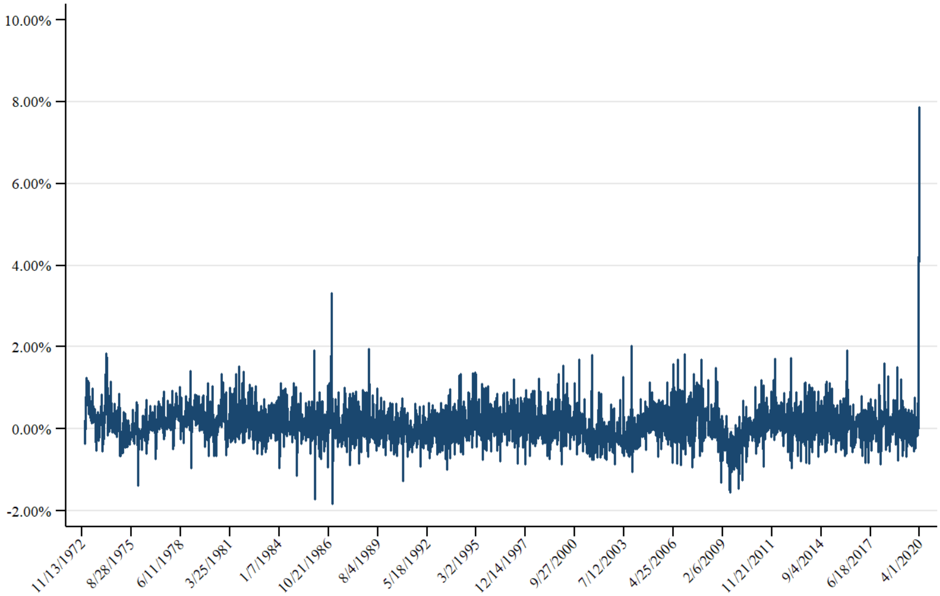

Figures are based on authors' calculations from data based on the H.8 data for all commercial banks, not seasonally adjusted.
See https://www.federalreserve.gov/releases/h8/current/h8.pdf

The three weeks in March 2020 are an unprecedented stress test on the ability of banks to supply liquidity. This 'stress test' was induced by the COVID-19 pandemic, which was unexpected to most firms and banks, was non-financial in nature, and affected all industries in the economy. In this paper, we study how bank characteristics and the characteristics of the markets in which they operate explain the cross-section of this explosion in lending. Our evidence suggests that all of the increase in lending occurred through drawdowns on existing credit commitments. Large banks experienced much more drawdowns than smaller ones. Anecdotal evidence suggests that the drawdowns came mainly from large firms, who typically borrow from the largest banks.[2] As

[2] For press accounts, see "Banks tolerate credit-line draws in coronavirus crisis — for now," *American Banker*, March 26, 2020 and "Credit-line drawdowns have peaked. Will banks get repaid?" *American Banker*, April 15, 2020.

Covid Economics 21, 22 May 2020: 134-167

a result, C&I lending grew much faster for banks with assets over $50 billion than for other banks (Figure 2). We do not find evidence that banks' financial condition before the onset of the crisis constrained their lending in this COVID-19 stress test. In other words, our results suggest that banks passed this stress test.

Figure 2: Cumulative C&I Loans Growth by Bank Size

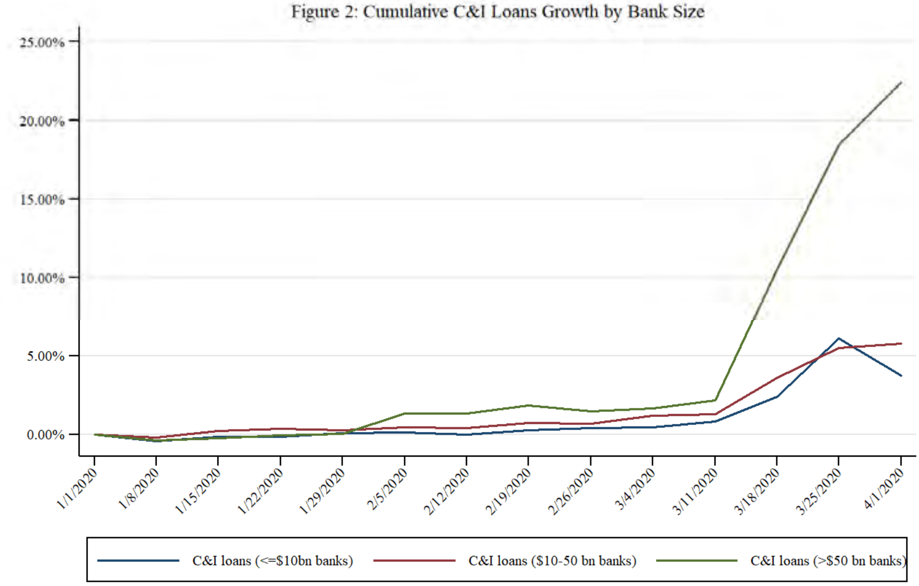

Figure 2 is based on authors' calculations from the FR 2644 data for all U.S. commercial banks, not seasonally adjusted. FR 2644 data require confidential treatment of institution-level data and any information that identifies the individual institutions that reported the data.

In particular, we study two questions in this paper: First, how has firm demand for bank liquidity responded to the onset of the COVID-19 crisis? Viral outbreaks varied substantially across localities, with some cities such as New York, New Orleans and Detroit facing major outbreaks and others facing more limited exposure. State and local governments responded differently to the pandemic, both in the intensity and in the timing of lockdowns and other measures aimed to slow the disease. For example, states like California initiated lockdowns early,

Covid Economics 21, 22 May 2020: 134-167

while other states such as Texas initiated these policies weeks later. We measure the size of local outbreaks using two strategies, one based on local employment declines in small firms, and the other by ex-post death rates from the COVID-19. We show much larger increases in lending at banks near large outbreaks.

Second, does bank financial condition constrain their ability to meet the unexpected increase in liquidity demands from their business clientele? To answer this question, we build off models developed by Cornett et al. (2011), who study a related phenomenon during the Financial Crisis of 2007–2009. Cornett et al. (2011) find that banks adjust to shocks to liquidity demands by reducing new credit origination, and that the changes in credit supply depend on bank financial constraints. Specifically, in the 2007-2009 period banks more reliant on core deposits, banks holding more liquid assets, and banks with more capital cut new lending less (increased lending more) than other banks. Our research tests whether or not these bank financial conditions have affected liquidity supply in response to the ongoing COVID-19 crisis.[3] One key difference between 2007-2009 and now is that the Federal Reserve has intervened more quickly and more massively this time. Another difference is that banks have amassed more capital due to the regulatory changes since 2008. Thus, we can assess whether these interventions helped alleviate the effect of bank financial constraints on lending during the COVID-19 crisis.

We show that the advent of the COVID-19 crisis explains the increase in lending, as banks located near areas with larger outbreaks experienced faster loan growth. Large banks with high levels of unused loan commitments to business experienced by far the largest increases in lending.

[3] Our model differs slightly from theirs in defining liquid assets. Cornett et al. define mortgage-backed and asset-backed securities as illiquid because these asset classes were at the center of the 2008 crisis. In this paper, we treat them the same as other securities (that is, we treat them as liquid assets).

Covid Economics 21, 22 May 2020: 134-167

These banks, however, were able to fund the liquidity demands due to the massive increase in deposits, which grew by about $1 trillion in aggregate during the crisis weeks, twice as much as the aggregate increase in lending. Our cross-bank regressions find no evidence that bank lending grew more at those financed more with stable deposits before the crisis. Similarly, we do not find that pre-crisis measures of asset liquidity explain increases in lending. These two 'non-results' suggest that concern about liquidity posed no constraint on banks, in stark contrast to what happened during the 2008 crisis. Moreover, we find no evidence that bank capital constrained their lending either. Again, in striking contrast to 2008, bank lending did not vary with capital in response to the COVID-19 crisis.

To develop our tests, we exploit two datasets: the quarterly *Call Reports* and weekly confidential FR 2644 data. The Q4, 2019 *Call Report* provides detailed measures of bank financial condition at the outset of the crisis, which we use to explain lending growth during Q1, 2020. Since lending exploded during the last three weeks of March (recall Figures 1 & 2), the bulk of the loan-growth variation during Q1, 2020 represents the effects of the liquidity shock during these three weeks. Hence, we exploit weekly growth in bank lending using the confidential FR 2644 data that underlie the Federal Reserve's H.8 releases. These are the only data that permit high-frequency analysis of the precise timing of the expansion of lending across banks. We construct all of our estimates within-bank (i.e., with bank fixed effects) to remove unobserved heterogeneity in bank lending patterns observed during normal (non-crisis) conditions.

Using the high frequency FR 2644 data, we control for lending patterns during normal periods based on the seven weeks between January 22 and March 11, after which lending took

Covid Economics 21, 22 May 2020: 134-167

off.[4] In a parallel set of tests, we validate our results using the lower frequency *Call Report* data, where we control for lending during normal periods based on the eight quarters leading up to the first quarter of 2020. *Call Report* data allow us to include all banks (rather than just the weekly reporting banks in FR 2644), and also allow us to model both on-balance sheet lending increases as well as total credit production (the sum of loans on balance sheet plus undrawn commitments). However, it does not allow us to pinpoint the exact timing of the liquidity shock. Our results across the two approaches are consistent, both in terms of statistical significance and economic magnitude.

Our identification strategy assumes no correlation between pre-crisis bank characteristics and the liquidity demand shock that occurs in March 2020. While we see no reason to assume otherwise – the pandemic and ensuing market panic was certainly a surprise to everyone – we show that our results are similar when we vary the set of variables capturing liquidity demand. In fact, our results are also similar (quantitatively) in models using the *Call Report* data, which omit local demand covariates.

Our paper contributes to the literature on banks' role as liquidity suppliers to firms. Earlier research suggests that combining deposits and off-balance sheet credit commitments creates diversification synergies, which allow banks to hold less cash (Kashap, Rajan and Stein, 2002). Gatev and Strahan (2006) argue that the synergy is especially powerful during periods of market stress because deposits flow into banks at the same time that borrower liquidity demands peak. Ivashina and Scharstein (2010) find evidence consistent with this latter mechanism during the 2008 crisis, although Acharya and Mora (2015) find that banks paid higher rates to attract deposits. In

[4] January 22 was the date of the first reported case of COVID-19 in the U.S. March 11 is the beginning of the rapid onset of loan drawdowns and corresponds with the World Health Organzation's declaration of a global pandemic.

Covid Economics 21, 22 May 2020: 134-167

this paper, we show that aggregate deposit inflows were more than enough to fund the increase in liquidity demands; these flows explain why pre-crisis deposits do not co-vary with lending across banks.

We also contribute to the emerging literature on the economic and financial consequences of the COVID-19 crisis. Many empirical papers study the stock market reaction to the pandemic, finding a strong response of equity prices to news about the virus and an increase in market volatility (Alfaro, Chari, Greenland, and Schott (2020); Baker, Bloom, Davis, Kost, Sammon, and Viratyosin (2020; Caballero and Simsek (2020)). Some studies compare how different types of stocks respond to the pandemic. Ding, Levine, Lin, and Xie (2020) find firms more exposed to the global supply chain fared worse, while Ramelli and Wagner (2020) find that exposure to international trade is also associated with poor stock price performance. Another set of studies focus on non-financial firms. Bartik, Bertrand, Cullen, Glaeser, Luca, and Stanton (2020), based on a survey small businesses, find a rapid onset of mass layoffs and concern by their surveyed firms about financial fragility. Several other authors study the early impact of the CARES Act and the Payroll Protection Program (Humphries, Neilson, and Ulyssea (2020); Granja, Makridis, Yannelis, and Zwick (2020); Cororaton and Rosen (2020)).

Like us, a number of studies focus on the effect of debt and liquidity on non-financial firms. Albuquerque et al. (2020) find stock returns at firms with high leverage ratios fared much worse during the crisis than those with less leverage, while Fahlenbrach, Rageth, and Stulz (2020) find that firms with more financial flexibility did better. De Vito and Gómez (2020) find that most firms would exhaust their cash holdings within two years, consistent with many firms relying on banks for liquidity. Our paper is most closely related to Acharya and Steffen (2020a), who document that access to bank credit lines during the COVID-19 crisis helped non-financial firms,

Covid Economics 21, 22 May 2020: 134-167

Covid Economics 21, 22 May 2020: 134-167

based on stock return patterns. Their paper studies the role of access to bank liquidity from the borrower perspective, whereas we study the problem from the bank (supply-side) perspective.

In another related note, Acharya and Steffen (2020b) apply models estimated from the pre-pandemic period to simulate the extent of credit line drawdowns at banks during the COVID-19 crisis. These models suggest that credit line usage increases when stock market returns are low (Berg et al. (2017)). Their analysis simulates aggregate loan drawdowns of $264 billion for the aggregate U.S. banking system, which is a little more than half of the increased lending seen in March 2020. They argue that banks overall are sufficiently well capitalized to accommodate this simulated demand for liquidity, but our data suggest the actual stress on banks has been substantially larger than the simulated one. Nevertheless, consistent with their conclusion, we find no evidence based on individual bank behavior that capital constrained their ability to meet this unprecedented demand for cash.

II. Empirical Methods, Data, and Results

Weekly Increases in Lending: Empirical Model & Data

The onset of the global COVID-19 virus pandemic initiated a market panic that led to a dramatic increase in firm drawdowns on existing credit lines. We exploit this increase in drawdowns in developing our empirical model, focusing on the three weeks from March 11 through April 1 as the period when liquidity demand spiked.[5] In our first empirical models, we

[5] We end the analysis on April 1 for several reasons. First, the rush to draw funds from pre-existing credit lines had abated by then. Second, we want the weekly analysis to be comparable to the analysis from the Q1, 2020 *Call Report* data. Third, government programs such as the Payroll Protection Program began in early April, which changed the main source of variation in bank lending and would require a different modeling approach.

Covid Economics 21, 22 May 2020: 134-167

use weekly data to construct the indicator variable *Crisis*, equal to one during these three weeks. To see how unusual this period is, Figure 3 illustrates the weekly growth in bank C&I loans from the beginning of 2020. The figure shows very clearly that these three weeks stand out from the early period. Moreover, the figure shows no unusual growth in other loans (e.g., real estate, consumer, etc.) during this time. Consistent with our interpretation, Acharya and Steffen (2020a) use data from S&P's *Loan Commentary and Data* to document that large public corporations drew about $225 billion on their bank lines during this period, which is only half of the increase in bank C&I lending in our data (about $480 billion). The difference probably reflects drawdowns by private firms.

Figure 3: Weekly Loan Growth, all Commercial Banks

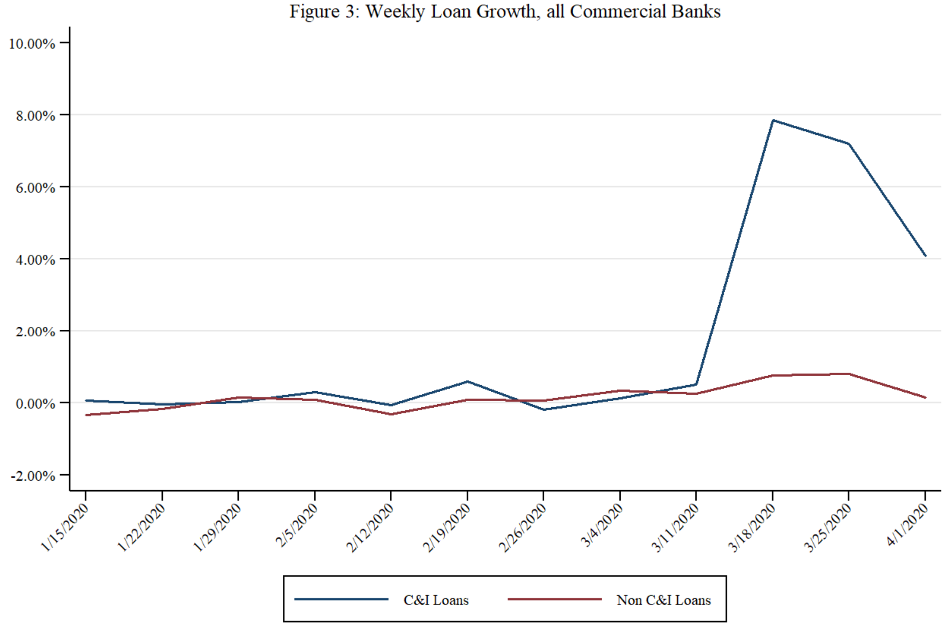

Figures are based on authors' calculations from data based on the H.8 data for all commercial banks, not seasonally adjusted. See https://www.federalreserve.gov/releases/h8/current/h8.pdf

We use this unexpected shock to liquidity demand to study whether bank financial condition affected their willingness to supply liquidity. Having such a bright-line increase in liquidity demand is nearly unique; it allows us to trace out how (or whether) financial condition constrained banks' ability to supply liquidity. The weeks after Lehman's bankruptcy in 2008 offer the only other similar situation. Cornett et al. (2011) show that both liquidity and capital affected bank liquidity supply then. We report similar tests, though as discussed in the introduction the increase in drawdowns during the COVID-19 crisis dwarfs that observed during the most intense weeks of the 2008 crisis.

