CENTRE FOR
ECONOMIC
POLICY
RESEARCH

CEPR PRESS

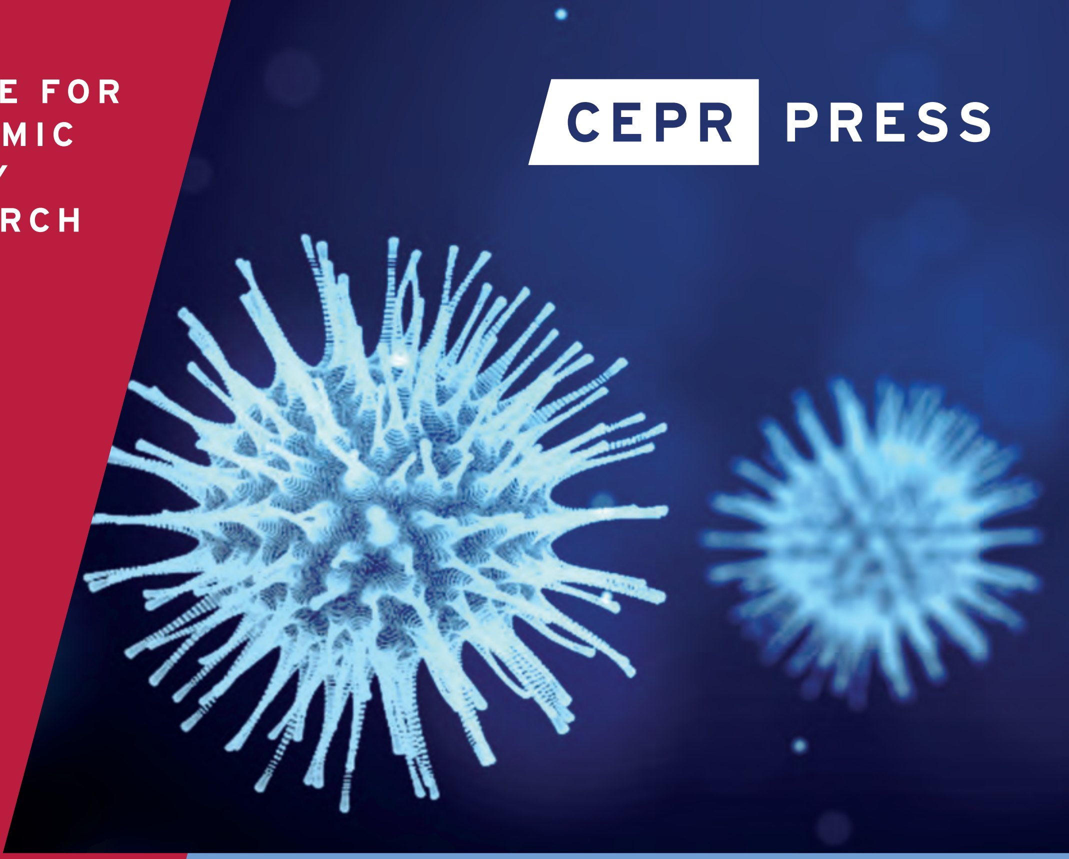

COVID ECONOMICS

VETTED AND REAL-TIME PAPERS

ISSUE 68
10 FEBRUARY 2021

MORTGAGE FORBEARANCE
Agostino Capponi, Ruizhe Jia and David Aaron Rios

SOCIAL REPERCUSSIONS
Philip Barrett and Sophia Chen

CHILDREN UNDER LOCKDOWN
Hugues Champeaux, Lucia Mangiavacchi, Francesca Marchetta and Luca Piccoli

VACCINE AUCTION
Romans Pancs

SHORT-TERM RENTAL MARKET
Elisa Guglielminetti, Michele Loberto and Alessandro Mistretta

HOUSEHOLD CASH MANAGEMENT
Fernando Alvarez and David Argente

CORPORATE CASH HOLDING IN JAPAN
Tomohito Honda and Iichiro Uesugi

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

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

Covid Economics
Vetted and Real-Time Papers

Issue 68, 10 February 2021

Contents

Covid Economics **Issue 68, 10 February 2021**

The effects of mortgage forbearance on refinancing: Evidence from the COVID-19 period[1]

Agostino Capponi,[2] Ruizhe Jia[3] and David Aaron Rios[4]

Date submitted: 6 February 2021; Date accepted: 7 February 2021

We argue that the COVID-19 foreclosure moratorium plays a crucial role in supporting refinancing activities, in addition to preventing foreclosures. We estimate that the moratorium prevented approximately 900,000 foreclosures lings and house price drops up to 9% from April to October 2020. Using loan-level data on GSE-backed mortgages, we find that the moratorium decreases the refinancing cost of households and relaxes their refinancing eligibility constraints. Our results imply that granting forbearance to households facing foreclosures has positive externalities on a broader range of households who intend to refinance. Mortgage forbearance can thus amplify the stimulative effect of monetary policy.

1 We would like to thank TianyuWang for the excellent research assistance. We are grateful to Tomasz Piskorski (discussant), Anton Korinek, Sumit Agarwal, Irena Vodenska (discussant), Chuan Du, Cynthia Balloch, and Levent Altinoglu for interesting discussions and valuable comments. We would also like to thank the seminar participants of the INET Virtual Seminar, and of the Australian Finance Banking Conference. Agostino Capponi acknowledges nancial support from the Global Risk Institute.
2 Associate Professor, Columbia University, Department of Industrial Engineering and Operations Research.
3 PhD Student, Columbia University, Department of Industrial Engineering and Operations Research.
4 Adjunct Assistant Professor, Columbia University, Department of Statistics.

Covid Economics 68, 10 February 2021: 1-46

1 Introduction

Foreclosures can have significant consequences on the real economy. A prominent example is the 2007-2008 financial crisis, where an unprecedented wave of foreclosures hit the economy. Foreclosures shocks led to large declines in house prices and residential investment (see, for instance, Mian, Sufi, and Trebbi (2015) and Piskorski and Seru (2020)). Furthermore, declines in house prices affected household borrowing through the housing collateral channel, and led to a decline in refinancing activities and consumption (see, for example, Cloyne et al. (2019)).

The recent COVID-19 pandemic imposed financial challenges on millions of homeowners in the United States, who struggled to pay their mortgages. To reduce the risk of widespread foreclosures, Congress passed the CARES Act and signed it into law on March 27, 2020. The Act offers mortgage forbearance to all mortgages backed by Ginnie Mae as well as the two government sponsored enterprises (GSEs), Fannie Mae and Freddie Mac, until June 30, and all in-progress foreclosure proceedings on such loans are paused. Later, the Federal Housing Finance Agency (FHFA) extended Fannie Mae and Freddie Mac's foreclosure moratorium until at least January 31, 2021; Federal Housing Administration (FHA) and U.S. Department of Veterans Affairs (VA) also extended the moratorium for FHA and VA guaranteed mortgages until at least Febuary 28, 2021. The moratorium grants American homeowners the option to defer mortgage payments without penalty.