Using the shock to liquidity demand in March 2020, we estimate models of weekly bank lending of the following form:

$$\Delta C\&I\ Loans_{i,t}/A_{i,Q4,2019} = \alpha_i + \beta^0 Crisis_t + \Sigma\beta^j Crisis_t{*}Bank\ Financial\ Condition^j_{i,Q4,\ 2019} + \Sigma\gamma^k Local\ Demand\ Conditions^k_{i,t} + \varepsilon_{i,t}\ . \quad (1)$$

The outcome in Equation (1) represents the weekly change in C&I lending from the Federal Reserve's FR 2644 dataset for bank *i* in week *t*, scaled by the bank's total assets from the end of the prior quarter.[6] We include the weeks from January 22 to April 1, 2020 for all of the domestic reporting banks, and set *Crisis* to one during the last three weeks of the sample. The FR 2644 data come from an authorized random stratified sample of weekly reporting banks.[7] We include a bank

[6] Results are similar using weekly loan growth as the outcome, although this variable contains large outliers. We prefer to normalize by beginning of period assets to eliminate the influence of outliers.

[7] The Federal Reserve reports the weekly aggregated balance sheet of U.S. banks at its website: https://www.federalreserve.gov/releases/h8/current/default.htm. We use the micro-data, which underlie these aggregates and were obtained through a confidential survey of depository institutions that requires confidential treatment of institution-level data and any information that identifies the individual institutions that reported the data.

fixed effect, α_i, in all of our models to remove bank-level heterogeneity, and cluster standard errors at the bank level throughout. Some of our tests also incorporate time effects.

We build j measures of bank financial condition based on the 2019 year-end bank *Call Reports*. As such, these measures are plausibly exogenous with respect to the COVID-19 pandemic (and any associated market panic), which was not declared by the World Health Organization until March 11, 2020 (which coincides with the beginning of our *Crisis* weeks). The bank financial measures vary only across banks (not over time), so the bank fixed effects fully absorb their direct effects in Equation (1). The β^j coefficients measure the impact of these conditions on lending during the three-week period in which firms were drawing down their credit lines, relative to their effects during the normal weeks that preceded it.

We include the following bank-level variables, from the Q4, 2019 bank *Call Reports*: 1) *Size*, equal to the log of total bank assets; 2) *Liquid asset*s, equal to non-interest bearing balances + interest-bearing balances + Federal funds sold + Repurchase agreement + held-to-maturity securities (at amortized cost) + available for sale securities (at fair value); 3) *Core deposits* (a measure of funding liquidity), equal to deposits in domestic offices minus deposits over $250,000; 4) *Tier 1 capital*; and, 5) *Unused commitments*, equal to undrawn commitments to business.[8] We normalize each of the on-balance sheet measures (other than log of assets) by total assets; for *Unused commitments*, we normalize by the sum of assets plus unused commitments. If banks are constrained by their asset liquidity, by the availability of stable funds, or by scarce capital, we would expect those factors to affect lending growth positively. Lending growth could be

[8] Our measure of capital is close to the regulatory Tier 1 Leverage ratio, although we use the year-end total assets rather than average total assets as the denominator for consistency with the other variables in the model.

Covid Economics 21, 22 May 2020: 134-167

constrained either by reducing new loan originations or by restricting access to liquidity under existing lines (as occurred during the 2008 crisis). In contrast, if banks hold substantial liquidity buffers, if their funding is sufficiently abundant, and if they operate sufficiently far from regulatory minimum capital ratios, then the effects of the pre-crisis financial conditions may not affect lending growth.

To capture bank-specific variation in exposure to local demand conditions, we incorporate two strategies. First, we control for the weekly growth in employment, as measured by total hours worked, at small firms located in the state in which each bank is headquartered.[9] These data come from *Homebase*, a software provider for small businesses to track employee working hours for scheduling and payroll. As of January 2020, the *Homebase* data cover about 60,000 small businesses across all 50 states. About 90 percent of their clients have fewer than 100 employees. *Homebase* covers only a small fraction of total state-level employment, but has concentration in the leisure, hospitality and retail trade sectors, which are the sectors most hard-hit by the COVID-19 crisis. Because this measure comes from very small firms, it is unlikely to be directly affected by the drawdown behavior, which was dominated by large firms (i.e., it acts as an exogenous measure of local exposure to the virus; employment patterns at large firms might be affected by the availability of liquidity, which is our outcome). As a second strategy, we measure the state-level COVID-19 deaths per capita through the beginning of May 2020, which gives a comprehensive, cross-sectional measure of the extent of the viral outbreak.[10] The bank fixed effect

[9] Alternatively, we use the average weekly growth in employment across states in a bank's branch network, weighted by the bank's deposits in each state, and find similar results.

[10] We use death after the end of our sample for two reasons. First, unlike the number of cases, deaths do not depend on the level of testing, which varies substantially across regions. Hence, it suffers less from measurement error. Second, deaths are a severely lagging indicator of the extent of the outbreak, so the total number of deaths in May represents a better measure of the magnitude of the viral outbreak in late March than would contemporaneous measures.

Covid Economics 21, 22 May 2020: 134-167

captures the cross-sectional effect of this variable, so we focus on its interaction with the *Crisis* indicator.[11]

Panel A of Table 1 reports summary statistics for the full sample of domestic, weekly reporting banks in FR 2644. We report the distribution for both the change in lending and weekly growth in hours worked split based on *Crisis*. Panels B-D of Table 1 report the data sorted into three size bins: banks with assets below $10 billion; banks with assets between $10 and $50 billion; and banks with assets above $50 billion.

Table 1 shows the dramatic shift that occurs during the crisis weeks. Pre-crisis, bank lending increases by only 0.01% of assets per week (with a standard deviation of 0.22% of assets); during the crisis weeks, bank lending increases by an order of magnitude more (to 0.1% of assets), while its standard deviation increases by about 70% (to 0.37%). These figures appear 'small' only because we normalize the change in C&I lending by the size of each bank's balance sheet from the end of the prior quarter. In fact, the changes in bank C&I lending during these three weeks are the largest since 1973 (when the FR 2644 data collection began).

Table 1 also shows that as lending explodes, small firm employment declines precipitously. During the pre-crisis period, weekly hours grew 1.07% per week but then falls by 17.23% during the crisis period (again, more than an order of magnitude more). The negative effect of the viral outbreak on local economic activity coincides with the explosion of bank lending because firms, anticipating future declines in cash flow, immediately draw funds from their bank credit lines when the effects of the pandemic become evident.

[11] Deaths come from Center for Systems Science and Engineering (CSSE), Johns Hopkins University. State population comes from World Population Review.

Covid Economics 21, 22 May 2020: 134-167

Panels B-D split the summary statistics based on bank size. This split shows that the largest banks faced, by far, the greatest increase in liquidity demand. Again, our measure normalizes the change in lending by the size of the lender's balance sheet, so the difference in lending in absolute terms across banks of different sizes is even more striking than at first glance (recall Figure 2). To be specific, the large banks experienced lending increases of 0.57% of assets per week (Panel B). In contrast, the small banks experienced lending increases of just 0.06% of assets (Panel D), and the medium-sized banks experience lending growth of 0.13% of assets (Panel C). Across all three bank-size bins, lending grew much faster during the crisis weeks, but this increase is most striking at the largest banks.

Weekly Loan Growth: Linking Lending to Bank Financial Conditions

Tables 2 and 3 report regressions from the weekly FR 2644 data, as in Equation (1). Table 2 reports pooled models, with all of the domestic reporting banks. Some specifications include only the loan-demand variables (*Crisis*, two lags of weekly growth in hours worked, their interactions with *Crisis*, and the state-level death rate from COVID-19 interacted with *Crisis*). We then report specifications that add the pre-crisis bank financial variables (*Size*, *Liquid assets*, *Core deposits*, *Tier 1 capital*, and *Unused commitments*) interacted with *Crisis*. We omit time effects from these models so the effects of *Crisis* and its interactions are well-identified (rather than being absorbed).

Table 3 then splits the analysis by bank size (<$10 billion in assets, $10-50 billion, and >$50 billion). In Table 3, we add the time effects to absorb fully the aggregate changes in lending patterns and thus allay concern about possible omitted variables related to liquidity demand. The

Covid Economics 21, 22 May 2020: 134-167

Covid Economics 21, 22 May 2020: 134-167

models in Table 3 allow us to draw inferences about the impact of banks' ex ante financial condition on liquidity supply.

As shown in Table 2, all three types of demand control have strong explanatory power for loan growth. From column (1): lending increases much faster during the crisis period (consistent with summary statistics). From column (2): the increase during the crisis period is greater in states with larger (lagged) declines in weekly hours worked at small firms. The first lag interacts negatively and significantly with *Crisis*, and the two lags are jointly statistically significant (p-value < 0.02). Similarly, the increase in lending is greater during the crisis weeks in states with more overall death per capita from COVID-19 (columns 4-6).

Table 2, columns 3, 5 and 6 introduce the pre-crisis bank-level variables. Consistent with the simple summary statistics, *Size* enters positively and significantly. Large banks face greater liquidity demands during the crisis weeks. Raising *Size* by one standard deviation (=1.9) increases weekly lending by about 0.04% of assets. The *Unused commitments* variable is by far the strongest predictor of weekly lending increases. This result validates the premise of the paper, which is that the increase in lending comes primarily from liquidity demands in which businesses draw funds from their pre-existing credit lines once the effects of the pandemic become clear. The economic magnitude of this variable is large: a standard deviation increase in *Unused commitments* (=0.04) comes with an increase in lending of 0.14% of assets, or about one-third of a standard deviation of lending growth during the crisis weeks (=0.37%; see Table 1, Panel A). We also find that both *Asset liquidity* and *Tier 1 capital* correlate positively with lending. These results suggest that financial condition may have constrained bank lending (although both variables are strongly correlated with bank size).

Covid Economics 21, 22 May 2020: 134-167

Table 3 separates the analysis by bank size. Two things stand out from this sample split. First, the effect of *Unused commitments*, while positive across all three samples, is much larger and more statistically significant for the largest banks. For banks with assets over $50 billion, a standard deviation increase in *Unused commitments* (=0.07; see Table 1, Panel B) leads to an increase in lending of 0.48% of assets (=0.07 x 0.068). This increase equals about two-thirds of a standard deviation of the outcome during the crisis week (=0.72% of assets; see Table 1, Panel B). Second, once we separate the sample by size, the effects of the three financial condition measures – *Liquid assets, Core deposits, and Tier 1 capital* – lose statistical power. We find no effect for the largest banks for any of these measures. For the medium-sized banks, *Tier 1 capital* enters with a marginally significant positive coefficient; for the small banks, *Asset liquidity* enters with a marginally significant positive coefficient. Taken together, these results offer almost no support for the idea that bank financial conditions constrained their ability to supply liquidity during the crisis weeks.

Quarterly Increases in Lending: Empirical Model & Data

The results from the weekly FR 2644 reporting banks have two limitations. First, we cannot separate new loan originations from drawdowns on pre-existing loan commitments because the change in loans on bank balance sheet equals the sum of net drawdowns on existing credit lines plus new originations. We have argued that the variation we observe in these data reflect drawdowns (not originations), but we cannot demonstrate that claim directly because we do not observe the off-balance sheet changes in the weekly data. Second, the FR 2644 sample does not include all banks.

To remedy these two defects, we look next at quarterly *Call Report* data. We have seen that the vast bulk of C&I lending changes during the Q1, 2020 come during the three crisis weeks in March. Hence, the total quarterly changes in lending (both on and off the balance sheet) will provide a good measure of banks' response to the crisis. So, we estimate the following two equations:

$$\Delta C\&I\ Loans_{i,t}/Assets_{i,t-1} = \gamma_t + \alpha_i + \Sigma\mu^j Bank\ Financial\ Condition^j_{i,t-1}$$

$$+ \Sigma\beta^j Crisis_t{*}Bank\ Financial\ Condition^j_{i,t-1} + \varepsilon_{i,t}\ , \qquad (2a)$$

and:

$$\Delta(C\&I\ Loans{+}Undrawn\ Commit.)_{i,t}/(Assets{+}Undrawn\ Commit.)_{i,t-1} = \gamma_t + \alpha_i + \Sigma\mu^j Bank$$

$$Financial\ Condition^j_{i,t-1} + \Sigma\beta^j Crisis_t{*}Bank\ Financial\ Condition^j_{i,t-1} + \varepsilon_{i,t}\ . \qquad (2b)$$

Equation (2a) is similar to Equation (1). We model the change in total C&I lending on the balance sheet of bank *i* in quarter *t*, normalized by lagged total assets. The sample includes all domestic banks and uses the eight quarters of 2018 and 2019 to pin down the 'normal' effect of bank condition on lending prior to the onset of the pandemic.[12] The time fixed effects, γ_t, capture the overall demand in each quarter and thus capture the overall shock observed in Q1, 2020. We leave out the location-specific measures of demand because these are not well-defined during the pre-crisis quarters. To construct standard errors, again we cluster by bank.

Unlike Equation (1), the bank fixed effects, α_i, do not fully absorb the 'normal' effects of the financial variables, as these exhibit within-bank variation over time. Hence, we include them in the regression, with their effect prior to the crisis captured by the μ^j coefficients. As in Equation

[12] In contrast, the weekly analysis uses the weeks of Q1, 2020 before the liquidity spike as the 'control' regime.

Covid Economics 21, 22 May 2020: 134-167

Covid Economics 21, 22 May 2020: 134-167

(1), the β^j coefficients capture the differential effect of bank financial condition during the crisis quarter relative to normal times. If bank financial conditions constrain their ability to accommodate the liquidity demand shock from the pandemic, the β^j coefficients will enter Equation (2a) with positive and significant effects. Equation (2b) allows us to estimate similar models using total credit production – the sum of on- and off-balance sheet lending commitments to businesses.

Table 4 reports summary statistics for the *Call Report* data, again looking at all banks first and then at banks in each of three size bins. We report these statistics separately for the crisis and pre-crisis quarters. Looking at the full sample (Panel A), on-balance sheet lending does not show any difference between Q1, 2020 and the earlier quarters. This masks very large differences, however, in the aggregates because the large banks were much more affected than the smaller ones. Consistent with the weekly data, large banks experienced much faster loan growth in Q1, 2020 than during the earlier quarters (Panel B). For them, loans grew by 0.2% of assets at the mean of the distribution prior to the crisis; in Q1, 2020, however, C&I loans grow 1.7% of assets at the mean. This difference, however, is *not* evident in total credit production (loans on balance sheets plus undrawn commitments), which is slightly lower during the crisis quarter for all banks and slightly higher for the largest banks. These patterns support our claim that credit-line drawdowns dominate changes in lending during the crisis.[13] Firms demand liquidity from pre-existing credit lines, as opposed to demanding new credit to facilitate growth or new investment. All else equal, each dollar drawdown leads to a dollar increase in loans on bank balance sheets but no change in total credit (=loans + undrawn commitments).

[13] The increase in lending was also much too massive and abrupt to have been driven by new loan originations, which require substantial time for negotiation of pricing and contract terms.

CEPR PRESS

Covid Economics 21, 22 May 2020: 134-167

Quarterly Increases in Lending: Results

Tables 5 & 6 report the estimates of Equations (2a) and (2b) using the nine quarters from the beginning of 2018 through the first quarter of 2020.

The results in Table 5 are fully consistent with those from Table 3, despite the fact that Table 5 uses all banks (rather than a subset), uses a different pre-crisis benchmark (all of 2018 and 2019, rather than the first weeks of Q1, 2020), and has no cross-state control for the level of the viral outbreak.

In particular, as in Table 3, the growth in lending in Q1, 2020 is best explained by the level of pre-existing business loan commitments. The effect of pre-existing commitments is much larger for the largest banks, and the magnitudes line up very closely with those from the weekly analysis. For example, the coefficient on *Unused commitments* for the largest banks equals 0.068 at weekly frequency (Table 3, column 2). This coefficient is 'turned on' for three weeks during Q1, 2020, so the total effect on lending over the quarter equals 3x0.068 = 0.204. This effect is very close to the coefficient estimated on the interaction of the *Crisis* x *Unused Commitments* from Table 5 of 0.195 (column 2). Moreover, as with the weekly data, there is very little evidence that bank financial conditions affect changes in lending. We estimate a positive effect of *Crisis x Tier 1 capital* for the small banks (column 4), but this effect is only marginally significant.

Table 6 reports estimates of Equation (2b), where the outcome captures total credit production to businesses. This variable is not affected by credit line drawdowns, only by overall changes in credit originations. We find no correlation between *Crisis x Unused Commitments* and credit production, either for the full sample or for any size-based subsample. Consistent with Table 5, there is little evidence that this broader measure of credit production is constrained by bank financial condition. Again, we see a marginally significant positive coefficient on *Crisis x Tier 1 capital* for the small banks, as well as marginally significant positive coefficient on *Crisis x Core deposits.* Neither of these, nor asset liquidity, has a significant effect on total credit production for the medium-sized or large banks.

Discussion

We have analyzed how banks have accommodated the unprecedented increase in liquidity demands in response to the ongoing COVID-19 pandemic. Pre-existing unused loan commitments explain the majority of the variation in lending, especially for large banks. Yet we see almost no correlation between bank financial strength and their willingness to bear this liquidity shock. The shock is the largest ever observed, going all the way back to 1973. It is larger than anything observed during the 2008 crisis, when financial condition of banks *did* constrain lending. How is this possible? We suspect that changes in the regulatory regime after the 2008 crisis, inflows of deposits, and the Federal Reserve's aggressive actions in response to the pandemic, explain our results.

Increases in bank liquidity has been a mechanical side effect of the massive expansion of the Federal Reserve's balance sheet from Quantitative Easing (QE), as it effectively expands the supply of excess reserves. The Fed began expanding the supply of reserves in September of 2019, and then announced a massive expansion of QE in response to the pandemic on March 15, 2020. At the same time, the Fed also expanded and reinstated lending programs to banks and other large financial institutions constructed during the 2008 crisis. The policy moves were timely – expanding liquidity supply from the central bank just as non-financial firms were drawing liquidity from their banks.

Covid Economics 21, 22 May 2020: 134-167

At the same time that liquidity supplied by the Fed expanded, it also expanded from private sources. Deposits flowed rapidly into banks, again at just the same time that liquidity demands spiked. Figure 4 graphs the aggregate flow of deposits into U.S. banks from the beginning of March 2020, compared to the increase in total C&I lending. As the figure shows, for every $1 of new lending, there is about $2 of additional deposits. Deposits increased by almost $1 trillion during the three weeks from March 11 to April 1. Hence, liquidity poured into banks at exactly the right time, both from the public sector and the private sector. This coincident increase in liquidity supply to banks, when most needed by their borrowers, is consistent with earlier episodes (e.g., Kashyap, Rajan and Stein, 2002; Gatev and Strahan, 2006).

Figure 4: Increases in C&I Loans and Deposits (Billions of $)

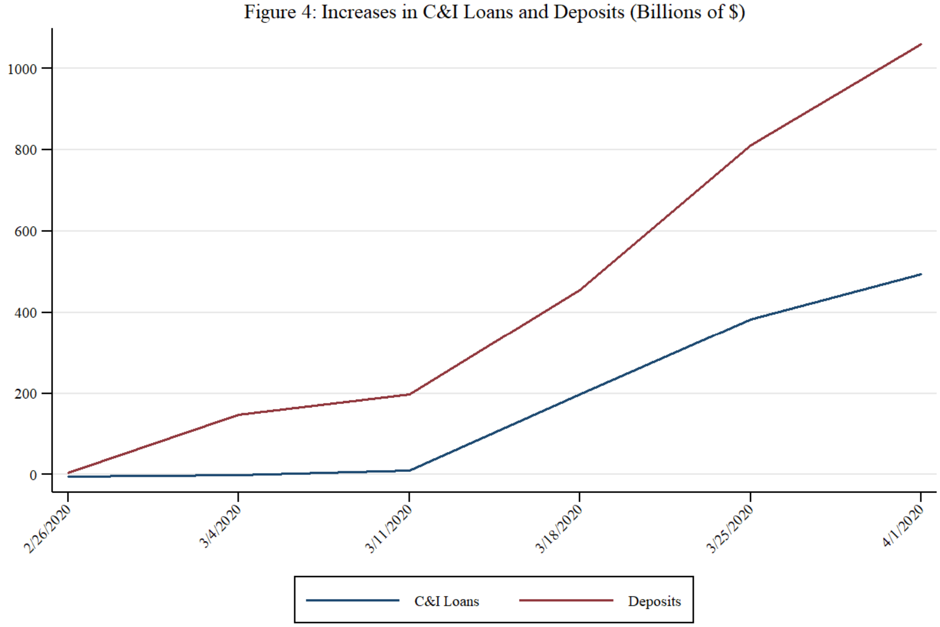

Figures are based on authors' calculations from data based on the H.8 data for all commercial banks, not seasonally adjusted. See https://www.federalreserve.gov/releases/h8/current/h8.pdf

Covid Economics 21, 22 May 2020: 134-167

Covid Economics 21, 22 May 2020: 134-167

Our results also suggest that capital did not constrain banks during this COVID-19 crisis. As we show, the largest banks faced by far the greatest increases in liquidity demands. But these banks also experienced the greatest increase in regulatory capital from the post-2008 changes in regulation. Innovations such as stress testing and additional capital buffers required for the systemically important financial institutions (SIFIs) moved the largest banks well above minimum capital requirements (Schneider, Yang, and Strahan, 2020). These regulatory innovations, while controversial, seem to have succeeded in building a sufficiently thick capital cushion to allow banks to bear this liquidity shock. That said, the longer-term effect of the increased credit exposure on bank solvency is impossible to assess until more time has passed.

III. CONCLUSION

We have shown that liquidity demands on the large U.S. banks reached unprecedented levels during late March 2020. Firms went to their banks for cash, drawing funds from pre-existing credit lines and loan commitments, in anticipation of massive declines in future cash flow from the oncoming economic shutdown. Large banks experienced the lion's share of these liquidity demands. Banks met the demand without running into binding financial constraints. We suggest two reasons for this. First, bank liquidity and bank solvency buffers were both substantially more robust before the COVID-19 crisis than they were before the 2008 crisis. Second, aggregate liquidity supply, from both the Federal Reserve and from depositors, flowed in at exactly the right time.

References

Acharya, V.V. and Steffen, S., 2020a. 'Stress tests' for banks as liquidity insurers in a time of COVID. CEPR COVID Economics.

Acharya, V.V. and Steffen, S., 2020b. The risk of being a fallen angel and the corporate dash for cash in the midst of COVID. VOX, CEPR Policy Portal.

Acharya, V.V. and Mora, N., 2015. A crisis of banks as liquidity providers. The journal of Finance, 70(1), pp.1-43.

Albuquerque, R.A., Koskinen, Y., Yang, S. and Zhang, C., 2020. Love in the Time of COVID-19: The Resiliency of Environmental and Social Stocks.

Alfaro, L., Chari, A., Greenland, A.N. and Schott, P.K., 2020. Aggregate and firm-level stock returns during pandemics, in real time (No. w26950). National Bureau of Economic Research.

Baker, S.R., Bloom, N., Davis, S.J., Kost, K.J., Sammon, M.C. and Viratyosin, T., 2020. The unprecedented stock market impact of COVID-19 (No. w26945). National Bureau of Economic Research.

Bartik, A.W., Bertrand, M., Cullen, Z.B., Glaeser, E.L., Luca, M. and Stanton, C.T., 2020. How are small businesses adjusting to COVID-19? Early evidence from a survey (No. w26989). National Bureau of Economic Research.

Berg, T., Saunders, A., Steffen, S. and Streitz, D., 2017. Mind the gap: The difference between US and European loan rates. The Review of Financial Studies, 30(3), pp.948-987.

Cornett, M.M., McNutt, J.J., Strahan, P.E. and Tehranian, H., 2011. Liquidity risk management and credit supply in the financial crisis. Journal of financial economics, 101(2), pp.297-312.

Cororaton, A. and Rosen, S., 2020. Public Firm Borrowers of the US Paycheck Protection Program. Available at SSRN 3590913.

De Vito, A. and Gómez, J.P., 2020. Estimating the COVID-19 cash crunch: Global evidence and policy. Journal of Accounting and Public Policy, p.106741.

Ding, W., Levine, R., Lin, C. and Xie, W., 2020. Corporate Immunity to the COVID-19 Pandemic (No. w27055). National Bureau of Economic Research.

Fahlenbrach, R., Rageth, K. and Stulz, R.M., 2020. How valuable is financial flexibility when revenue stops? Evidence from the COVID-19 crisis. Fisher College of Business Working Paper, (2020-03), p.007.

Gatev, E. and Strahan, P.E., 2006. Banks' advantage in hedging liquidity risk: Theory and evidence from the commercial paper market. The Journal of Finance, 61(2), pp.867-892.

Covid Economics 21, 22 May 2020: 134-167

Granja, J., Makridis, C., Yannelis, C. and Zwick, E., 2020. Did the Paycheck Protection Program Hit the Target? (No. w27095). National Bureau of Economic Research.

Humphries, J.E., Neilson, C. and Ulyssea, G., 2020. The evolving impacts of COVID-19 on small businesses since the CARES Act.

Ivashina, V. and Scharfstein, D., 2010. Bank lending during the financial crisis of 2008. Journal of Financial economics, 97(3), pp.319-338.

Kashyap, A.K., Rajan, R. and Stein, J.C., 2002. Banks as liquidity providers: An explanation for the coexistence of lending and deposit - taking. The Journal of finance, 57(1), pp.33-73.

Ramelli, S. and Wagner, A.F., 2020. Feverish stock price reactions to covid-19.

Schneider, T.I., Strahan, P.E. and Yang, J., 2020. Bank Stress Testing: Public Interest or Regulatory Capture? (No. w26887). National Bureau of Economic Research.

Covid Economics 21, 22 May 2020: 134-167

CEPR PRESS

Covid Economics 21, 22 May 2020: 134-167

Table 1: Summary Statistics for Weekly Lending

This table reports summary statistics for the weekly change in commercial and industrial (C&I) lending by banks and hours worked by small firms from January 22 through April 1, 2020. Lending data are from the Federal Reserve FR 2644 data and hours are from Homebase. In addition, we report bank characteristics for banks in the sample as of Q4, 2019 from the bank Call Reports. All variables, except Assets and Log(Assets), are winsorized at 1%.

Panel A: All Banks	N	Mean	Std Dev
	(1)	(2)	(3)
Full Sample (January 22 - April 1, 2020)			
Bank Assets (in $thousands)	8,234	21,029,000	141,830,000
Unused C&I comm./(Unused C&I comm.+Assets)	8,234	0.05	0.04
Weekly Change in C&I Loans/Assets	8,234	0.0004	0.0029
Liquid assets/Assets	8,234	0.27	0.14
Core deposits/Assets	8,234	0.77	0.08
Tier 1 capital/Assets	8,234	0.11	0.03
%Change in weekly hours	8,234	-5.60	12.64
State COVID death per capita as of May 4 (in p.p.)	8,234	0.02	0.03
Log(Assets)	8,234	14.15	1.90
Pre-Crisis (January 22 - March 10, 2020)			
Weekly ΔC&I/Assets	5,235	0.0001	0.0022
%Change in weekly hours	5,235	1.07	2.23
Crisis (March 11 - April 1, 2020)			
Weekly ΔC&I/Assets	2,999	0.0010	0.0037
%Change in weekly hours	2,999	-17.23	14.72

CEPR PRESS

Covid Economics 21, 22 May 2020: 134-167

Table 1: Summary Statistics for Weekly Lending (Cont'd)

Panel B: Large Banks (>$50 billion)	**Pre-Crisis (January 22 - March 10, 2020)**			**Crisis (March 11 - April 1, 2020)**		
	N	Mean	Std Dev	N	Mean	Std Dev
	(1)	(2)	(3)	(4)	(5)	(6)
Unused C&I comm./(Unused C&I comm.+Assets)	270	0.11	0.07	156	0.11	0.07
Weekly Change in C&I Loans/Assets	270	0.0001	0.0014	156	0.0057	0.0072
Liquid assets/Assets	270	0.31	0.16	156	0.31	0.16
Core deposits/Assets	270	0.75	0.09	156	0.75	0.09
Tier 1 capital/Assets	270	0.09	0.02	156	0.09	0.02
%Change in weekly hours	270	0.84	1.94	156	-16.96	14.97
State COVID death per capita as of May 4 (in p.p.)	270	0.03	0.04	156	0.03	0.04
Log(Assets)	270	19.00	0.97	156	19.00	0.97
Panel C: Medium-Sized Banks ($10-50 billion)	**Pre-Crisis (January 22 - March 10, 2020)**			**Crisis (March 11 - April 1, 2020)**		
	N	Mean	Std Dev	N	Mean	Std Dev
	(1)	(2)	(3)	(4)	(5)	(6)
Unused C&I comm./(Unused C&I comm.+Assets)	511	0.06	0.04	292	0.06	0.04
Weekly Change in C&I Loans/Assets	511	0.0002	0.0015	292	0.0013	0.0027
Liquid assets/Assets	511	0.22	0.10	292	0.22	0.10
Core deposits/Assets	511	0.74	0.09	292	0.74	0.09
Tier 1 capital/Assets	511	0.10	0.02	292	0.10	0.02
%Change in weekly hours	511	0.93	2.14	292	-17.80	14.96
State COVID death per capita as of May 4 (in p.p.)	511	0.02	0.03	292	0.02	0.03
Log(Assets)	511	16.85	0.43	292	16.85	0.43

CEPR PRESS

Table 1: Summary Statistics for Weekly Lending (Cont'd)

Panel D: Small Banks (<=$10 billion)	**Pre-Crisis (January 22 - March 10, 2020)**			**Crisis (March 11 - April 1, 2020)**		
	N	Mean	Std Dev	N	Mean	Std Dev
	(1)	(2)	(3)	(4)	(5)	(6)
Unused C&I comm./(Unused C&I comm.+Assets)	4,454	0.04	0.04	2,551	0.04	0.04
Weekly Change in C&I Loans/Assets	4,454	0.0001	0.0023	2,551	0.0006	0.0033
Liquid assets/Assets	4,454	0.28	0.14	2,551	0.28	0.14
Core deposits/Assets	4,454	0.78	0.07	2,551	0.78	0.07
Tier 1 capital/Assets	4,454	0.11	0.03	2,551	0.11	0.03
%Change in weekly hours	4,454	1.10	2.25	2,551	-17.19	14.68
State COVID death per capita as of May 4 (in p.p.)	4,454	0.02	0.03	2,551	0.02	0.03
Log(Assets)	4,454	13.54	1.24	2,551	13.55	1.24

Covid Economics 21, 22 May 2020: 134-167

Table 2: Explaining Weekly Lending Growth

This table reports panel regressions of the weekly change in bank C&I loans from January 22 to April 1, 2020 (from the Federal Reserve's FR 2644 data) on bank financial condition measures from Q4, 2019 *Call Reports*. *Crisis*is an indicator variable equal one for the weeks between March 11 and April 1. We capture local credit demand with *Crisis* and its interactions with state-level growth in hours worked by small firms (from *Homebase*), along with the extent of death from COVID-19. F-stat and p-value at the bottom of the table report tests on whether coefficients on *Crisis * %Change in weekly hours (t-2 to t-1)* and *Crisis * %Change in weekly hours (t-3 to t-2)* sum up to zero. Standard errors are clustered at bank level. T-statistics are reported in parentheses. '*' denotes significance at the 10% level, '**' the 5% level, and '***' the 1% level.