In this paper, we argue that the foreclosure moratorium embedded in the CARES Act not only prevents a large wave of foreclosures but also plays a crucial role in supporting refinancing activities. Figure 1 describes the key economic mechanism behind our analysis: preventing foreclosures stabilizes house prices, which then by collateral effects, decreases the refinancing cost of households and relaxes their refinancing eligibility constraints. To evaluate the impact of the mortgage forbearance program on refinancing activities through this mechanism, there are three important, yet non-trivial questions that we must answer — how many foreclosures has the forbearance prevented? What is the impact of forbearance on house prices? How sensitive are refinancing costs and refinancing eligibility of households to house price changes?

Conducting a quantitative analysis of the mechanism outlined above and answering these three key questions is important for several reasons. First, during periods of crisis such as COVID-19, mortgage refinancing is a crucial channel for households to benefit from the stimulative effect of monetary policies. In the US, house makes up around two thirds of the median household total wealth, and fixed rate mortgages are the most dominant type of debt for the household sector (Iacoviello (2011), Campbell (2012)). As a result, in adverse economic conditions, the central bank supports mortgage refinancing through interest rate cuts and quantitative easing, which then facilitates the provision of liquidity and credit to households (see, for instance, Di Maggio, Kermani, and Palmer (2019), Keys et al. (2014)). Refinancing activities have a large impact on the real economy, especially during crisis periods. For example, refinancing reduces the interest payments of mortgagors and significantly lowers their default probability (see Agarwal et al. (2015), Agarwal et al. (2017) and Fuster and Willen (2017)). Moreover, refinancing provides liquidity to households, which effectively increases their consumption and have a large stimulus effect on the real economy (see Agarwal et al. (2015), Mian and Sufi (2011), Cloyne, Ferreira, and Surico (2019), Hurst and Stafford (2004), Canner, Dynan, and Passmore (2002) and Di Maggio et al. (2017)). However, foreclosure may dampen this stimulative effect when it is mostly needed. Second, it is important for policy makers to understand whether mortgage forbearance can support the pass-through of monetary policy through the refinancing channel. If this is the case, debt forbearance policies which target households facing foreclosures can have positive externalities on a broader range of households who intend to obtain liquidity or credit through refinancing. Such a benefit needs to be taken into consideration when designing policy intervention during crisis times.

Figure 1: The economic mechanism through which foreclosure moratorium affects refinancing activities

To study the first question, we need to evaluate the counterfactual amount of foreclosures in the absence of a forbearance program. This task is challenging for two main reasons. First, since the

Covid Economics 68, 10 February 2021: 1-46

mortgage forbearance is nationwide, a compelling quasi-experiment using regional segmentation is difficult to design. Second, mortgage forbearance also provides incentives for some households to be strategically delinquent (see, for instance, Mayer et al. (2014)) and postpone their monthly mortgage payments by 6-12 months. It is thus important for the counterfactual analysis to account for this strategic behavior. We circumvent those two difficulties by (i) rationally identifying mortgagors that could be strategically delinquent, and (ii) proposing a Markov transition model to create the counterfactual.

We use monthly, loan-level delinquency data to identify the mortgagors who are experiencing financial difficulty, and are thus unable to pay their monthly mortgages. We start by using two criteria to identify strategically delinquent mortgagors. The first criteria is based on a net present value (NPV) calculation of the future mortgage payments. We exploit the requirement that a mortgage in forbearance is not eligible for refinancing for a long period of time after which the loan becomes current again. This means that a strategic mortgagor has to make a choice between refinancing and strategic delinquency. The second criteria is based on observed Voluntary Prepayment Rates (VPRs) of mortgages with similar coupon and maturity. When the VPRs are sufficiently high, it means that loans with similar coupon and maturity are often prepaid, suggesting that the benefit of refinancing them is higher than strategic delinquency. We then estimate a Markov transition model for delinquent mortgages. Using the estimated transition rate from delinquencies to foreclosures along with the delinquent loans identified as non strategically delinquent in the first step, we find that the foreclosure moratorium has prevented approximately 900,000 foreclosures filings in the months from April through October 2020.

Next, to address the second question, we evaluate how foreclosure would have impacted house prices in the absence of a moratorium. It has been well documented that preventing foreclosures can stop the transmission and amplification of financial shocks through the housing market (see Calomiris, Longhofer, and Miles (2013), Guren and McQuade (2020), Arslan, Guler, and Taskin (2015), and Chatterjee and Eyigungor (2009)). These studies analyze the two-way channel between foreclosures and house prices: a surge in foreclosures causes a drop in house prices, and a drop in house prices precipitates more foreclosures. Due to this feedback loop, an initial surge of foreclosures could lead to a persistent drop in house prices. The findings of these studies imply that the forbearance would stabilize house prices by shutting down the spillover effect of foreclosures. We

Covid Economics 68, 10 February 2021: 1-46

confirm this stabilizing effect using monthly data and a structural vector autoregression (SVAR) model similar to the one in Calomiris, Longhofer, and Miles (2013). Our counterfactual estimate suggests that without the foreclosure moratorium, house prices would have dropped in the months from April through October 2020, respectively by 0.7%, 1.3%, 1.7%, 3.5%, 5.4%, 7.2%, and 9.4%.

To provide an answer to the third question, we investigate how the above-mentioned declines in house prices would have further affected households through the refinancing channel. Using GSE loan-level data as well as GSE eligibility requirements and the loan-level pricing adjustment (LLPA) table, we document that a decline in house prices, which reduces the value of home equity and increases households leverage, can negatively impact refinancing activities in three different ways through the collateral channel: tightening the refinancing eligibility constraint, lowering equity extraction, and increasing refinancing cost of households.

The first serious impact of house price declines on refinancing activities is that a significant portion of households will no longer be eligible for refinancing, because lenders normally require a minimal amount of home equity for mortgagors to refinance. Using loan-level data on GSE-backed, 30-year fixed-rate refinance loans originated since the start of the pandemic, we estimate that 3.3% of total non-cashout refinance loans (about 60,000) would have been ineligible to refinance through standard GSE programs in the absence of forbearance. The second consequence of a decline in house prices is less equity extraction from households as mortgagors need to maintain a minimum level of home equity after cash-out refinance. We show that the foreclosure moratorium allowed around 145,000 households (roughly 22% of total cash-out refinance) to extract around \$15,000 more, on average, from their home equity. The third negative effect of a house price decline is a much higher refinancing cost for households. Our counterfactual analysis shows that, in the absence of forbearance, the house price decline would have greatly increased the refinancing cost of households through changes in their LTV ratios. We estimate that 37% of total mortgagors who refinanced from April to October 2020 (about 900,000) would have higher LLPA fees, which on average cost those homeowners an extra amount of around \$1,400 in the form of an upfront fee, or equivalently an extra \$5,600 in interest payment over the life of the loan. We also estimate that 14.5% of the total households with non-cashout refinance (about 260,000) would have been required to purchase private mortgage insurance (PMI) due to high LTV ratios, resulting in additional cost with estimates ranging from \$1,450 to \$2,900.

Covid Economics 68, 10 February 2021: 1-46

Covid Economics 68, 10 February 2021: 1-46

We then explore the cross-sectional heterogeneity of the increase in refinancing cost. We find that the group of mortgagors with low credit scores would have been more severely impacted in the absence of forbearance, and roughly 50% of them would have incurred an extra refinancing cost, resulting in either an additional up-front fee higher than \$3,700, or an average extra interest payment of \$15,500 over the life of the loan.