	Weekly Change in C&I Loans/Assets					
	(1)	(2)	(3)	(4)	(5)	(6)
Crisis	0.000858***	0.000136	-0.00517***	0.000631***	-0.00422***	-0.00505***
	(9.210)	(1.403)	(3.442)	(5.533)	(2.908)	(3.393)
%Change in hours $_{(t-2\ to\ t-1)}$		0.00217	0.00138			0.000928
		(1.274)	(0.830)			(0.557)
Crisis * %Change in hours $_{(t-2\ to\ t-1)}$		-0.00634***	-0.00548***			-0.00493***
		(3.422)	(3.074)			(2.784)
%Change in hours $_{(t-3\ to\ t-2)}$		0.00224	0.00166			0.00123
		(1.301)	(0.980)			(0.738)
Crisis * %Change in hours $_{(t-3\ to\ t-2)}$		-0.00254	-0.00187			-0.00137
		(1.406)	(1.064)			(0.800)
Crisis * Log(Assets)			0.000197***		0.000181***	0.000182***
			(3.348)		(3.025)	(3.044)
Crisis * Liquid assets/Assets			0.00240***		0.00235***	0.00236***
			(3.862)		(3.853)	(3.850)
Crisis * Core deposits/Assets			-0.000633		-0.000770	-0.000609
			(0.547)		(0.677)	(0.535)
Crisis * Tier 1 capital/Assets			0.00747*		0.00674*	0.00718*
			(1.785)		(1.729)	(1.826)
Crisis * Unused C&I comm./(Unused C&I comm.+Assets)			0.0323***		0.0325***	0.0324***
			(6.042)		(6.284)	(6.236)
Crisis * State COVID death per capita				0.0128**	0.00970**	0.00686*
				(2.015)	(2.465)	(1.788)
Bank FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	No	No	No	No	No
F-stat		7.33	5.46			4.18
p-value		0.01	0.02			0.04
Observations	8,234	8,234	8,234	8,234	8,234	8,234
R-squared	0.140	0.157	0.222	0.144	0.208	0.223

Table 3: Explaining Weekly Lending Growth, by Bank Size

This table reports panel regressions of the weekly change in bank C&I loans from January 22 to April 1, 2020 (from the Federal Reserve's FR 2644 data) on bank financial condition measures from Q4, 2019 Call Reports. We capture local credit demand with Crisis, which equals one for the weeks between March 11 and April 1, and its interactions with state-level growth in hours worked by small firms (from Homebase), along with the extent of death from COVID-19. F-stat and p-value at the bottom of the table report tests on whether coefficients on Crisis * %Change in weekly hours (t-2 to t-1) and Crisis * %Change in weekly hours (t-3 to t-2) sum up to zero. Standard errors are clustered at bank level. T-statistics are reported in parentheses. '*' denotes significance at the 10% level, '**' the 5% level, and '***' the 1% level.

	Weekly Change in C&I Loans/Assets					
	Large Banks (>$50 billion)		Medium-Sized Banks ($10-50 billion)		Small Banks (<$10 billion)	
	(1)	(2)	(3)	(4)	(5)	(6)
%Change in hours $_{(t-2\ to\ t-1)}$	-0.00130	-0.00217	0.000575	0.000146	0.00260	0.00227
	(0.205)	(0.334)	(0.130)	(0.0325)	(1.330)	(1.166)
Crisis * %Change in hours $_{(t-2\ to\ t-1)}$	-0.0124	-0.00854	-0.00121	8.72e-05	-0.00707***	-0.00612**
	(1.471)	(0.962)	(0.262)	(0.0180)	(2.745)	(2.397)
%Change in hours $_{(t-3\ to\ t-2)}$	0.00213	0.000977	0.000389	-0.000165	0.000848	0.000525
	(0.286)	(0.132)	(0.0722)	(0.0306)	(0.448)	(0.280)
Crisis * %Change in hours $_{(t-3\ to\ t-2)}$	-0.00119	0.00295	-0.00266	-0.000839	-0.00147	-0.000397
	(0.104)	(0.261)	(0.449)	(0.141)	(0.593)	(0.166)
Crisis * Log(Assets)	-5.45e-05	-9.83e-06	0.000369	0.000353	2.69e-05	1.47e-05
	(0.150)	(0.0288)	(0.921)	(0.901)	(0.352)	(0.189)
Crisis * Liquid assets/Assets	0.00205	0.00141	0.00141	0.00114	0.00115*	0.00112*
	(1.087)	(0.635)	(0.559)	(0.473)	(1.819)	(1.829)
Crisis * Core deposits/Assets	0.00400	0.00259	0.000865	0.000715	-0.00150	-0.00149
	(0.940)	(0.528)	(0.472)	(0.394)	(1.098)	(1.106)
Crisis * Tier 1 capital/Assets	0.0216	0.0114	0.0182*	0.0159	0.00302	0.00288
	(0.540)	(0.292)	(1.685)	(1.579)	(0.622)	(0.618)
Crisis * Unused C&I comm./(Unused C&I comm.+Assets)	0.0659***	0.0680***	0.0177***	0.0180***	0.0222***	0.0223***
	(14.96)	(14.32)	(2.723)	(2.820)	(3.115)	(3.203)
Crisis * State COVID death per capita		0.0181*		0.00883		0.00608
		(1.759)		(1.539)		(1.326)
Bank FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
F-stat	0.65	0.10	0.18	0.01	4.21	2.59
p-value	0.43	0.75	0.67	0.94	0.04	0.11
Observations	426	426	803	803	7,005	7,005
R-squared	0.756	0.760	0.325	0.328	0.135	0.136

Table 4: Summary Statistics for Quarterly Lending

This table reports summary statistics for the quarterly change in commercial and industrial (C&I) lending by banks from Q1, 2018 to Q1 2020. Data are from the bank Call Reports. All variables, except Assets and Log(Assets), are winsorized at 1%.

Panel A: All Banks	Pre-Crisis (2018-2019)			Crisis (Q1, 2020)		
	N	Mean	Std Dev	N	Mean	Std Dev
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Assets)	43,443	12.47	1.48	4,483	12.55	1.51
Core deposits/Assets	43,443	0.77	0.13	4,483	0.77	0.13
Tier 1 capital/Assets	43,443	0.13	0.09	4,483	0.13	0.09
C&I comm./(C&I comm.+ Assets)	43,443	0.03	0.03	4,483	0.03	0.03
Liquid assets/Assets	43,443	0.30	0.17	4,483	0.30	0.17
ΔC&I loans/Lagged Assets	43,443	0.002	0.009	4,483	0.002	0.009
Δ(C&I Loans+Unused Commit.)/(Lagged C&I comm.+ Assets)	43,443	0.002	0.011	4,483	0.002	0.010

Panel B: Large Banks (>$50 billion)	Pre-Crisis (2018-2019)			Crisis (Q1, 2020)		
	N	Mean	Std Dev	N	Mean	Std Dev
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Assets)	336	18.93	0.93	39	19.00	0.98
Core deposits/Assets	336	0.72	0.15	39	0.73	0.12
Tier 1 capital/Assets	336	0.10	0.02	39	0.09	0.02
C&I comm./(C&I comm.+ Assets)	336	0.09	0.06	39	0.09	0.06
Liquid assets/Assets	336	0.33	0.18	39	0.30	0.18
ΔC&I loans/Lagged Assets	336	0.002	0.006	39	0.017	0.015
Δ(C&I Loans+Unused Commit.)/(Lagged C&I comm.+ Assets)	336	0.003	0.009	39	0.004	0.011

Table 4: Summary Statistics for Quarterly Lending (Cont'd)

Panel C: Medium-Sized Banks ($10-50 billion)	**Pre-Crisis (2018-2019)**			**Crisis (Q1, 2020)**		
	N	Mean	Std Dev	N	Mean	Std Dev
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Assets)	743	16.79	0.45	88	16.83	0.45
Core deposits/Assets	743	0.73	0.09	88	0.74	0.09
Tier 1 capital/Assets	743	0.10	0.02	88	0.10	0.03
C&I comm./(C&I comm.+ Assets)	743	0.06	0.04	88	0.06	0.04
Liquid assets/Assets	743	0.24	0.14	88	0.24	0.14
ΔC&I loans/Lagged Assets	743	0.003	0.008	88	0.007	0.008
Δ(C&I Loans+Unused Commit.)/(Lagged C&I comm.+ Assets)	743	0.005	0.010	88	0.002	0.008

Panel D: Small Banks (<=$10 billion)	**Pre-Crisis (2018-2019)**			**Crisis (Q1, 2020)**		
	N	Mean	Std Dev	N	Mean	Std Dev
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Assets)	42,364	12.34	1.25	4,356	12.41	1.25
Core deposits/Assets	42,364	0.77	0.13	4,356	0.77	0.13
Tier 1 capital/Assets	42,364	0.13	0.09	4,356	0.13	0.09
C&I comm./(C&I comm.+ Assets)	42,364	0.03	0.03	4,356	0.03	0.03
Liquid assets/Assets	42,364	0.30	0.17	4,356	0.30	0.17
ΔC&I loans/Lagged Assets	42,364	0.002	0.009	4,356	0.002	0.009
Δ(C&I Loans+Unused Commit.)/(Lagged C&I comm.+ Assets)	42,364	0.002	0.011	4,356	0.002	0.010

Table 5: Explaining Quarterly Lending Growth

This table reports panel regressions of the quarterly change in bank C&I loans from Q1, 2018 to Q1, 2020 (nine quarters). All data are from Call Reports. Crisis is an indicator variable equal to one for Q1, 2020. All explanatory variables are from the end of the prior quarter. Standard errors are clustered at bank level. T-statistics are reported in parentheses. '*' denotes significance at the 10% level, '**' the 5% level, and '***' the 1% level.

	ΔC&I Loans/Lagged Assets			
	All banks	Large Banks (> $50 billion)	Medium-Sized Banks ($10-50 billion)	Small Banks (< $10 billion)
	(1)	(2)	(3)	(4)
Liquid assets/Assets	0.0254***	0.0216	0.0456***	0.0253***
	(13.49)	(0.894)	(4.103)	(13.49)
Crisis * Liquid assets/Assets	0.000825	0.0107	-0.00386	-0.000291
	(1.077)	(1.565)	(0.758)	(0.385)
Core deposits/Assets	0.00714***	-0.00445	-0.0114	0.00703***
	(2.769)	(0.243)	(0.986)	(2.717)
Crisis * Core deposits/Assets	0.00267	0.00282	0.00101	0.00244
	(1.632)	(0.417)	(0.121)	(1.453)
Tier 1 capital/Assets	0.0124*	0.124	-0.104*	0.0147**
	(1.690)	(1.298)	(1.755)	(2.028)
Crisis * Tier 1 capital/Assets	0.00550**	0.107	0.0398	0.00438*
	(2.236)	(0.913)	(1.404)	(1.746)
C&I comm./(C&I comm.+ Assets)	0.122***	0.0964*	0.246***	0.123***
	(11.63)	(1.682)	(3.858)	(11.57)
Crisis * C&I comm./(C&I comm.+ Assets)	0.0247***	0.195***	0.0856***	0.0151**
	(3.633)	(10.28)	(4.104)	(2.144)
Log(Assets)	-0.00547***	-0.0179***	-0.0278***	-0.00490***
	(4.096)	(3.524)	(4.541)	(3.870)
Crisis * Log(Assets)	0.000615***	0.00202*	0.00517**	0.000230*
	(5.254)	(1.707)	(2.566)	(1.815)
Observations	47,926	374	830	46,712
R-squared	0.221	0.630	0.330	0.218
Bank FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes

Table 6: Explaining Quarterly Growth in Total Credit Production

This table reports panel regressions of the quarterly change in bank C&I loans plus unused loan commitments to businesses, from Q1, 2018 to Q1, 2020 (nine quarters). All data are from *Call Reports*. *Crisis* is an indicator variable equal to one for Q1, 2020. All explanatory variables are from the end of the prior quarter. Standard errors are clustered at bank level. T-statistics are reported in parentheses. '*' denotes significance at the 10% level, '**' the 5% level, and '***' the 1% level.

	Δ(C&I Loans+Unused Commit.)/(Lagged C&I comm.+ Assets)			
	All banks	Large Banks (> $50 billion)	Medium-Sized Banks ($10-50 billion)	Small Banks (< $10 billion)
	(1)	(2)	(3)	(4)
Liquid assets/Assets	0.0137***	-0.0209	0.0363**	0.0135***
	(5.495)	(0.526)	(2.095)	(5.400)
Crisis * Liquid assets/Assets	-0.000780	-0.00259	-0.000623	-0.000932
	(0.835)	(0.217)	(0.100)	(0.979)
Core deposits/Assets	0.00553*	0.00133	-0.0110	0.00521
	(1.673)	(0.0482)	(0.665)	(1.569)
Crisis * Core deposits/Assets	0.00353*	0.00886	-0.00125	0.00340*
	(1.794)	(0.875)	(0.140)	(1.660)
Tier 1 capital/Assets	0.0195	0.0425	-0.0865	0.0217*
	(1.569)	(0.377)	(1.151)	(1.747)
Crisis * Tier 1 capital/Assets	0.00576*	-0.0145	0.0253	0.00573*
	(1.891)	(0.0773)	(0.818)	(1.821)
C&I comm./(C&I comm.+ Assets)	-0.317***	-0.237**	-0.0406	-0.321***
	(19.16)	(2.578)	(0.409)	(19.17)
Crisis * C&I comm./(C&I comm.+ Assets)	-0.00860	0.00722	-0.0111	-0.00855
	(1.070)	(0.248)	(0.430)	(1.001)
Log(Assets)	-0.00463**	-0.0232***	-0.0320***	-0.00400*
	(2.231)	(2.818)	(3.516)	(1.920)
Crisis * Log(Assets)	3.34e-05	0.00205	0.00348	3.84e-05
	(0.254)	(1.411)	(1.140)	(0.242)
Observations	47,926	374	830	46,712
R-squared	0.247	0.257	0.264	0.248
Bank FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes

Optimal control of an epidemic through social distancing[1]

Thomas Kruse[2] and Philipp Strack[3]

Date submitted: 17 May 2020; **Date accepted:** 20 May 2020

We analyze how to optimally engage in social distancing (SD) in order to minimize the spread of an infectious disease. We identify conditions under which the optimal policy is single-peaked, i.e., first engages in increasingly more social distancing and subsequently decreases its intensity. We show that the optimal policy might delay measures that decrease the transmission rate substantially to create 'herd-immunity' and that engaging in social distancing sub-optimally early can increase the number of fatalities. Finally, we find that optimal social distancing can be an effective measure in substantially reducing the death rate of a disease.

Covid Economics 21, 22 May 2020: 168-193

1 We want to thank Stefan Ankirchner and Aleh Tsyvinski for helpful discussions. We plan on updating this paper frequently, the current version of the paper can be found at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3581295.
2 Professor, Giessen University.
3 Associate Professor, Yale University and CEPR Research Fellow.

Covid Economics 21, 22 May 2020: 168-193

1 Introduction

This paper analyzes how to optimally engage in measures to contain the spread of an infectious disease. We formalize this question in the context of a standard model from epidemiology, the Susceptible-Infected-Recovered (SIR) model (Kermack and McKendrick, 1927). This model divides the population into three groups susceptible, infected and recovered, and people transition from one group into another at given exogenously specified rates depending on the size of each sub-population. We extend this model by allowing an additional parameter controlled by the planer that affects the rate at which the disease is transmitted. We think of this parameter as capturing political measures such as social distancing, and the lockdown of businesses, schools, universities and other institutions. While such measures reduce the spread of the disease, they often come at a substantial economic and social cost. We model this trade-off by considering a planner who faces convex cost in the number of infected (capturing the number of people whose death is caused by the disease) and the reduction in transmission rate (capturing the cost of shutting down society).

Our analysis identifies several features of any optimal policy. First, whenever a constant fraction of those who are infected dies, the optimal policy is single peaked in the sense that first the measures to reduce the transmission rate are escalated until some point in time, and after this point in time these measures are reduced. Second, if the cost of reducing the transmission rate is linear, meaning that closing half of society for two days is equally costly as closing all of society for one day, only the most extreme policies are used. Either, the planner imposes the maximal possible lockdown or no restrictions at all. Intuitively, the planner can achieve a greater effect by imposing a more extreme policy for a shorter time and thus does not find it optimal to use intermediate policies. These results imply that for linear cost the optimal policy has a simple structure and consists of three phase: first it imposes no restrictions then it imposes as many restrictions as possible, and finally in the third phase imposes no restrictions at all. This result drastically simplifies the search for an optimal policy as the planner has to only optimize over the start and end time of the social distancing period. We furthermore show that in this case the number of infected peaks at most twice under the optimal policy.

We then calibrate our model to the current Covid-19 epidemic to illustrate some further insights. We first characterize the optimal timing of the social distancing period given that the planner has access to a certain budget of days of social distancing. We find that the optimal social distancing is often substantially delayed. For example, if the planer has a budget of 100 days of social distancing in the next 360 days after 0.1% of the population are infected it is optimal to delay social distancing by 50 days. This initial period of letting the disease spread uncontrolled is useful as it creates "herd immunity" and thereby reduces the overall severity of the epidemic. We show

Covid Economics 21, 22 May 2020: 168-193

by an example that the benefit of herd immunity is so strong that sometimes more social distancing can *increase* the number of people that die from the epidemic. We show that in this example more people die when social distancing is imposed from day 0-100 compared to day 50-100. As this example suggests, benefit of optimally timing social distancing measures is often large and we illustrate this by comparing social distancing in the first t days after 0.1 percent are infected to t days of optimally timed social distancing. Finally, we quantify the optimal amount of social distancing. We find that for parameters commonly used to describe the spread of Covid-19 that when one assumes a value of a life of 10 million and that social distancing reduces the transmission rate by 60% that the optimal policy starts social distancing almost immediately and maintains it for around 300 days.

Related Literature Our theoretical results extend the literature on the optimal control of an infectious disease (for an overview see chapter 5 in Wickwire, 1977). There are three policy tools commonly used to control an infectious disease: 1) immunization, 2) testing and isolation of infected individuals and 3) lockdown measures that lead to a reduction in the contact rate for the whole population.

Most of the preceding literature has focused on immunization and selective isolation measures. Abakuks (1973) considers the question of how to optimally isolate infectious population if infectious population can be instantaneously isolated. Abakuks (1972, 1974) determine the optimal vaccination strategy in the same framework. Morton and Wickwire (1974) and Wickwire (1975) extend the previous work on vaccination and isolation by considering flow controls. Behncke (2000) considers more general functional forms and Hansen and Day (2011) allows for hard bounds on the control, while considering vaccination and isolation policies simultaneously. The general insight from this literature is that for linear cost of vaccination/testing the optimal policy switches from vaccinating/testing the population at the maximal feasible intensity until some point in time to vaccinating/testing no one after that point in time.

While the previously discussed literature has analyzed vaccination and isolation policies, little is known about optimal lockdown policies. In general, the analysis of optimal lockdown policies in the standard SIR model is a challenging problem. For that reason other current research has focused on mathematically easier to handle variants of the model (Gonzalez-Eiras and Niepelt, 2020).[1] To the best of our knowledge the first article that discusses the optimal social distancing or

[1] Gonzalez-Eiras and Niepelt (2020) adopt a model from Bohner et al. (2019) that assumes that there is no contact between infectious population and population that was previously infected (see the discussion in Bohner et al. (2019) on page 2) and are able to obtain closed form solutions for the optimal policy in the two polar cases in which, either every infected person dies, or no one dies from the disease.

Covid Economics 21, 22 May 2020: 168-193

lockdown policies in the SIR model is chapter 4 in Behncke (2000). In a model without terminal cost, this paper observes that the optimal policy depends only on the shadow price difference between infected and susceptible.[2]

Our paper contributes to the literature on control of an infectious diseases by deriving properties of the optimal lockdown policy about which little has been known before. We show that the gain from reducing the transmission rate, solves an ODE, which allows us to reduce the dimension of the dynamics of the problem and provide insights into the structure of the optimal policy. We provide a sufficient condition (which is also necessary for linear cost) such that it is optimal to use the minimal/maximal transmission rate at every point in time. For the case of linear cost we show that the optimal control is quasi-convex, i.e., the transmission rate first decreases and then increases. We provide further conditions such that the optimal solution is extremal, i.e., uses only the maximal and minimal transmission rate. In this case we show that the number of infected peaks at most twice under the optimal policy.

Finally, our paper also relates to the recent literature that numerically studies optimal policies for the current epidemic of Covid-19 in the context of SIR models (Alvarez et al., 2020; Kissler et al., 2020; Toda, 2020; Acemoglu et al., 2020). Alvarez et al. (2020) numerically characterize the optimal lockdown policy for the current covid pandemic in a similar SIR model. The paper Kantner (2020) considers an extended SEIR model and numerically analyzes social distancing policies that minimize disease-related deaths while establishing a desired degree of herd immunity at the same time. Kissler et al. (2020) numerically compare various lockdown policies and allow for seasonality effects. Toda (2020) estimates the transmission rate in the context of an SIR model with fixed transmission rate for various countries, compares various SD policies numerically, and considers asset prices during an epidemic. Acemoglu et al. (2020) consider a SIR model with different age groups and numerically analyze the effect of lockdown policies that target groups based on age. While it is not a goal of this paper to make any recommendations for the current Covid-19 epidemic we hope that the formal analysis and insights into the structure of the optimal policy this paper contributes will be useful in the rapidly evolving discussion of how to optimally react to the Covid-19 epidemic (Atkeson, 2020; Barro et al., 2020; Dewatripont et al., 2020; Piguillem et al., 2020; Stock, 2020; de Walque et al., 2020).

[2]This insight generalizes to our analysis (see Proposition 1).

CEPR PRESS

Covid Economics 21, 22 May 2020: 168-193

2 The Evolution of an Epidemic

The SIR Model To model the spread of an infectious disease we rely on a basic model from epidemiology, the *Susceptible Infected Recovered* (SIR) model introduced in Kermack and McKendrick (1927). We divide society into three groups: susceptible s, infected i, and the rest which is either immune to the disease as they recovered from it or died. We denote by $s(t)$ the fraction of the population that is healthy, but susceptible to disease at time t, and by $i(t)$ the fraction of the population that is infected. The SIR model assumes the number of people that gets infected, by a single infected person is deterministic proportional to the fraction of society $s(t)$ that is still susceptible to the disease. Intuitively, if only a small fraction of society is susceptible to the disease it is unlikely that an infected person meets a susceptible person. The mass of healthy people that become infected during dt thus equals

$$\beta(t)i(t)s(t)\,,$$

where the transmission rate $\beta(t)$ captures both how infectious the disease, as well as measures society has taken to influence the speed at which the disease spreads (like social distancing). Infected become non-infected, by either recovering from the disease, or dying of it at rate $\gamma > 0$, such that during a short time span dt, the fraction of infected is reduced by $\gamma i(t)$.The susceptible and infected populations $(s(t), i(t))_t$ thus for every $t \in [0, \infty)$ evolve according to the following dynamics

$$\begin{aligned} s'(t) &= -\beta(t)i(t)s(t), & s(0) &= s_0, \\ i'(t) &= \beta(t)i(t)s(t) - \gamma i(t), & i(0) &= i_0, \end{aligned} \tag{1}$$

where $s_0, i_0 \in (0,1)$ are given initial values satisfying $s_0 + i_0 \leq 1$.

Control of the Transmission Rate The time-dependent transmission rate $\beta\colon [0,\infty) \to B$ takes values in an compact interval $B = [\underline{b}, \overline{b}] \subset (0,\infty)$. We denote by $\overline{b}$ the maximal transmission rate and by $\underline{b}$ the minimal transmission rate that can be achieved through some policy measures. The set of admissible controls $\mathscr{B}$ consists of all measurable functions $\beta\colon [0,\infty) \to B$.

We introduce two cost functions $v\colon [0,1] \to [0,\infty)$ and $c\colon B \to [0,\infty)$. The cost $v(i)$ measures the number of people that die per unit of time if a share $i \in [0,1]$ of the population is infected. We suppose that $v(0) = 0$, that v is convex, continuously differentiable and strictly increasing. Convexity of v captures the fact that the probability of dying from the disease might be higher if a large share of the population is infected and the hospital system is overwhelmed. We note that v can

Covid Economics 21, 22 May 2020: 168-193

not only capture the people who die of the disease directly, but also those who die because other medical conditions remain untreated as an indirect consequence of the disease. We assume that a vaccine arrives at time T but no cure is available at that point in time. As after the comprehensive vaccination of the population no new infected would be added, the share of infected would evolve according to $i'(t) = -\gamma i(t)$ after time T and thus be given by $i(t) = i(T)e^{-\gamma(t-T)}$. The share of the population that would die after the arrival of the vaccine in this case would thus be given by

$$\bar{v}(i(T)) = \int_T^\infty v\left(i(T)e^{-\gamma(t-T)}\right) dt = \int_0^{i(T)} \frac{v(z)}{z\gamma} dz. \tag{2}$$

The cost function c captures the economic and social cost of measures taken to reduce the transmission rate. For example if social distancing measures are imposed which require the closure of most businesses this comes at a substantial economic cost. We only make minimal assumption on c and assume that it is convex and continuous, and without loss normalize the cost associated with the highest transmission rate to zero, $c(\overline{b}) = 0 > c(b)$ for all $b \in [\underline{b}, \overline{b})$.

In our model the planner trades-off the number of people who die as a direct (or indirect) consequence of the disease with the economic and social cost of reducing the transmission rate and thus aims at minimizing the cost functional

$$J(\beta) = \int_0^T v(i(t)) + c(\beta(t))\, dt + \bar{v}(i(T)) \tag{3}$$

over $\beta \in \mathscr{B}$. A policy β^* is *optimal* if it minimizes J over $\mathscr{B}$

$$\beta^* \in \arg\min_{\beta \in \mathscr{B}} J(\beta). \tag{4}$$

3 The Optimal Policy

The next result shows existence of an optimal policy and provides necessary conditions that any solution of the optimal control problem (4) must satisfy.

Proposition 1. *An optimal policy exists. Let $\beta^* \in \mathscr{B}$ be such an optimal strategy and denote by $s^*, i^* \colon [0,T] \to [0,1]$ the associated state processes satisfying (1). Then there exists a function*

Covid Economics 21, 22 May 2020: 168-193

$\eta^*: [0,T] \to \mathbb{R}$ with $\eta^*(T) = \frac{v(i^*(T))}{\gamma i^*(T)}$ such that for almost all $t \in [0,T]$ it holds

$$\begin{aligned} (\eta^*)'(t) &= \eta^*(t)\beta^*(t)i^*(t) - v'(i^*(t)) + \frac{v(i^*(t)) + c(\beta^*(t)) - \min_{b\in B}\left[\frac{1}{\gamma}v(i^*(T))s^*(T)b + c(b)\right]}{i^*(t)}, \\ (\beta^*)(t) &\in \arg\min_{b\in B}\left[\eta^*(t)i^*(t)s^*(t)b + c(b)\right]. \end{aligned} \tag{5}$$

Moreover, we have $\eta^*(t) > 0$ for all $t \in [0,T]$.

The proof of Proposition 1 relies on a sequence of auxiliary results we establish in the appendix using standard arguments from control theory that can, e.g., be found in Clarke (2013). The existence of an optimal policy follows as the convexity of c and B ensures compactness of the policy space which leads to the existence of an optimal policy. Pontryagin's optimality principle then yields that for every optimal policy there exist two Lagrange multipliers λ_1^*, λ_2^* such that the optimal control is only a function of these multipliers. As the transmission rate controls how fast susceptible population becomes infected, the optimal control can be determined from $\eta^* = \lambda_2^* - \lambda_1^* > 0$ according to (5).[3] This Lagrange multiplier $\eta^*(t)$ has a clear interpretation as marginal increase in the cost from infecting susceptible population. The fact that $\eta^*(t) > 0$ reflects the fact that the planner always benefits from having fewer infected. Finally, Proposition 1 goes beyond the Pontryagin maximum principle as it shows that the two Lagrange multipliers λ_1^*, λ_2^* can be summarized in a single Lagrange multiplier η^* whose dynamics can be expressed as an ODE (independent of λ_1, λ_2) thereby effectively reducing the dimension of the problem by 1. This, simplification of the problem allows us to explicitly characterize features of the optimal policy later.

Gain from Reducing the Transmission Rate An important quantity is the *gain from reducing the transmission rate* at time t along the optimal path which we define as

$$g^*(t) = \eta^*(t)i^*(t)s^*(t).$$

Proposition 1 implies that $g^*: [0,T] \to (0,\infty)$ is strictly greater zero and completely determines the optimal control through (5). To simplify notation we define

$$M(g) = \min_{b\in B}[g\,b + c(b)].$$

[3] A similar observation is made in Behncke (2000).

Covid Economics 21, 22 May 2020: 168-193

Our next result characterizes how the gain from reducing the transmission rate evolves over time.

Proposition 2 (Properties of the Gain from Reducing the Transmission Rate)**.**

(i) The gain from reducing the transmission rate $g^*: [0,T] \to \mathbb{R}_+$ *evolves according to*

$$(g^*)'(t) = -\gamma g^*(t) + s^*(t)\left(M(g^*(t)) - M(g^*(T)) - [v'(i^*(t))i^*(t) - v(i^*(t))]\right), \qquad (6)$$

with terminal condition $g^*(T) = \frac{1}{\gamma}[v(i^*(T))s^*(T)]$.

(ii) For all $t \in [0,T)$ *the gain* g^* *satisfies the following bounds*

$$0 < g^*(T) < g^*(t) < \frac{1}{\gamma}\left[v(1) + v'(1)\left(e^{(T-t)\gamma} - 1\right)\right]. \qquad (7)$$

(iii) Any optimal control β^* *is non-increasing in* g^*

$$g^*(t) > g^*(t') \Rightarrow \beta^*(t) \le \beta^*(t').$$

Equation (6) characterizes how the incentive to reduce the transmission rate evolves over time. The first term $-\gamma g^*(t)$ shows that the incentive is exponentially decaying at the rate at which the infected die or become cured. Intuitively, when there are fewer infected the benefits from reducing the transmission rate is smaller. The term $M(g^*(t)) - M(g^*(T))$ quantifies how much the objective function of the planner increases as a function of $g^*(t)$. This term is strictly positive, and increasing in $g^*(t)$. Finally, the third term $v'(i^*(t))i^*(t) - v(i^*(t))$ is non-negative, increasing in i^* and a measure of the increase in cost due to the convexity of v.[4] The term implies that the gain from reducing the transmission rate falls more quickly when more people are currently infected. The last part of the proposition states that the incentive to reduce the transmission rate is minimal at the final time T. This implies that the optimal control is maximal at the terminal time T.

Corollary 3. *Any optimal control* $\beta^* \in \mathscr{B}$ *satisfies* $\beta^*(t) \le \beta^*(T)$ *for all* $t \in [0,T]$.