We use our counterfactual to quantify the impact of mortgage forbearance on aggregate consumption through the refinancing channel. Our back-of-the-envelope calculation suggests that the additional home equity extracted and the refinancing cost saved due to the moratorium can increase households' consumption by at least 6.4 billion. This suggests that mortgage forbearance assists the delivery of liquidity and credit to households through the refinancing channel, which then supports the real economy during crisis times.

Our findings suggest that, without forbearance, the foreclosure shocks would not only affect the group of households which are unable to pay their mortgage in adverse economic times, but also have negative externalities on a wide range of households which borrow against their home equity. As a result, mortgage forbearance that prevents foreclosures can amplify the stimulative effect of monetary policy through the refinancing channel. This highlights the importance of coordinating household debt forbearance and monetary policies.

Literature Review. Our study is related to literature on the economic impact of foreclosure prevention policies. The early study of Alston (1984) analyzes how the state-level farm foreclosure moratorium legislation around 1930s affected interest rate and loan supply. Clauretie and Herzog (1990) analyze the impact of state foreclosure moratorium on loan losses. Collins and Urban (2018) explore how the New Jersey foreclosure moratorium affected lender's repayment. Gerardi, Lambie-Hanson, and Willen (2013) investigate whether the foreclosure-delay program in Massachusetts improved borrowers' outcomes. Mitman (2016) uses an equilibrium model to investigate whether the Home Affordable Refinance Program (HARP) reduced foreclosures. Agarwal et al. (2015) analyze the efficacy of HARP. They show that by helping under-collateralized mortgage borrowers refinance, HARP led to less foreclosures, more consumption, and faster house price recovery. It is worth remarking that the mechanism in their study works in the opposite direction to ours. One potential explanation is that mortgage forbearance is a remedy imposed ex-ante to directly avoid foreclosure and support refinancing. In contrast, to reverse the situation ex-post, i.e., after foreclosure shocks occurred, corrective actions should target homeowners who have not yet defaulted but could be delinquent in the future, by providing favorable conditions to refinance their mortgages. Agarwal et al. (2017) discuss the extent to which the Affordable Modification Program (HAMP) in 2008 induced lower foreclosure rates, smaller house price declines, and raised consumption. By exploiting the regional heterogeneity of the program implementation, they show that foreclosure prevention through loan modifications is indeed associated with lower declines in house prices and more durable spending. Both Eberly and Krishnamurthy (2014) and Piskorski and Seru (2018) argue about the importance of debt relief in preventing foreclosure, and design mortgage contracts that can automatically guarantee a reduction in payments during crisis times. Agarwal et al. (2020a) study borrowers' responses to mortgage forbearance during the COVID-19 period, by examining the communication records between borrowers and servicers. Emmons (2020) estimates that mortgage forbearance has prevented 500,000 foreclosures in Q4 2020. Cherry et al. (2021) analyze extensively the debt forbearance policies during the COVID-19 period, including mortgage, auto loans, student loans, and so on. They illustrate that apart from the forbearance mandated in the CARES Act, the private sector also provides substantial debt forbearance. They also argue that borrower selection in requesting forbearance results in better targeted policy. Our contribution relative to these studies is to document and quantify how mortgage forbearance could prevent foreclosure shocks from being amplified and transmitted to a wide range of households through the refinancing channel. Our findings suggest that foreclosure prevention policies can facilitate the pass-through of lower interest rate to households that intend to refinance. Hence, this benefit should be taken into consideration for the evaluation of the program's success.

We build on existing literature which has investigated the relationship between foreclosure and house prices. Several studies (see, for example Anenberg and Kung (2014), Campbell, Giglio, and Pathak (2011), Mian, Sufi, and Trebbi (2015), Harding, Rosenblatt, and Yao (2009) and Hartley (2010)) have shown that foreclosure affects house prices negatively through multiple channels. Mayer, Pence, and Sherlund (2009) show that a decline in house prices also lead to more mortgage defaults. Guren and McQuade (2020), Calomiris, Longhofer, and Miles (2013), Arslan, Guler, and Taskin (2015), and Chatterjee and Eyigungor (2009) explore the feedback mechanism between house prices and foreclosures. They show that not only more foreclosures impose a downward

Covid Economics 68, 10 February 2021: 1-46

CEPR PRESS

pressure on house prices, but also that declining house prices lead to more foreclosures. The studies of Chatterjee and Eyigungor (2009), Arslan, Guler, and Taskin (2015), and Guren and McQuade (2020), develop theoretical models to explore policy interventions aimed at reducing the spillover effect of foreclosures. In addition to providing empirical evidence supporting the theoretical predictions of these studies, we also examine how the spillover effect of foreclosures could impact mortgage refinancing of the entire household sector, which has important implications on borrowing activities of the real economy.

Our results also contribute to the stream of literature that has analyzed how changes in house prices influence refinancing. Mian and Sufi (2011) show that homeowners tend to extract their home equity when house prices go up. Bhutta and Keys (2016) show that house price growth amplifies the effect of interest rate cut on equity extraction, and vice versa. Mian and Sufi (2014) illustrate that when house prices increase, households in low-income zip codes are more likely to cash out their home equity relative to households in high-income zip codes. Beraja et al. (2018) exploit regional heterogeneity of house prices during the 2008 financial crisis, and show that regions with larger house price declines tend to have less refinancing activities after interest rate cuts. Cloyne et al. (2019) use administrative mortgage data from the United Kingdom to show that house prices affect refinancing activities through the collateral channel. We also refer to Amromin, Bhutta, and Keys (2020) for a survey on the factors that can impede refinancing activities, including credit impairment and insufficient home equity. We contribute to this strand of literature in multiple aspects. On one hand, our paper studies and specifies how house prices affect refinancing activities through the collateral channel at a micro level. We quantify how household's refinancing eligibility, equity extraction, and refinancing cost would have been impacted by a decline in house prices by means of our loan-level counterfactual analysis. On the other hand, our study emphasizes that in order to support refinancing, it is important to prevent large foreclosure shocks so to avoid negative spillover effects on house prices. Our findings suggest that debt forbearance policy aimed at preventing foreclosures can support the transmission of monetary policies through the mortgage market. In addition, our findings also complement those of Agarwal et al. (2020b) who show that mortgage refinancing during the COVID-19 pandemic may have the unintended consequence of contributing to wealth inequality.

The rest of the paper is organized as follows. We provide institutional details of the mortgage

Covid Economics 68, 10 February 2021: 1-46

market in Section 2. We describe the data set in Section 3. In Section 4, we present our main results on the impact of forbearance on refinancing activities through prevention of foreclosures and stabilization of house prices. In Section 5, we provide a robustness analysis. We conclude in Section 6.

2 Institutional Structure of the Mortgage Market

In this section, we provide institutional details. Section 2.1 gives a brief overview of the mortgage market. Section 2.2 discusses the delinquency and foreclosure process, as well as mortgage refinancing. Section 2.3 examines mortgage forbearance along with changes in delinquencies and foreclosures observed during the COVID-19 pandemic period.

2.1 Overview of the Mortgage Market

Mortgages are the second largest sector of the U.S. fixed-income security market behind U.S. Treasury debt. The total outstanding unpaid balance of American mortgages on 1-4 family homes is currently roughly \$11.2 trillions, according to the Federal Reserve Z1 table released in June 2020.[1] Most of these mortgages are fixed rate, that is, the monthly coupon and payment are constant, and some are adjustable rate mortgages (ARMS), which retain monthly payments but whose interest payment is reset periodically. The most common term for mortgages is 30 years (amounting to 360 payments in total), but other terms do exist and 15 years is the second most common. It has been estimated that more than about 90% of the U.S. mortgages are 30-year, fixed-rate mortgages (see Campbell (2012)).