An immediate consequence of Proposition 1 and Proposition 2 is that (i^*, s^*, g^*, β^*) solve the

[4]For example suppose that the percentage infected who dies grows linearly in the fraction of the population that is infected $v(i) = ki^2$. In this case $v'(i^*(t))i^*(t) - v(i^*(t)) = ki^2$.

following system of equations for almost all $t \in [0,T]$:

$$\begin{aligned}
(s^*)'(t) &= -\beta^*(t)i^*(t)s^*(t),\\
(i^*)'(t) &= \beta^*(t)i^*(t)s^*(t) - \gamma i^*(t),\\
(g^*)'(t) &= -\gamma g^*(t) + s^*(t)\left[M(g^*(t)) - M(g^*(T)) - [v'(i^*(t))i^*(t) - v(i^*(t))]\right],\\
s^*(0) &= s_0, \quad i^*(0) = i_0, \quad g^*(T) = \frac{v(i^*(T))s^*(T)}{\gamma},\\
\beta^*(t) &\in \arg\min_{b\in B}\left[g^*(t)b + c(b)\right].
\end{aligned} \tag{8}$$

We throughout restrict attention to optimal policies where $\beta^*(t) \in \arg\min_{b\in B}[g^*(t)b + c(b)]$ is satisfied for *all* $t \in [0,T]$. This assumption is inconsequential in the sense that changing the policy β^* on a set of measure zero does not affect the paths of susceptible and infected (s^*, i^*).

We note that (8) is a system of coupled equations where s^*, i^* run forward in time from the initial condition $(s^*(0), i^*(0)) = (i_0, s_0)$ and g^* runs backward in time ending with the terminal condition $g^*(T) = \frac{1}{\gamma}v(i^*(T))s^*(T)$. As $(g^*)'(t)$ depends on $g^*(T)$ the gain g^* is *not* characterized by an ordinary differential equation. We denote by $c'(\underline{b}+) = \lim_{b\searrow\underline{b}} \frac{c(b)-c(\underline{b})}{b-\underline{b}}$ the right-derivative of c at $\underline{b}$ and by $c'(\overline{b}-) = \lim_{b\nearrow\overline{b}} \frac{c(\overline{b})-c(b)}{\overline{b}-b}$ the left-derivative of c at $\overline{b}$ (which exist due to convexity of c). Due to the forward backward nature of the coupled system existence and uniqueness is in general not guaranteed. In our case existence follows from the existence of an optimal policy which we established in Proposition 1. While we conjecture that (8) admits a unique solution if c and v are sufficiently regular, we were so far not able to identify sufficient regularity conditions. We thus state our next result under the assumption that (8) admits a unique solution. The result establishes sufficient conditions that ensure that using the minimal resp. the maximal transmission rate over the whole time interval $[0,T]$ is the optimal policy.

Lemma 4 (Minimal or Maximal Transmission Rate is Optimal)**.** *Suppose that the solution* (i^*, s^*, g^*, β^*) *to* (8) *is unique.*

(i) Let $\beta \in \mathscr{B}$ *satisfy* $\beta(t) = \underline{b}$ *for all* $t \in [0,T]$ *and let* $s, i\colon [0,T] \to [0,1]$ *be the associated paths satisfying* (1)*. If*

$$\frac{v(i(T))s(T)}{\gamma} \geq -c'(\underline{b}+), \tag{9}$$

then $\beta \equiv \underline{b}$ *is an optimal control.*

(ii) Let $\beta \in \mathscr{B}$ *satisfy* $\beta(t) = \overline{b}$ *for all* $t \in [0,T]$ *and let* $s, i\colon [0,T] \to [0,1]$ *be the associated paths*

Covid Economics 21, 22 May 2020: 168-193

satisfying (1). *If*

$$\frac{1}{\gamma}\left[e^{T\gamma}v(i(T))s(T)+(v'(1)-v(1))\left(e^{T\gamma}-1\right)\right]\leq -c'(\overline{b}-), \tag{10}$$

then $\beta\equiv\overline{b}$ *is an optimal control. The condition*

$$\frac{1}{\gamma}\left[v(1)+v'(1)\left(e^{T\gamma}-1\right)\right]\leq -c'(\overline{b}-) \tag{11}$$

is sufficient for (10).

The condition (9) identified by Lemma 4 is surprisingly simple: Observe that $\frac{v(i(T))}{\gamma i(T)}$ is the fraction of the infected that will die from the disease at time T.[5] This number is multiplied by the fraction of population that is infected and by the fraction of population that is susceptible at time T and compared to the marginal cost of further reducing the transmission rate. For example, suppose that a disease kills 5% of the infected $v(i)=0.05i\gamma$. Furthermore, suppose that c is linear and shutting down 1% of the economy reduces the transmission rate by 1%. If the planner values one life at 10 million dollar, which corresponds to 148 US per capita GDPs this leads to a cost of $-c'\equiv\frac{1}{148}$. Together, this yields that the minimal transmission rate is optimal if the number of susceptible and infected at time T under the minimal transmission rate satisfies $i(T)s(T)\geq\frac{1}{0.05\times 148}=0.135$.

3.1 Linear Costs

In this section we impose additional linearity assumptions on the cost to provide further insight into the structure of the optimal policy. Again we suppose that $\beta^*\in\mathscr{B}$ is an optimal control and denote by $s^*,i^*\colon[0,T]\to[0,1]$ the associated state processes.

Our first type of result assumes that v is linear, which means that the fraction of infected that die from the disease is independent of the total fraction of the population that is infected at any point in time. This assumption rules out capacity effects that arise from the overload of the medical system. It is thus a reasonable assumption if the number of infected is kept within levels that do not overburden the health system.

Proposition 5. *Suppose that v is linear and let $\beta^*\in\mathscr{B}$ be an optimal control.*[6]

(i) There exists $t^*\in[0,T]$, such that β^* is non-increasing on $[0,t^*]$ and non-decreasing on $[t^*,T]$, i.e., β^* is quasi-convex.

[5]We note that as $v(i)$ is the rate at which the deceased population increases and γi is the rate at which the infected population decreases it follows that $\frac{v(i)}{\gamma i}$ is the probability that an infected person dies.

[6]There exists $\alpha>0$ such that $v(i)=\alpha i$

Covid Economics 21, 22 May 2020: 168-193

(ii) The fraction of infected population $i^*(\cdot)$ is strictly log-concave on $[0,t^*]$ and thus admits at most one local maximum on $[0,t^*]$.

Proposition 5 establishes that any optimal policy is single peaked, in the sense that the measures to decrease the transmission rate are first escalated until some point in time and then reduced over time. Any policy where a reduction in measures is followed by an increase is suboptimal. Furthermore, in the initial period where SD measures are escalated and the transmission rate falls the fraction of infected is single peaked.

Our next result establishes that if both costs v and c are linear then the optimal policy involves only the two most extreme controls. The assumption that the cost c of measures that reduce the transmission rate is linear has a simple interpretation in the context of social distancing: Shutting down half of the economy for two days is equally costly as shutting down the whole economy for a single day.[7] While we think that there is no normative reason for this assumption we think of it as a natural baseline for the analysis.

Proposition 6. *Suppose that v and c are linear.*[8] *Then for any optimal control β^* there exists $0 \leq t_1^* \leq t_2^* \leq T$ such that for a.e. $t \in [0,T]$*

$$\beta^*(t) = \begin{cases} \overline{b} & \text{for } t \in [0,t_1^*) \\ \underline{b} & \text{for } t \in [t_1^*,t_2^*] \\ \overline{b} & \text{for } t \in (t_2^*,T] \end{cases} .$$

Furthermore, the number of infected i^* under an optimal policy has at most two local maxima on $[0,T]$.

Proposition 6 drastically simplifies the search for an optimal policy as it implies that any optimal policy is characterized by the two points in time (t_1^*,t_2^*). Note that the proposition does not rule out that any of the intervals is empty. In particular, reducing the transmission rate by the maximal amount at every point in time as well as taking no measures at all to reduce the transmission rate can be optimal. The main insight of the proposition is that under plausible assumption it is never optimal to use intermediate measure for a longer time (i.e. closing only parts of the economy) as doing so is dominated by implementing maximal measures for a shorter time. The second part of the proposition establishes that under the optimal policy the number of infected peaks at most twice.

[7]This implicitly assumes that the transmission rate β depends linearly on the shut down of the economy.

[8]There exists $\alpha,\delta > 0$ such that $v(i) = \alpha i$ and $c(\beta) = \delta(\overline{b} - \beta)$.

Covid Economics 21, 22 May 2020: 168-193

CEPR PRESS

Covid Economics 21, 22 May 2020: 168-193

4 An Illustration

We next illustrate how our results can be used to derive policy advice for fighting an epidemic. In this illustration we aim to choose parameters in line with the COVID-19 pandemic. An important disclaimer is that at the current point in time there is substantial uncertainty about the true parameters governing the spread of COVID-19 which substantially influence the optimal policies.

Parametric Assumptions We assume that the average length of an infection equals 18 days ($\gamma = 1/18$) and that the social planner has access to two policies $0 < \underline{b} < \overline{b}$ corresponding to *social distancing* (SD) $\underline{b}$ and *no social distancing* (NSD) $\overline{b}$. We set $\overline{b}$ to 0.16 in line with a reproduction rate of R_0 of $\overline{b}/\gamma = 2.88$. We assume that enacting social distancing reduces the number of contacts by 60% and set $\underline{b} = 0.4\overline{b}$ consistent with $R_0 = \underline{b}/\gamma = 1.152$. The fraction of infected that dies equals 0.8%, below 20 times[9] the critical care bed capacity $\kappa = 0.000347$[10] and then grows linearly such that if 20% of the population is simultaneously infected 5% of infected die[11]

$$v(i) = (\gamma i) \times \left[0.008 + \frac{0.042(\gamma i - 20\kappa)}{\gamma 0.2 - 20\kappa} \mathbf{1}_{\gamma i \geq 20\kappa} \right].$$

Throughout our simulations we assume that at day zero, 0.1% of the population is infected. Finally, we assume that a vaccine for the disease arrive in one year (360 days) and is immediately deployed widely such that no one gets infected by the disease afterwards.

The Optimal Timing of Social Distancing We begin by analysing the optimal timing of social distancing. In order to do so we first suppose that the planner has a fixed budget of days of social distancing and answer the question during which time period he optimally engages in social distancing. We only consider policies that consist of three subsequent periods, first NSD, follows by a period of SD, and a period of NSD. For example, consider the case where the planer has a budget of 100 days of SD. In this case the optimal policy is to start social distancing on day 48 and end it on day 148. As one can see in the left graph of Figure 1 this leads to a substantially

[9]This is we implicitly assume that 5% of infections are sufficiently severe that they need hospitalization and access to critical care.

[10]The number of cricitcal care beds per population equals $\kappa = 0.000347$ for the US, $\kappa = 0.000292$ for Germany, and $\kappa = 0.000125$ for Italy. See https://www.sccm.org/getattachment/Blog/March-2020/United-States-Resource-Availability-for-COVID-19/United-States-Resource-Availability-for-COVID-19.pdf?lang=en-US.

[11]Note, that $v(i)$ aims not only at capturing the people who die directly as a consequence of the disease, but also those who die as they do not have access to critical care as a consequence of the overloaded medical system. Our assumption implies that if 50% of the population is simultaneously infected the death rate increases to 11.5%. We

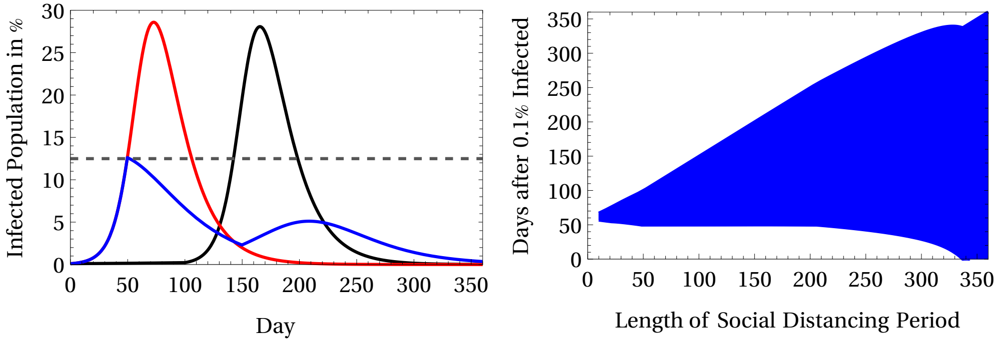

Figure 1: On the left: The number of infected over time without SD (in red), with SD from day 0-100 (in black), and for the optimal SD period from day 48 to day 148 in blue. On the right: the optimal period of SD as a function of the length of SD. The dashed line marks the point beyond which the fraction of infected who die increases due to an overload of the medical system.

flatter curve of infected over time than social distancing in the first 100 days (in black) or no social distancing (in red). Interestingly, the effect of suboptimal social distancing is marginal in the sense that while it initially reduces the number of infected substantially, it essentially only delays the peak of infected, but does not substantially flatten it. This leads to a substantial reduction in the implied death rate within a year: 0.6% under optimal social distancing, 4.6% with social distancing in the first hundred days, and 4.8% without social distancing.

We next analyse how the optimal timing of social distancing depends on the length of social distancing. As one can see in the right graph of Figure 1 it is optimal to delay social distancing beyond the date where 0.1% of the population is infected. For example even if it is optimal for the planner to engage in 300 days of social distancing within the next year it is only optimal to start social distancing after 25 days. This observation might be surprising as it implies that if it is not optimal to maintain permanent social distancing (until the arrival of a vaccine/cure), then it is optimal to delay the period of SD.

The Value of Social Distancing Whether or not the planner wants to engage in SD is an orthogonal question to the optimal timing of SD. To study this question we plot in Figure 2 the death rate withing a year as a function of the number of days of social distancing. As one can see in the figure social distancing can be an effective measure to prevent the death of population. For example, 50 days of optimally timed social distancing (from day 50 to day 100) reduce the death rate by

note that these are extremely pessimistic assumptions if the number of infected is substantially underestimated.

Covid Economics 21, 22 May 2020: 168-193

Covid Economics 21, 22 May 2020: 168-193

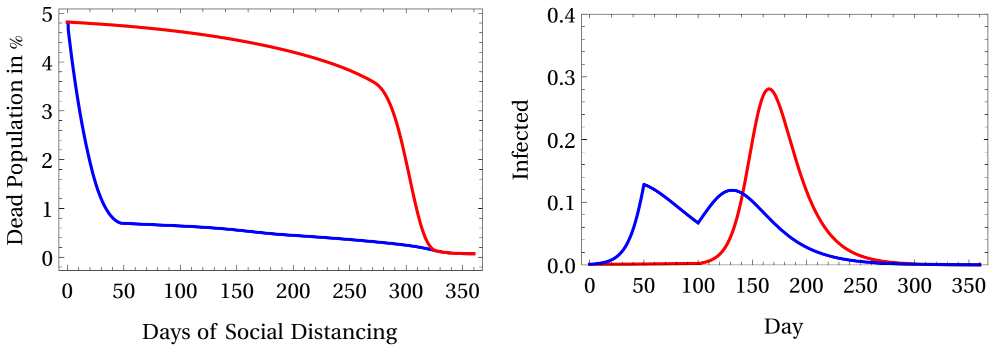

Figure 2: On the Left: Fraction of the population that dies within 360 days of 0.1% infected as a function of the length of social distancing for the optimal timing of social distancing (in blue) and social distancing starting at day 0 (in red). On the right: infected over time, if SD is exercises from day 50-100 (in blue) and from day 0-100 (in red).

roughly 4%. The figure however shows that without optimal timing SD is much less effective and to achieve an equal reduction in the example one needs more than 300 days of SD.

We asses the value of optimal social distancing in prevented dead where we plot the percentage of dead population prevented per day of SD as a function of the number of days the planner engages in SD. The initial efficacy of social distancing is around 0.2% per day of SD and then decreases to 0.05% around 100 days. If one assigns a value of around 10 million dollar to a statistical life as it is typically estimated in the literature[12] then one life corresponds to around 148 US per capita GDPs[13], which implies that the planner should be willing to endure a day of SD to save 0.0019% of the population. Thus, for the commonly assumed value of a statistical life the planner should engage in constant SD until a vaccine or cure is found. Figure 3 shows that this conclusion is robust and stays valid even if one assigns just a tenth of the commonly assumed value to a life, i.e. 1 million $.

Social Distancing can Lead to more Dead We next illustrate how suboptimally timed social distancing can actually *increase* the number of fatalities as a consequence of the epidemic. A particular example of this is shown in the right graph of Figure 2 which shows that social distancing from day 50-100 can lead to a substantially *flatter* curve than SD from day 0-100. The reason for this perhaps surprising phenomenon is that by not engaging in SD early, many more people are

[12] See for example Viscusi and Aldy (2003).