Mortgage loans can be held by the banks or other institutions that originate them, or they can be sold into the secondary market. When sold, the loans are often "wrapped" with additional credit protection. In the event of a default, the real estate backing the mortgage (the collateral) can be seized and sold to pay off the remaining debt. Any shortfall is a loss for the debt holder, not the borrower. Unlike other countries, in the United States the borrower is not personally responsible

[1] The Z1 table is accessible at the Federal Reserve website (https://www.federalreserve.gov/releases/z1/20200611/z1.pdf). Mortgages can be made on any property. Homes built for a single family, or with 1, 2 or 3 extra apartments are referred to as 1-4 family homes. Above 4, it becomes a "multifamily" residence. The total outstanding U.S. residential mortgage debt on 1-4 family homes can be found in L.217.

Covid Economics 68, 10 February 2021: 1-46

CEPR PRESS

for the debt, i.e., there is no recourse. Defaults occur most often when the homeowner cannot sell the house above the debt level, and it is not uncommon for these shortfalls to occur. There are multiple ways mortgages can be wrapped; private mortgage insurance, deal structure, and, most importantly through the Federal Government, Ginnie Mae, or either of the two government sponsored enterprises, Fannie Mae and Freddie Mac.

Ginnie Mae provides US Government backing to roughly 18% of the outstanding balance. In addition, the two GSEs, Fannie Mae and Freddie Mac, provide credit backing to another 44% of the outstanding balance. The remaining balance is owned as raw loans (with no additional credit enhancement), or securitized through as non agency deals. These non agency deals played a significant role in the 2008 financial crisis, but have been a much smaller segment of the market in 2020.

2.2 Delinquencies, Foreclosures, and Refinancing

Mortgage payments are monthly, beginning one month after the loan is issued and continuing until the original loan amount, net of interest, is paid. Delinquencies occur when a homeowner is unable, or unwilling, to make mortgage payments. If the mortgagor misses a payment, then the loan is reported to be 30-day delinquent as of the calendar end of month. If the mortgagor is 30-day delinquent and then misses a second payment, he is counted as 60-day delinquent, and so forth. Mortgagors who have missed 3 or more consecutive payments are counted as 90+ delinquent. At any time, the mortgagor may make up for all or a few of the missing payments to bring the loan out of delinquency or to reduce the severity of its status. Under normal market conditions, a loan will go from 30-day, to 60-day, to 90-day delinquent and then into the foreclosure process. The mortgage servicer will repossess and sell the loans collateral, and use the proceeds to pay the bond holder.

When the interest rate decreases significantly, households who hold fixed-rated mortgages can refinance them, that is, paying off their existing mortgage with a new loan at a lower interest rate. Failing to refinance leads to a large amount of foregone savings, which on average costs households \$160 per month, or \$45,000 (unadjusted) over the remaining life of the loan (see Keys, Pope, and Pope (2016) for details).

There are two different types of refinancing: non-cashout refinance and cashout refinance. A

Covid Economics 68, 10 February 2021: 1-46

CEPR PRESS

cashout refinance replaces the homeowner's existing mortgage with a new mortgage whose loan amount is larger than the unpaid balance of the existing mortgage. The difference is cashed out by the homeowner, and used at his discretion. In contrast, with a non-cashout refinance, the homeowner borrows an amount which does not exceed the current remaining balance plus any additional closing cost.

Not all loans are eligible for a refinance through standard GSE programs. GSEs impose requirements, such as a maximum LTV ratio, which need to be met.[2]. Loans which do not satisfy these requirements can only refinance through other special programs, such as High LTV Refinance Option (HIRO). Additionally, there are several costs associated with refinancing, including closing fees, insurance cost, and points. Points are upfront fees charged to pay off the loan-level pricing adjustment fee (LLPA) imposed by GSE in exchange for a lower interest rate from the lender.[3] A LLPA is an up-front risk-based fee assessed to mortgagors with a conventional mortgage.

2.3 Mortgage Forbearance during the COVID 19 Pandemic

The CARES Act (see Consumer Financial Protection Bureau (2020)) provides federal relief for the United States in response to the economic impact of the COVID-19 pandemic. Section 4022 of the Act focuses on providing financial assistance to homeowners who lost income because of the pandemic. Mortgages with federal direct backing, i.e., Ginnie, or with indirect backing by the government sponsored enterprises, Fannie Mae or Freddie Mac, have been granted a moratorium on foreclosures. Foreclosures on Ginnie backed mortgages are restricted from occurring in the period from March 18 till February 28, 2021, and foreclosures on Fannie Mae or Freddie Mac backed mortgages are prohibited until January 31, 2021. At present, the homeowner is expected to make the missing payments when the period ends, but that aspect of the bill is currently up for debate, with some politicians calling for partial forgiveness.

As COVID-19 spreads across the United States, mortgage delinquencies surge. Figure 2 provides a historical perspective and shows the national serious delinquency rate[4] and foreclosure rate in United States since 1990. It is evident from the Figure that prior to the implementation of

[2]See the GSE Eligibility Matrix at `https://singlefamily.fanniemae.com/media/20786/display`

[3]According to the report of Fannie Mae and Freddie Mac Guarantee Fee Review in 2015, 25 basis point upfront charge is approximately equivalent to 5 basis points of the current mortgage rate.

[4]Serious delinquency is defined as more than 90 days delinquent.

Covid Economics 68, 10 February 2021: 1-46

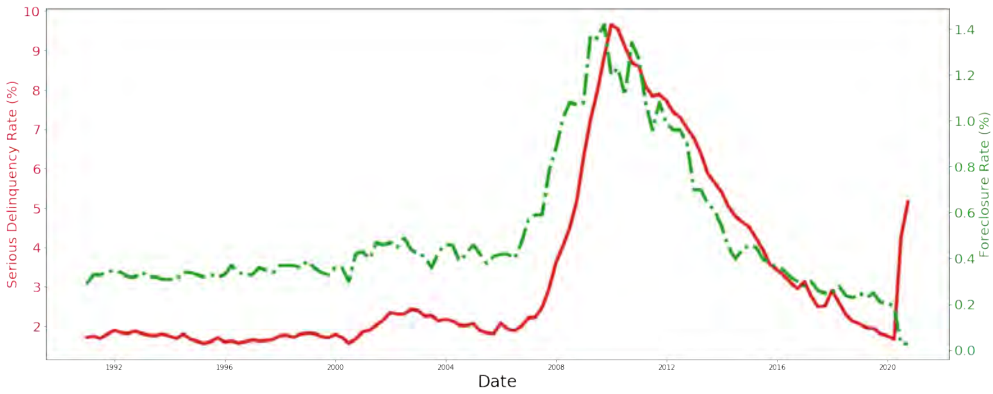

Figure 2: This plot shows the quarterly series of serious delinquency rates (solid) and foreclosure rates (dashed) in US. Source: Mortgage Bankers Association

mortgage forbearance in March 2020, there exists a strong positive correlation between serious delinquency and foreclosure rates: the correlation between the two series is 0.89. Moreover, after the implementation of the CARES Act in March, there has been no observance of a surge in foreclosures following the surge of delinquencies. In particular, serious delinquencies rates in Q2 2020 reached the highest level since the 2008 financial crisis, but, at the same time, foreclosure rates were at almost the lowest level since 1990. These observations support the intended objective of the foreclosure moratorium: foreclosures exhibit a large drop despite the large increase in delinquencies.

3 Data

We use three categories of data in our analysis: delinquency and foreclosure data, house prices data, and agency mortgage-backed security loan-level data.

House Prices: We use the seasonally adjusted House Price Index (HPI) from FHFA[5]. The FHFA HPI is a weighted, repeat-sales index which measures average price changes in repeat sales

[5]Our choice of using the FHFA House Price Index instead of the Case-Shiller Index is due to the following reasons. First, our loan-level counterfactual analysis is based on GSE-backed loans. FHFA's valuation data are calculated directly from conforming mortgages provided by GSEs, which is well suited for estimating the counterfactual LTV ratios of GSE-backed loans without the foreclosure moratorium. Second, the monthly Case-Shiller Index is based on a 3-month moving average, whereas the FHFA HPI reflects the house prices in each month. Our counterfactual analysis is based on monthly house prices rather than moving averages.

Covid Economics 68, 10 February 2021: 1-46

CEPR PRESS

on the same properties. It covers all single-family properties whose mortgages have been purchased or securitized by Fannie Mae or Freddie Mac since January 1975. This index is published with a two-month lag. At present, the last month of release was October 2020.

Foreclosures and Delinquencies: We use the U.S. Home Foreclosure Filings Total from Realty Trac Inc. to count properties at different stages of the foreclosure process. Those stages include receiving a notice, auction, and actual foreclosure, in each month from 2005-01 till 2020-10.

We also use information about delinquency and foreclosures by Ginnie Mae. For each loan backed by Ginnie Mae, the loan-level data records its delinquency status in each month since 2013-11, such as current, 30-day delinquent, 60-day delinquent, and 90-day delinquent. For each month, the dataset also provides the next payment code for every loan. This allows determining whether in the following month the loan exited the delinquency state due to foreclosure, repurchase, or loss mitigation. Moreover, the dataset indicates whether a loan is in forbearance for each month. Apart from the delinquency and foreclosure status, the dataset also provides detailed, monthly information on the loan and borrower characteristics (coupon, maturity, location, LTV ratio, FICO scores). The source of this data are the third-party vendors eMBS and Recursion.

Agency Mortgage-Backed Security Loan-Level Data: Our analysis of refinancing is based on GSE-backed, 30-year fixed-rate, refinance loans originated from April 2020 to October 2020. We source this data from Recursion. According to the report released by the Urban Institute in Oct 2020[6], 65.2% of the total originations in the second quarter of 2020 is GSE-backed, and 68% to 72% of them are refinance loans. There are approximately two million observations in our dataset.

Our Loan-level data provides information for each loan, with a code to map it to a pool. The data include static information about the origination of the loan, such as state, origination month, credit score, loan-to-value (LTV) ratio, debt to income ratio, coupon rate, loan size, originator, and loan purpose (refinancing, cash-out refinancing, new purchases, or others). Additionally, it includes information that changes monthly, such as the remaining balance, delinquency status, and mortgage servicer.

Table 1 provides summary statistics of our sample. The distribution of LTV ratio and credit score across homeowners is especially important for our study. Panels 3a and 3b in Figure 3 illustrate

[6]See `https://www.urban.org/research/publication/housing-finance-glance-monthly-chartbook-october-2020/view/full_report`

Covid Economics 68, 10 February 2021: 1-46

CEPR PRESS

the distribution of LTV ratio among cash-out and non-cashout refinance loans, respectively. Most of the loans have LTV ratios between 60% and 80%. In particular, all cashout refinance loans have LTV ratios lower than 80% because this is the standard eligibility requirements imposed by GSE. Panels 3c and 3d in Figure 3 illustrate the distribution of credit score among cashout and non-cashout refinance loans, respectively. Most of the loans have a credit score above 740.

Table 1: This table provides summary statistics for GSE-backed, 30-year, fixed-rate, refinance loans originated from April 2020 to October 2020. The first column presents statistics of non-cashout refinance loans, and the second column presents statistics of cash-out refinance loans. LTV = loan to value .

	Cash-out Refi	Non-Cashout Refi
Item Count	646,031	1,813,536
Weighted Average Gross Coupon	3.34	3.12
Weighted Average Maturity	356.46	355.93
Weighted Average Original Loan Size ($)	347,258.14	366,824.62
Orig LTV Weighted Average	64.82	69.07
Orig LTV Weighted Standard Deviation	13.30	14.79
The 10th percentile of LTV ratio	45	48
The 25th percentile of LTV ratio	57	60
The 50th percentile of LTV ratio	68	72
The 75th percentile of LTV ratio	75	80
The 90th percentile of LTV ratio	80	87
Credit Score Weighted Average	754.35	764.58
Credit Score Weighted Standard Deviation	41.41	37.83
The 10th percentile of Credit Score	706	692
The 25th percentile of Credit Score	740	724
The 50th percentile of Credit Score	772	760
The 75th percentile of Credit Score	794	787
The 90th percentile of Credit Score	807	802

Covid Economics 68, 10 February 2021: 1-46

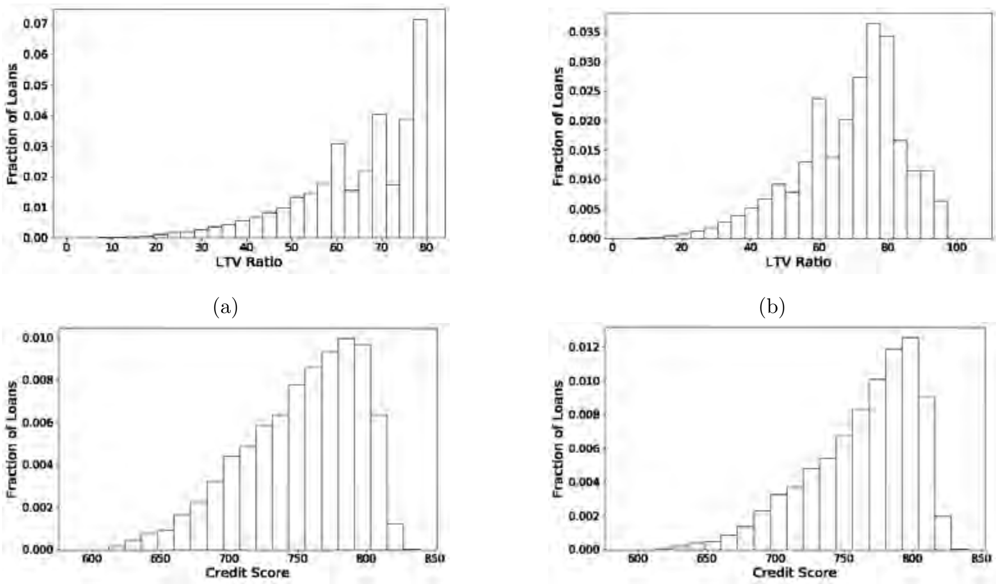

(a) (b)

(c) (d)

Figure 3: Distribution of the LTV ratio and credit score for GSE-backed, 30-year, fixed-rate, refinance loans originated from April 2020 to October 2020. Panels (a) and (c) report the distribution of LTV ratio and credit score for cash-out refinance. Panels (b) and (d) report the distribution of LTV ratio and credit score for non-cashout refinance

Covid Economics 68, 10 February 2021: 1-46

4 The Impact of Forbearance on Foreclosures, House Prices, and Refinancing

In this section, we analyze the implications of the foreclosure moratorium on the housing and mortgage market. Specifically, we provide an answer to the three main questions highlighted in the introduction — how many foreclosures has the forbearance prevented? What is its impact on house prices? How sensitive are refinancing costs and refinancing eligibility of households to house price changes?