[13] See https://en.wikipedia.org/wiki/List_of_countries_by_GDP_(PPP)_per_capita.

Covid Economics 21, 22 May 2020: 168-193

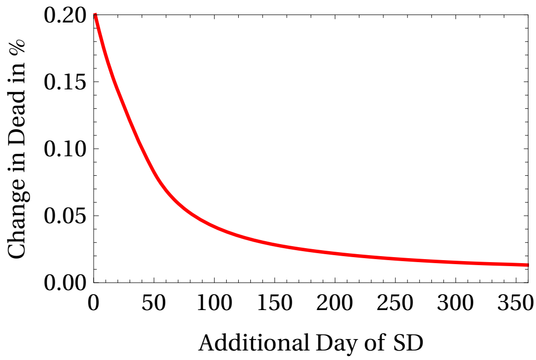

Figure 3: Prevented dead population per day of optimal social distancing.

infected before day 100 which are then later on immune. The "herd-immunity" that is created this way leads to a substantially lower peak in infections and fewer dead (0.7% of the population vs 4.6%).

5 Extensions

5.1 Random Arrival of a Vaccine

In this section we introduce a variant of the model of Section 2 where the time until a vaccine or cure is available is random. More formally, we let $\tau\colon \Omega \to [0,T]$ be a bounded random variable on a probability space $(\Omega, \mathscr{F}, \mathbb{P})$. We assume that τ has a continuous density function $p\colon [0,T] \to [0,\infty)$ and we denote by $F(t) = \mathbb{P}[\tau \le t] = \int_0^t p(s)\,ds$ its distribution function[14]. The expected costs of a strategy $\beta \in \mathscr{B}$ are given by

$$J(\beta) = \mathbb{E}\left[\int_0^\tau v(i(t)) + c(\beta(t))\,dt\right]. \qquad (12)$$

These can be transformed to

$$J(\beta) = \mathbb{E}\left[\int_0^T \mathbf{1}_{[0,\tau)}(s)(v(i(t)) + c(\beta(t)))\,dt\right] = \int_0^T (1 - F(t))(v(i(t)) + c(\beta(t)))\,dt. \qquad (13)$$

We obtain the following variant of Proposition 1[15].

[14] Without loss of generality we assume that T is the smallest upper bound of τ, i.e., $F(t) < 1$ for all $t < T$.
[15] See (Clarke, 2013, Corollary 22.6) for a statement of the maximum principle for time-dependent payoffs.

CEPR PRESS

Proposition 7. *An optimal policy exists. Let $\beta^* \in \mathscr{B}$ be such an optimal strategy and denote by $s^*, i^* \colon [0,T] \to [0,1]$ the associated state processes satisfying (1). Then there exists a function $\eta^* : [0,T] \to \mathbb{R}$ with $\eta^*(T) = 0$ such that for almost all $t \in [0,T]$ it holds*

$$(\eta^*)'(t) = \frac{(1-F(t))(v(i^*(t)) + c(\beta^*(t)) - v'(i^*(t))i^*(t)) - \int_t^T (v(i^*(s)) + c(\beta^*(s)))p(s)\,ds}{i^*(t)} + \eta^*(t)\beta^*(t)i^*(t) \quad (14)$$

and

$$\beta^*(t) \in \arg\min_{b \in B} \left[\eta^*(t)i^*(t)s^*(t)b + (1-F(t))c(b)\right]. \quad (15)$$

Moreover, we have $\eta^*(t) > 0$ for all $t \in [0,T)$.

6 Conclusion

We derived the optimal policy for social distancing during an epidemic. Our analysis revealed several features of the optimal policy. For cost linear in the number of infected, the optimal policy consists of two phases, a first phase where the measures taken to decrease the transmission rate are escalated and then a second phase where these measures are reduced. Furthermore, if the cost of reducing the transmission rate is linear, the optimal policy is always extreme. At any point in time either social distancing is carried out to the maximal extend possible or not at all. The intuitive reason for this result is that more extreme measures over a shorter time horizon are more effective than less extreme measures over a longer horizon. We illustrated through an example that the effectiveness of social distancing depends crucially on its optimal timing. Within the context of this example optimal social distancing is often substantially delayed in order to generate herd immunity. Engaging in more, but too early social distancing can increase the peak number of infected and thereby the fatalities from the disease.

Covid Economics 21, 22 May 2020: 168-193

Covid Economics 21, 22 May 2020: 168-193

A Appendix

We first prove a sequence of auxiliary results which together imply Proposition 1. Recall that the function $M\colon \mathbb{R} \to \mathbb{R}$ was defined by $M(y) = \min_{b\in B}[yb + c(b)]$, $y \in \mathbb{R}$ and note that M is non-decreasing.

Lemma 8. *An optimal policy β^* exists that solves (4). Let $s^*, i^*\colon [0,T] \to [0,1]$ be the state processes associated with an optimal control satisfying (1). Then there exist absolutely continuous functions $\lambda_1^*, \lambda_2^*\colon [0,T] \to \mathbb{R}$ which satisfy for almost all $t \in [0,T]$ the dynamics*

$$\begin{aligned} (\lambda_1^*)'(t) &= (\lambda_1^*(t) - \lambda_2^*(t))\beta^*(t)i^*(t), & \lambda_1^*(T) &= 0, \\ (\lambda_2^*)'(t) &= (\lambda_1^*(t) - \lambda_2^*(t))\beta^*(t)s^*(t) + \gamma\lambda_2^*(t) - v'(i^*(t)), & \lambda_2^*(T) &= \frac{v(i^*(T))}{\gamma i^*(T)}, \end{aligned} \tag{16}$$

and the optimality condition

$$\beta^*(t) \in \operatorname*{arg\,min}_{b\in B} \left[(\lambda_2^*(t) - \lambda_1^*(t))i^*(t)s^*(t)b + c(b)\right]. \tag{17}$$

Moreover, for all $t \in [0,T]$ we have

$$M\big(\lambda_2^*(t) - \lambda_1^*(t))i^*(t)s^*(t)\big) - \gamma\lambda_2^*(t)i^*(t) + v(i^*(t)) = M\left(\frac{v(i^*(T))s^*(T)}{\gamma}\right). \tag{18}$$

Proof of Lemma 8. Suppose an optimal policy β^* exists. The existence of λ_1^*, λ_2^* that satisfy (16), (17) and (18) follows from the Pontryagin principle (see, e.g., Clarke, 2013, Theorem 22.2 and Corollary 22.3). We show the existence of an optimal policy by verifying the conditions of Theorem 23.11 in Clarke (2013).

(a) $g(t,(s,i)) = \begin{pmatrix} -is \\ +is \end{pmatrix}$ which implies that $|g(t,(s,i))| \le 2|is| < 2$.

(b) $B = [\underline{b}, \overline{b}]$ is closed and convex by definition.

(c) The sets $E = \{(s_0, i_0)\} \times \mathbb{R}_+$ and $Q = [0,T] \times [0,1]^2$ are closed and $\ell(s_0, i_0, s_T, i_T) = \bar{v}(i_T)$ is lower semicontinuous.

(d) The running cost $\beta \mapsto v(i) + c(\beta)$ is convex as c is convex. Furthermore, $v(i) + c(\beta) \ge 0$.

(e) The projection set is given by $\{(s_0, i_0)\}$ and thus bounded.

(f) As $\beta \in B$ it follows that $|\beta| \le \overline{b}$. This verifies (f) (ii).

Moreover, the constant control $\beta(t) = \overline{b}$ has finite costs. We have hence verified that there exists an optimal policy. □

Covid Economics 21, 22 May 2020: 168-193

In the following we suppose that $\beta^* \in \mathscr{B}$ is an optimal control and denote by $s^*, i^* \colon [0,T] \to [0,1]$ the associated state processes satisfying (1). Moreover, we denote by $\lambda_1^*, \lambda_2^* \colon [0,T] \to \mathbb{R}$ the Lagrange variables from Lemma 8. Note that compactness of B and continuity of c ensure that for all $t \in [0,T]$ the function $b \mapsto [\lambda_2^*(t) - \lambda_1^*(t)]i^*(t)s^*(t)b + c(b)$ attains its minimum on B. By (17) this minimum is attained by $\beta^*(t)$ for almost all $t \in [0,T]$. By potentially changing β^* on a set of measure zero we suppose in the sequel that $\beta^*(t)$ attains the minimum for all $t \in [0,T]$ (i.e., (17) holds for all $t \in [0,T]$). Note that this change does not affect the trajectories of s^*, i^*, λ_1^* and λ_2^*.

We introduce the new Lagrange variable

$$\eta^*(t) = \lambda_2^*(t) - \lambda_1^*(t). \tag{19}$$

The variable $\eta^*(t)$ has a clear interpretation: it measures the marginal change in the cost with respect to infecting susceptible population. Intuitively speaking, $\eta^*(t)$ measures the additional cost if one additional person is infected at time t given the optimal policy is used. Note that by (17) at each time t the optimal control $\beta^*(t)$ depends on $\lambda_1^*(t)$ and $\lambda_2^*(t)$ only through their difference $\eta^*(t) = \lambda_2^*(t) - \lambda_1^*(t)$.

Lemma 9. *Let $\beta^* \in \mathscr{B}$ be an optimal control and suppose that the optimality condition (17) holds for all $t \in [0,T]$. Suppose that $t_0 \in [0,T]$ satisfies $\eta^*(t_0) \le 0$. Then it holds that $\lim_{t \to t_0} \beta^*(t) = \beta^*(t_0) = \overline{b}$.*

Proof of Lemma 9. First note that the assumption $\eta^*(t_0) \le 0$ ensures that the function $b \mapsto \eta^*(t)i^*(t)s^*(t)b + c(b)$ attains its global minimum on B at $\overline{b}$. Hence (17) implies that $\beta^*(t_0) = \overline{b}$. Next let (t_n) be a sequence such that $t_n \to t_0$ as $n \to \infty$. Suppose by contradiction that there exists a subsequence such that $\lim_{n\to\infty} \beta^*(t_n) =: b_0 < \overline{b}$. Next note that (17) ensures for all $n \in \mathbb{N}$ that (recall that $c(\overline{b}) = 0$)

$$\eta^*(t_n)i^*(t_n)s^*(t_n)\beta^*(t_n) + c(\beta^*(t_n)) \le \eta^*(t_n)i^*(t_n)s^*(t_n)\overline{b}. \tag{20}$$

This implies that

$$\eta^*(t_n)i^*(t_n)s^*(t_n) \ge \frac{c(\beta^*(t_n))}{\overline{b} - \beta^*(t_n)}. \tag{21}$$

Taking the limit $n \to \infty$ yields the contradiction

$$0 = \lim_{n\to\infty} \eta^*(t_n)i^*(t_n)s^*(t_n) \ge \lim_{n\to\infty} \frac{c(\beta^*(t_n))}{\overline{b} - \beta^*(t_n)} = \frac{c(b_0)}{\overline{b} - b_0} > 0. \tag{22}$$

Therefore, we have $\lim_{t\to t_0} \beta^*(t) = \overline{b} = \beta^*(t_0)$. □

Covid Economics 21, 22 May 2020: 168-193

The next result shows that the cost of additional infected η^* is characterized by an ordinary differential equation (ODE) that does not depend on λ_1^* and λ_2^*. Moreover, we show that both λ_1^* and λ_2^* can be recovered from η^*.

Lemma 10. *The variable η^* solves*

$$(\eta^*)'(t) = \eta^*(t)\beta^*(t)i^*(t) - v'(i^*(t)) + \frac{v(i^*(t)) + c(\beta^*(t)) - M\left(\frac{v(i^*(T))s^*(T)}{\gamma}\right)}{i^*(t)} \tag{23}$$

with terminal condition $\eta^*(T) = \frac{v(i^*(T))}{\gamma i^*(T)}$. Conversely, suppose that $i, s, \beta, \eta \colon [0,T] \to \mathbb{R}$ satisfy

$$\begin{aligned} s'(t) &= -\beta(t)i(t)s(t), \quad s(0) = s_0, \\ i'(t) &= \beta(t)i(t)s(t) - \gamma i(t), \quad i(0) = i_0, \\ \eta'(t) &= \eta(t)\beta(t)i(t) - v'(i(t)) + \frac{v(i(t)) + c(\beta(t)) - M\left(\frac{v(i(T))s(T)}{\gamma}\right)}{i(t)}, \quad \eta(T) = \frac{v(i(T))}{\gamma i(T)}, \\ \beta(t) &\in \arg\min_{b\in B}\left[\eta(t)i(t)s(t)b + c(b)\right], \end{aligned} \tag{24}$$

then

$$\begin{aligned} \lambda_1(t) &= \frac{1}{\gamma}\left[\eta(t)\beta(t)s(t) + \frac{v(i(t)) + c(\beta(t)) - M\left(\frac{v(i(T))s(T)}{\gamma}\right)}{i(t)}\right] - \eta(t), \\ \lambda_2(t) &= \frac{1}{\gamma}\left[\eta(t)\beta(t)s(t) + \frac{v(i(t)) + c(\beta(t)) - M\left(\frac{v(i(T))s(T)}{\gamma}\right)}{i(t)}\right] \end{aligned} \tag{25}$$

solves (16).

Proof of Lemma 10. First note that it follows from (18) and (17) that

$$\begin{aligned} -\eta^*(t)\beta^*(t)i^*(t)s^*(t) + \gamma\lambda_2^*(t)i^*(t) &= -M(\eta^*(t)i^*(t)s^*(t)) + \gamma\lambda_2^*(t)i^*(t) + c(\beta^*(t)) \\ &= v(i^*(t)) + c(\beta^*(t)) - M\left(\frac{v(i^*(T))s^*(T)}{\gamma}\right). \end{aligned} \tag{26}$$

Covid Economics 21, 22 May 2020: 168-193

Then (16) implies that

$$
\begin{aligned}
(\eta^*)'(t) &= (\lambda_2^*)'(t) - (\lambda_1^*)'(t) = -\eta^*(t)\beta^*(t)s^*(t) + \gamma\lambda_2^*(t) - v'(i^*(t)) + \eta^*(t)\beta^*(t)i^*(t) \\
&= \eta^*(t)\beta^*(t)i^*(t) - v'(i^*(t)) + \frac{v(i^*(t)) + c(\beta^*(t)) - M\left(\frac{v(i^*(T))s^*(T)}{\gamma}\right)}{i^*(t)}.
\end{aligned} \tag{27}
$$

Next suppose that s, i and η solve (24) and that λ_1 and λ_2 are given by (25). Observe that it holds that $\lambda_1(T) = 0$ and $\lambda_2(T) = \frac{v(i(T))}{\gamma i(T)}$. Next note that the envelope theorem ensures that

$$
\frac{\partial}{\partial t}\left[\eta(t)i(t)s(t)\beta(t) + c(\beta(t))\right] = \frac{\partial}{\partial t}\min_{b\in B}\left[\eta(t)i(t)s(t)b + c(b)\right] = \beta(t)\frac{\partial}{\partial t}\left[\eta(t)i(t)s(t)\right]. \tag{28}
$$

Then it holds that

$$
\begin{aligned}
\lambda_2'(t) &= \frac{\partial}{\partial t}\left[\frac{\eta(t)\beta(t)i(t)s(t) + v(i(t)) + c(\beta(t)) - M\left(\frac{v(i(T))s(T)}{\gamma}\right)}{\gamma i(t)}\right] \\
&= \frac{\beta(t)\frac{\partial}{\partial t}\left[\eta(t)i(t)s(t)\right] + v'(i(t))i'(t)}{\gamma i(t)} \\
&\quad - \frac{\left(\eta(t)\beta(t)i(t)s(t) + v(i(t)) + c(\beta(t)) - M\left(\frac{v(i(T))s(T)}{\gamma}\right)\right)i'(t)}{\gamma(i(t))^2} \\
&= \frac{1}{\gamma}\left(\beta(t)\eta'(t)s(t) + \beta(t)\eta(t)s'(t)\right) \\
&\quad + \frac{i'(t)}{\gamma i(t)}\left(v'(i(t)) - \frac{v(i(t)) + c(\beta(t)) - M\left(\frac{v(i(T))s(T)}{\gamma}\right)}{i(t)}\right) \\
&= \frac{1}{\gamma}\left(\beta(t)\eta'(t)s(t) + \beta(t)\eta(t)s'(t)\right) + \frac{i'(t)}{\gamma i(t)}\left(\eta(t)\beta(t)i(t) - \eta'(t)\right) \\
&= \frac{\beta(t)\eta(t)}{\gamma}\left(s'(t) + i'(t)\right) + \frac{\eta'(t)}{\gamma}\left(\beta(t)s(t) - \frac{i'(t)}{i(t)}\right) \\
&= -\beta(t)\eta(t)i(t) + \eta'(t).
\end{aligned}
$$