In Section 4.1, we use a Markov transition model to quantify the counterfactual level of foreclosures in the absence of intervention. In Section 4.2, we estimate the impact of mortgage forbearance on house prices. In Section 4.3, we analyze the mechanism through which forbearance not only benefits directly households facing foreclosures, but also indirectly those who intend to refinance. In Section 4.4, we conduct a sensitivity analysis and show that a slight decline in house prices can have a large negative impact on refinancing activities.

4.1 Counterfactual Level of Foreclosures

We conduct a counterfactual analysis to assess what the amount of foreclosures would have been in the absence of forbearance. There are two main challenges imposed by such an analysis. First, we need to account for the strategic behavior of mortgagors, because forbearance may create perverse incentives for households who may strategically suspend their monthly payments for 6-12 months, and repay the missed mortgage payments in the future. Second, because mortgage forbearance is mandated nationwide, it is nontrivial to exploit regional segmentation and design a compelling quasi-experiment. Hence, a model is required to generate the counterfactual. It may be argued that the segmentation in the mortgage markets between conforming loan and jumbo loans can be used to design a quasi-experiment, especially when the forbearance mandated in the CARES Act is only for conforming loans. However, jumbo segment may not be a good candidate for a control group. This is because the loan and borrower characteristics for jumbo loans are very different from those of conforming loans. Moreover, as shown in Cherry et al. (2021), private lenders for jumbo loans also provide substantial mortgage forbearance. Thus, both segments are treated by forbearance and have a much lower level of foreclosure activities than before the pandemic. As a

Covid Economics 68, 10 February 2021: 1-46

result, this segmentation may not provide enough variation.

We address the first concern identifying mortgagors who are likely to be strategically delinquent based both on a NPV calculation and on empirically observed prepayment rates. We then address the second concern by using a finite-state Markov transition model to generate the counterfactual foreclosures.

For each Ginnie Mae backed, 90-day delinquent mortgage in the forbearance state, we calculate the NPV of the future mortgage payments at the time the loan enters forbearance, if the mortgagor chooses to be strategically delinquent and postpones mortgage payments by one year. We then compare it with the NPV of the future mortgage payments if the mortgagor refinances at the prevailing interest rate at the time the loan enters forbearance. We discount using the 10-year treasury rate in each month and assume an up-front cost equal to 4% of the total loan size[7]. A mortgage in forbearance is not eligible for loan refinancing, and it will typically not be available for a refinance at least for one or two years[8] after the loan becomes current again. As a result, those who can save more money from refinance have little incentive to be strategically delinquent on their mortgage payments. Hence, we identify those loans as being very unlikely to enter the forbearance program because of strategic delinquency. This suggests that even without forbearance, those loans would still have been delinquent. [9]

Moreover, to verify that our methodology is consistent with observed prepayment rates, we examine the VPRs of mortgages, whose coupon and maturity are similar to the loan identified as not strategically delinquent in the first step of the procedure, at the month it enters forbearance. When the VPRs are sufficiently high, it means that loans with similar characteristics to those identified as not strategically delinquent are indeed often prepaid, suggesting that the benefit

[7]Given that the average loan size is over $300,000, a 4% up-front cost is a conservative parameter choice compared to existing literature, such as Keys, Pope, and Pope (2016), which assumes a transaction cost equal to $2000 + 1% of the loan size.

[8]FHA has relaxed this requirement to 3-6 months around August 2020. However, this does not affect the loans which are part of our calculation, because they have already entered forbearance before this change occurred. Moreover, availability to refinance does not guarantee that lenders will be willing to lend to mortgagors at normal rates if they ever enter forbearance.

[9]This calculation does not apply to adjustable-rate mortgages (ARM) as it automatically adjusts the mortgage rate when interest rates go down, resulting in smaller savings from refinance. To be conservative, we assume that all ARMs in forbearance may be strategically delinquent.

Covid Economics 68, 10 February 2021: 1-46

of refinancing them is higher than being strategically delinquent. Instead, if the corresponding annualized VPRs are less than 20%[10], we remove this loan from the set of those identified as not strategically delinquent.

Based on our NPV calculation and observed prepayment rates, we estimate that around 90.6% of the mortgagors who are behind their mortgage payments are unlikely to be strategically delinquent. Restricting our sample to those households allows for a better estimates of how many household would still face foreclosure in the absence of forbearance. Admittedly, our restriction is imperfect. For example, some households may not make their financial decisions fully rationally.

We then use a finite-state Markov transition model to estimate how many foreclosures would have occurred without forbearance. We first collect the monthly delinquency state transition data for all Ginnie Mae backed mortgages in the period from November 2013 till October 2020. In every month, each loan then belongs to one the following states of the Markov chain: performing, 30-day, 60-day, 90-day delinquent, foreclosed, repurchased, paid-off, and loss mitigation. In particular, the states of foreclosures, repurchased, paid-off, and loss mitigation result in the removal of the loan from the pool, i.e., those states are absorbing states in the Markov chain. The specification of the Markov transition model is provided in Appendix A.

Figure 4 illustrates transition from delinquency to foreclosure in the Markov transition model. Panel 4a illustrates that without forbearance, loans that are 90-day delinquent can transit to foreclosures with a certain transition probability. In contrast, panel 4b shows that, with forbearance, loans that are 90-day delinquent will no longer be foreclosed. Therefore, to construct the counterfactual estimation of foreclosures in the absence of forbearance using the Markov transition model, we need to (1) identify 90-day delinquent loans that are not strategically delinquent and would remain delinquent without forbearance and (2) estimate the transition rate from 90-day delinquency to foreclosure.

As we have already identified loans which are not strategically delinquent, we next estimate the transition rate from 90-day delinquency to foreclosure in the absence of forbearance. To control for loan-level characteristics, for each month we divide loans into different groups according to their

[10]After identifying strategically delinquent loans based on the NPV calculation, the results in the second step are not sensitive to the thresholds choices. This suggests that the NPV calculation is consistent with the observed prepayments rates.

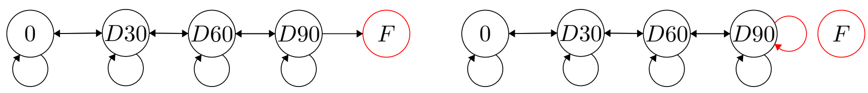

(a) Without Forbearance

(b) With Forbearance

Figure 4: Panel 4a illustrates that without forbearance, loans that are 90-day delinquent (D90) can transit to foreclosures (F) with a certain transition probability. Panel 4b illustrates that, with forbearance, loans that are 90-day delinquent will no longer be foreclosed. Therefore, to construct the counterfactual estimation of foreclosures in the absence of forbearance, we need to identify 90-day delinquent loans that are not strategically delinquent and estimate the transition rate from 90-day delinquency to foreclosure.

current LTV ratios (30, 60, 80, 100, 120), credit scores (600, 660, 740), Ginnie Mae programs (FHA, VA, PIH, RD), and states (judicial/non-judicial). Following Jarrow, Lando, and Turnbull (1997), for each group we estimate the transition probability from 90-day delinquency to foreclosures using the sample prior to the implementation of forbearance in the following ways:

$$p_{(D90,F)} = \frac{\sum_t N_{(D90,F)}(t)}{\sum_t N_{D90}(t)},$$

where $N_{(D90,F)}(t)$ is the number of transitions from 90-day delinquency to foreclosure in month t, $\sum_t N_{(D90,F)}(t)$ is the total number of loans that transited from 90-day delinquency to foreclosure prior to the implementation of forbearance, and $\sum_t N_{D90}(t)$ is the total number of loans in the state 90-day delinquency before the implementation of mortgage forbearance.

We then estimate the expected foreclosures in the absence of forbearance using our estimated transition rate from 90-day delinquency to foreclosure and the number of delinquencies identified as non strategic. The details of the calculation are reported in Appendix A. Figure 5 plots the estimated number of foreclosures with and without the foreclosure moratorium. In the absence of intervention, foreclosures would have shot up since April whereas in the same period, the actual foreclosures have declined to the lowest level in the past 15 years. Table 2 compares our counterfactual estimation of foreclosures with the actual foreclosures. We find that in the period from April to October 2020, the foreclosure moratorium has prevented approximately 900,000 foreclosure filings that might have occurred without intervention. Clearly, the forbearance assisted those homeowners whose homes could have been foreclosed upon to avoid both financial costs and significant long-term non-pecuniary costs (see Diamond, Guren, and Tan (2020)).

Covid Economics 68, 10 February 2021: 1-46

It is worth remarking that our estimates are conservative. First, we cannot rule out the possibility that, without the foreclosure moratorium, the transition rate from 90-day delinquency to foreclosure would have increased relative to the period before COVID-19 due to the economic impact of the pandemic. Second, serious delinquent mortgages could be repurchased by mortgage servicers from Ginnie Mae, and the foreclosure of such mortgages would not be reported afterward. This would lead to an underestimation of the transition rate from delinquency to foreclosure without the moratorium. Third, when we identify the strategically delinquent loans, our parameter choices, such as the 4% up-front refinancing cost, are conservative compared to other studies. Given that our main objective is to study how forbearance may prevent foreclosure shocks from propagating to a wider range of households and impact refinancing activities as well as aggregate consumption, it is conceivable to be conservative regarding the magnitude of the foreclosure shock.

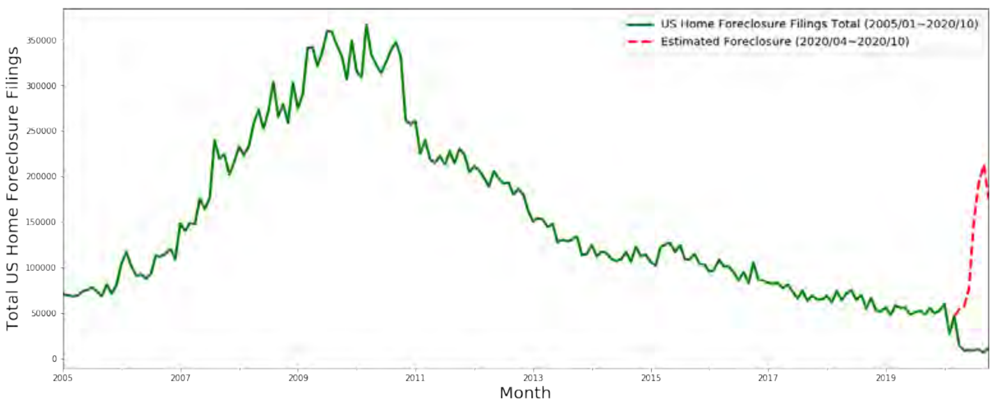

Figure 5: Number of estimated total US foreclosures filings with the foreclosure moratorium of the CARES Act (red dashed) and in the hypothetical scenario of no moratorium (green solid).

4.2 The Impact of Forbearance on House Prices

In the previous sections, we have analyzed the extent to which forbearance serves to avoid foreclosures. There is however, another important mitigating effect of forbearance, which is to stop foreclosure shocks from transmitting to the housing and mortgage market. In this section, we quantify the impact of forbearance on house prices.

The studies of Calomiris, Longhofer, and Miles (2013), Arslan, Guler, and Taskin (2015), and

Table 2: This table reports the counterfactual estimates of foreclosures from April to October in the hypothetical scenario of no foreclosure moratorium. For comparison purposes, the actual foreclosures from April to October are also reported in the table.

	Foreclosures:							
	Apr	May	Jun	Jul	Aug	Sep	Oct	Total
Counterfactuals	55,037	58,957	79,759	157,406	201,880	222,665	183,555	959,259
Actual Foreclosures	14,148	8,767	9,247	8,892	9,889	6,872	11,673	69,488

Covid Economics 68, 10 February 2021: 1-46

Chatterjee and Eyigungor (2009) document that an unanticipated increase in foreclosures leads to a flood of forced liquidations, which depresses house prices. Conversely, a drop in house prices precipitates more foreclosures. As a result of this self-reinforcing feedback loop, without any intervention, an initial surge of foreclosures turns into a lasting problem, i.e., a persistent drop in house price growth and a persistent increase in foreclosures, like the one observed in the years between 2007 and 2011.

The moratorium prevents foreclosures, and as a result, temporarily shuts down the negative spillovers to house prices. In order to estimate how much house prices would have dropped after foreclosure shocks and without intervention, we construct an econometric model to measure house price fluctuations associated only with exogenous shocks to foreclosures. We consider a five-variable VAR model, similar to the one in Calomiris, Longhofer, and Miles (2013). The biggest difference between their model and ours is that we use a standard VAR based on monthly data from February 2005 to March 2020, while theirs is a Panel VAR based on quarterly data. We deviate from their model and consider a higher frequency for two main reasons: (i) quarterly frequency is too coarse for our counterfactual analysis on refinancing, whereas monthly frequency allows us to calculate the counterfactual decline in house prices for each month from April 2020 to October 2020 without foreclosure moratorium, and the corresponding counterfactual LTV ratios of refinance loans originated in those months; and (ii) we can assess the robustness of our model by comparing implied house price declines at a quarterly frequency with those estimated from their model (see Section 5 for the comparison). The detailed specification of the SVAR model is reported in Appendix B.

We use the SVAR model to estimate how much house prices would have declined without the foreclosure moratorium. We assume that the moratorium only exerts structural shocks specific to foreclosures. Using our counterfactual estimation of foreclosures done in Section 4.1 and actual foreclosures, we can identify the magnitude of shocks that the moratorium imposes on foreclosures every month. We then use the identified shocks to calculate their impact on house prices. The details of the calculation are included in the Appendix C. We find that without the foreclosure moratorium, house prices would have dropped by 0.7%, 1.3%, 1.7%, 3.5%, 5.4%, 7.2%, 9.4% relative to the realized values from April through October 2020.[11] Figure 6 compares the actual house prices with their counterfactual estimates. In the hypothetical scenario of no foreclosure moratorium, the counterfactual House Price Index exhibits a large, persistent drop since April. In contrast, the actual House Price Index shows a persistent increase during the same period. This comparison is consistent with our claim that the foreclosure moratorium prevents a tidal wave of foreclosures and stabilizes house prices.