Therefore we obtain that

$$
\begin{aligned}
\lambda_2'(t) &= \frac{v(i(t)) + c(\beta(t)) - M\left(\frac{v(i(T))s(T)}{\gamma}\right)}{i(t)} - v'(i(t)) = \gamma\lambda_2(t) - \eta(t)\beta(t)s(t) - v'(i(t)) \\
&= (\lambda_1(t) - \lambda_2(t))\beta(t)s(t) + \gamma\lambda_2(t) - v'(i(t)).
\end{aligned}
$$

Similarly, λ_1 satisfies

$$\lambda_1'(t) = \lambda_2'(t) - \eta'(t) = [\eta'(t) - \eta(t)\beta(t)i(t)] - \eta'(t) = (\lambda_1(t) - \lambda_2(t))\beta(t)i(t). \qquad \square$$

Lemma 11 (More Infected are Costly)**.** *The function η^* satisfies $\eta^*(t) > 0$ for all $t \in [0,T)$.*

Proof of Lemma 11. Suppose that there exists $t \in [0,T)$ such that $\eta^*(t) \leq 0$. Then Lemma 9 shows that $\beta^*(t) = \bar{b}$. Moreover, Lemma 9 ensures that β^* is continuous at t and hence η^* is differentiable at t. Then (23) shows (recall that $c(\bar{b}) = 0$)

$$(\eta^*)'(t) = \eta^*(t)\beta^*(t)i^*(t) + \frac{v(i^*(t)) - M\left(\frac{v(i^*(T))s^*(T)}{\gamma}\right) - v'(i^*(t))i^*(t)}{i^*(t)}. \tag{29}$$

Since v is convex and since $M\left(\frac{v(i^*(T))s^*(T)}{\gamma}\right) > 0$ we thus obtain that

$$(\eta^*)'(t) < \eta^*(t)\beta^*(t)i^*(t) \leq 0 \tag{30}$$

We conclude from the terminal condition $\eta^*(T) = \frac{v(i^*(T))}{\gamma i^*(T)} > 0$ that $\eta^*(t) > 0$ for all $t \in [0,T)$. $\square$

Since $(\lambda_1^*)'(t) = -\eta^*(t)\beta^*(t)i^*(t)$ we obtain from Lemma 11 that λ_1^* is decreasing in time and, in particular, that λ_1^* is non-negative. This means that a marginal increase of the susceptible population (while keeping the infected population constant) marginally increases the costs. This marginal effect decreases over time and vanishes at time T.

Combining the results of Lemma 8, Lemma 9, Lemma 10 and Lemma 11 proves Proposition 1.

Proof of Proposition 2. Recall that $g^*(t) = \eta^*(t)i^*(t)s^*(t)$, $t \in [0,T]$. We compute $(g^*)'(t)$

$$\begin{aligned}
(g^*)'(t) &= (\eta^*)'(t)i^*(t)s^*(t) + (\eta^*(t))(i^*)'(t)s^*(t) + \eta^*(t)i^*(t)(s^*)'(t) \\
&= [(\lambda_1^*(t) - \lambda_2^*(t))\beta^*(t)s^*(t) + \lambda_2^*(t)\gamma - v'(i^*(t)) - (\lambda_1^*(t) - \lambda_2^*(t))\beta^*(t)i^*(t)]i^*(t)s^*(t) \\
&\quad + (\lambda_2^*(t) - \lambda_1^*(t))s^*(t)[\beta^*(t)i^*(t)s^*(t) - \gamma i^*(t)] \\
&\quad - (\lambda_2^*(t) - \lambda_1^*(t))i^*(t)\beta^*(t)i^*(t)s^*(t) \\
&= (\lambda_1^*(t) - \lambda_2^*(t))\beta^*(t)[-(s^*(t))^2 i^*(t) + (i^*(t))^2 s^*(t) + (s^*(t))^2 i^*(t) - (i^*(t))^2 s^*(t)] \\
&\quad + [\gamma\lambda_2^*(t) - v'(i^*(t))]i^*(t)s^*(t) - \gamma(\lambda_2^*(t) - \lambda_1^*(t))s^*(t)i^*(t) \\
&= [\gamma\lambda_1^*(t) - v'(i^*(t))]i^*(t)s^*(t).
\end{aligned} \tag{31}$$

Covid Economics 21, 22 May 2020: 168-193

Covid Economics 21, 22 May 2020: 168-193

Using (31) and the fact that $g^*(t) = (\lambda_2^*(t) - \lambda_1^*(t))i^*(t)s^*(t)$ we have that

$$(g^*)'(t) = \gamma\lambda_1^*(t)i^*(t)s^*(t) - v'(i^*(t))i^*(t)s^*(t) = -\gamma g^*(t) + \gamma\lambda_2^*(t)i^*(t)s^*(t) - v'(i^*(t))i^*(t)s^*(t).$$

Next (18) implies that

$$(g^*)'(t) = -\gamma g^*(t) + s^*(t)\left[M(g^*(t)) - M\left(\frac{v(i^*(T))s^*(T)}{\gamma}\right) + v(i^*(t)) - v'(i^*(t))i^*(t)\right].$$

Using that $g^*(T) = \eta^*(T)i^*(T)s^*(T) = \frac{v(i^*(T))s^*(T)}{\gamma}$ we obtain

$$(g^*)'(t) = -\gamma g^*(t) + s^*(t)\left[M(g^*(t)) - M(g^*(T)) - [v'(i^*(t))i^*(t) - v(i^*(t))]\right].$$

We next argue that we have that $g^*(t) > g^*(T) > 0$ for all $t \in [0,T)$. Let $t \in [0,T)$ and suppose that $g^*(t) \leq g^*(T)$. Since M is non-decreasing we have $M(g^*(t)) \leq M(g^*(T))$. Convexity of v ensures that $v'(i^*(t))i^*(t) - v(i^*(t)) \geq 0$. Moreover, by Lemma 11 we have $g^*(t) > 0$ and consequently (6) implies that $(g^*)'(t) < 0$. This contradicts continuity of g^* and therefore proves $g^*(t) > g^*(T) > 0$ for all $t \in [0,T)$.

Next, we verify the upper bounds in (7). To this end let $Q\colon [0,1] \to \mathbb{R}$ satisfy $Q(i) = v'(i)i - v(i)$. Note that convexity of v implies that Q is non-decreasing. Note that (6), monotonicity of M and the fact that $g^*(t) > g^*(T)$ ensure that

$$\begin{aligned}(g^*)'(t) &= -\gamma g^*(t) + s^*(t)\left(M(g^*(t)) - M(g^*(T)) - [v'(i^*(t))i^*(t) - v(i^*(t))]\right)\\ &> -\gamma g^*(t) - s^*(t)Q(i^*(t))\\ &\geq -\gamma g^*(t) - Q(1).\end{aligned}$$

Gronwall's lemma shows that

$$g^*(t) < e^{(T-t)\gamma}g^*(T) + \frac{Q(1)}{\gamma}\left(e^{(T-t)\gamma} - 1\right).$$

Using that $g^*(T) = \frac{v(i^*(T))s^*(T)}{\gamma} \leq \frac{v(1)}{\gamma}$ and $Q(1) = v'(1)1 - v(1)$ we get

$$g(t) < \frac{1}{\gamma}\left[e^{(T-t)\gamma}v(i^*(T))s^*(T) + (v'(1) - v(1))\left(e^{(T-t)\gamma} - 1\right)\right] \leq \frac{1}{\gamma}\left[v(1) + v'(1)\left(e^{(T-t)\gamma} - 1\right)\right]. \tag{32}$$

Finally, take two points in time $t, t' \in [0,T]$. Then (17) shows that

$$g^*(t)\beta^*(t) - c(\beta^*(t)) \le g^*(t)\beta^*(t') - c(\beta^*(t')) \text{ and } g^*(t')\beta^*(t') - c(\beta^*(t')) \le g^*(t')\beta^*(t) - c(\beta^*(t)). \tag{33}$$

Adding these two inequalities yields that

$$(g^*(t) - g^*(t'))(\beta^*(t) - \beta^*(t')) \le 0 \tag{34}$$

Thus, $g^*(t) > g^*(t')$ implies $\beta^*(t) \le \beta^*(t')$. □

Proof of Corollary 3. The statement follows directly from parts (i) and (ii) of Proposition 2. □

Proof of Lemma 4. We first prove Item (i). By Lemma 8 there exists an optimal control $\beta^* \in \mathscr{B}$. It follows from (6) that the associated processes (s^*, i^*, g^*) satisfy (8). By assumption this solution is unique. Next, let g solve

$$g'(t) = -\gamma g(t) + s(t)\left[M(g(t)) - M(g(T)) - [v'(i(t))i(t) - v(i(t))]\right] \tag{35}$$

with terminal condition $g(T) = \frac{v(i(T))s(T)}{\gamma}$. By the same arguments as in the proof of Proposition 2 it follows that $g(t) \ge g(T) = \frac{v(i(T))s(T)}{\gamma} \ge -c'(\underline{b}+)$. Convexity of c ensures that $\beta(t) = \underline{b} \in \arg\min_{b \in B}[g(t)b + c(b)]$. Consequently the processes (s,i,g) satisfy (8) and hence coincide with (s^*, i^*, g^*). This proves optimality of β.

The proof of Item (ii) goes along the same lines and uses the upper bounds in (32). □

Proof of Proposition 5. We first prove part (i) of the proposition. By (31) we have that $(g^*)'(t) = [\gamma\lambda_1^*(t) - v'(i^*(t))]i^*(t)s^*(t)$. Since $v'(i) = \alpha$ we have by Lemma 11 that

$$\frac{\partial}{\partial t}[\gamma\lambda_1^*(t) - v'(i^*(t))] = -\gamma\eta^*(t)\beta^*(t)i^*(t) < 0\,. \tag{36}$$

This together with (31) shows that $(g^*)'$ changes its sign at most once and that this change (if existent) is from positive to negative. It follows from (36) that $(g^*)'$ can not be equal to zero on any interval. This implies that g^* is first strictly increasing and then strictly decreasing, i.e., strictly quasi-concave. Since g^* is strictly quasi-concave we obtain from Proposition 2 (ii) that β^* is quasi-convex, i.e. there exists t^* such that β^* is non-increasing on $[0,t^*]$ and non-decreasing on $[t^*, T]$.

We next proof part (ii) of the proposition. Let $t^* \in [0,T]$ be such β^* is non-increasing on $[0,t^*]$.

Covid Economics 21, 22 May 2020: 168-193

It holds for a.e. $t \in [0,T]$

$$\frac{\partial}{\partial t}\log(i^*(t)) = \frac{i'(t)}{i(t)} = \beta^*(t)s^*(t) - \gamma \tag{37}$$

The facts that β^* is non-increasing on $[0,t^*]$ and that s^* is strictly decreasing on $[0,T]$ imply that the function $t \mapsto \beta^*(t)s^*(t)$ is strictly decreasing on $[0,t^*]$. It follows that $\log(i^*)$ is strictly concave $[0,t^*]$. As log-concavity implies quasi-concavity, i^* has at most on local maximum on $[0,t^*]$. $\square$

Proof of Proposition 6. Let $v(i) = \alpha i$ and $c(\beta) = \delta(\overline{b} - \beta)$. It follows from (17) that any optimal control satisfies

$$\beta^*(t) = \begin{cases} \overline{b} & \text{if } g(t) < \delta \\ \underline{b} & \text{if } g(t) > \delta\,. \end{cases}$$

As argued in the proof of Proposition 5, g' is strictly quasi-convex which implies the result with $t_1^* = \inf\{t \geq 0 \colon g(t) \geq \delta\}$ and $t_2^* = \sup\{t \leq T \colon g(t) \geq \delta\}$.

We next argue that the fraction of infected population i^* admits at most two local maxima. We have that any optimal policy β^* is non-increasing on the interval $[0,t_2^*]$. It follows from Proposition 5 that i^* has at most one local maximum on $[0,t_2^*]$. Since β^* is constant equal to $\overline{b}$ on $(t_2^*,T]$ it follows similarly that i^* has at most one local maximum on $(t_2^*,T]$. $\square$

References

A. Abakuks. *Some optimal isolation and immunisation policies for epidemics.* PhD thesis, University of Sussex, 1972.

A. Abakuks. An optimal isolation policy for an epidemic. *Journal of Applied Probability*, 10(2): 247–262, 1973.

A. Abakuks. Optimal immunisation policies for epidemics. *Advances in Applied Probability*, 6 (3):494–511, 1974.

D. Acemoglu, V. Chernozhukov, I. Werning, and M. D. Whinston. A multi-risk sir model with optimally targeted lockdown. Technical report, National Bureau of Economic Research, 2020.

F. E. Alvarez, D. Argente, and F. Lippi. A simple planning problem for covid-19 lockdown. *Covid Economics*, 14, 2020.

A. Atkeson. What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios. Technical report, National Bureau of Economic Research, 2020.

Covid Economics 21, 22 May 2020: 168-193

R. J. Barro, J. F. Ursua, and J. Weng. The coronavirus and the great influenza epidemic-lessons from the. 2020.

H. Behncke. Optimal control of deterministic epidemics. *Optimal control applications and methods*, 21(6):269–285, 2000.

M. Bohner, S. Streipert, and D. F. Torres. Exact solution to a dynamic sir model. *Nonlinear Analysis: Hybrid Systems*, 32:228–238, 2019.

F. Clarke. *Functional analysis, calculus of variations and optimal control*, volume 264. Springer Science & Business Media, 2013.

D. de Walque, J. Friedman, R. Gatti, and A. Mattoo. How two tests can help contain covid-19 and revive the economy, 2020.

M. Dewatripont, M. Goldman, E. Muraille, and J.-P. Platteau. Rapidly identifying workers who are immune to covid-19 and virus-free is a priority for restarting the economy. *VOX CEPR Policy Portal*, 23, 2020.

M. Gonzalez-Eiras and D. Niepelt. On the optimal “lockdown” during an epidemic. Technical report, 2020.

E. Hansen and T. Day. Optimal control of epidemics with limited resources. *Journal of mathematical biology*, 62(3):423–451, 2011.

M. Kantner. Beyond just ”flattening the curve”: Optimal control of epidemics with purely non-pharmaceutical interventions. *arXiv Preprint 2004.09471*, 2020.

W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721, 1927.

S. M. Kissler, C. Tedijanto, M. Lipsitch, and Y. Grad. Social distancing strategies for curbing the covid-19 epidemic. *medRxiv*, 2020.

R. Morton and K. H. Wickwire. On the optimal control of a deterministic epidemic. *Advances in Applied Probability*, 6(4):622–635, 1974.

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

Covid Economics 21, 22 May 2020: 168-193

J. H. Stock. Data gaps and the policy response to the novel coronavirus. Technical report, National Bureau of Economic Research, 2020.

A. A. Toda. Susceptible-infected-recovered (sir) dynamics of covid-19 and economic impact. *arXiv preprint arXiv:2003.11221*, 2020.

W. K. Viscusi and J. E. Aldy. The value of a statistical life: a critical review of market estimates throughout the world. *Journal of risk and uncertainty*, 27(1):5–76, 2003.

K. Wickwire. A note on the optimal control of carrier-borne epidemics. *Journal of Applied probability*, 12(3):565–568, 1975.

K. Wickwire. Mathematical models for the control of pests and infectious diseases: a survey. *Theoretical population biology*, 11(2):182–238, 1977.

Covid Economics 21, 22 May 2020: 168-193