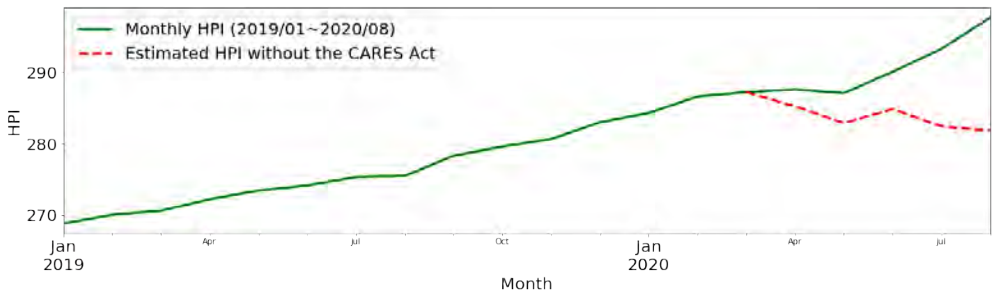

Figure 6: The actual House Price Index (green solid line), and the counterfactual estimate of the House Price Index without the foreclosure moratorium (red dashed line).

[11]To put those numbers into context, we observe that house prices have increased by 1% from February to March 2020. Without such an increase, house prices in March would have been 1% lower than the realized ones. As a result, prices would have stayed 1% lower relative to their realization every month since March. Moreover, house prices rose 8% from August 2019 to August 2020, a growth which would have been completely eroded without intervention.

CEPR PRESS

4.3 The Impact of Forbearance on Refinancing

In this section, we analyze the impact of mortgage forbearance on refinancing activities. What is the mechanism through which forbearance, whose mandate is to prevent foreclosure shocks, supports refinancing activities? By stabilizing house prices, forbearance prevents a wide range of homeowners' LTV ratios from increasing. This in turn, by collateral effect, increases the number of homeowners eligible for refinancing, leads to more equity extraction through cash-out refinancing, reduces the up-front fee for refinancing, and lowers interest payments. In this way, forbearance not only benefits the group of households facing foreclosures, but also exerts positive externalities on all households who intend to borrow against their home equity. Using our loan-level data set, we conduct a counterfactual analysis on how forbearance impacts homeowners' refinancing cost, equity extraction, and refinancing eligibility. We show that forbearance helps pass credit and liquidity to homeowners through the refinancing channel during the COVID-19 period.

As shown in Section 4.2, house prices would have decreased in the absence of intervention. Using the counterfactual house prices from each month, we can calculate the counterfactual LTV ratios for all 30-year, fixed-rate, GSE-backed refinance loans originated from April 2020 to October 2020. Without forbearance, the homeowners' LTV ratios would have increased because of the decline in home equity. As discussed next, due to the collateral effect, this has negative impact on households who seek to refinance along three main dimensions

Firstly, a significant portion of mortgagors would have been ineligible to refinance in the absence of forbearance. Without the moratorium, among all 30-year, fixed-rate, GSE-backed refinance loans originated from April 2020 to October 2020, 3.3% of non-cashout refinance loans (about 60,000) would have not satisfied the GSE Standard Eligibility Requirements due to high LTV ratios. That is, those loans' LTV ratios would have been above the maximum for traditional conventional refinances, and they could have been refinanced only through special programs such as Fannie Mae High-LTV Refinance Option (HLRO). We apply the Fannie Mae Eligibility Matrix[12] to every GSE-backed non-cashout refinance loans originated since April and check if it would still be eligible for conventional refinances in the hypothetical scenario of no forbearance and declining house prices. The ineligibility

[12] The Eligibility Matrix provides the comprehensive LTV, CLTV, and HCLTV ratios requirements for conventional first mortgages eligible for delivery to Fannie Mae (same eligibility will apply to Freddie Mac).

Covid Economics 68, 10 February 2021: 1-46

CEPR PRESS

may potentially make those homeowners forgo a large amount of interest savings from refinancing. In other words, forbearance stabilizes house prices, makes a large number of homeowners eligible for traditional conventional refinances, and saves them money from interest payments. According to Agarwal et al. (2017) and Fuster and Willen (2017), refinancing significantly lowers their probability to default in the future and prevents more foreclosure shocks. Secondly, there would have been much less equity extraction from mortgagors in the absence of forbearance. Among all 30-year, fixed-rate, GSE-backed cash-out refinance loans originated since April, around 145,000 of them (about 22%) would have been forced to decrease the amount of their cashed-out home equity, and on average, they would have decreased the amount of equity cashed out by around $15,000. This is because without forbearance, deteriorating housing prices would lead to losses in home equity for homeowners, who would be forced to decrease their home equity extraction. For each GSE cash-out refinance loan, this decrease in equity extraction is calculated by taking the positive part of the difference between the total amount of principal borrowed and the maximum amount of principal[13] that the homeowner could have borrowed in the hypothetical scenario of no forbearance and declining home equity. For example, a homeowner with a 1-unit, principal residence worth $300,000 borrows $240,000 thorough cash-out refinance to pay off the remaining unpaid principal of $200,000 from their existing mortgage, and then takes out $40,000 in cash. However, if house prices were to decline to $280,000 in the scenario without intervention, this homeowner could only borrow at most $224,000 thorough cash-out refinance. This is because the current maximum LTV ratio eligible for standard conventional cash-out refinance is 80% (for 1-unit, principal residence). As a result, the homeowner could only extract $24,000 in cash from their home equity, i.e., $16,000 less than in the scenario where house prices are stabilized through forbearance. Moreover, if like in Di Maggio, Kermani, and Palmer (2019) and Mian and Sufi (2011), we assume that the marginal propensity to consume out of cashed-out equity is 1, the decline in equity extraction for the affected households would imply a decrease in total consumption by about $2.2 billion. In other words, forbearance stabilizes house prices, allows a large number of homeowners to extract more cash from their home equity, helps passing liquidity to those households, and increases their

[13]Based on the characteristics of each loan (LTV, credit score, etc...), we use the Eligibility Matrix to find the corresponding maximum LTV ratio above which the loan does not qualify for a standard conventional cash-out refinance.

Covid Economics 68, 10 February 2021: 1-46

CEPR PRESS

consumption.

Thirdly, without forbearance, the refinancing cost for homeowners would increase drastically, especially for those with low credit score. The underlying reason is that homeowners would have higher LTV ratios, making it riskier to guarantee their loans, so higher loan-level pricing adjustment (LLPA) fees would be charged by GSE. The fee would be passed to the homeowners, either in the form of up-front fee (so-called points), or in the form of higher coupon. We use the LLPA Matrix[14] for all the conventional refinance loans originated since April, and calculate the difference between the current LLPA fees and the LLPA fees that GSE would have charged in the hypothetical scenario of no forbearance. We find that 37% of the homeowners who refinanced between April and October 2020 (about 900,000) would be charged a higher LLPA fee, and the increase in upfront fee would be around $1,400 on average. If the LLPA fees are rolled into interest payment, this would be equivalent to an increase of $5,600 in interest payments over the life of the loan. According to the report of Fannie Mae and Freddie Mac Guarantee Fee Review in 2015[15], a 25 basis points upfront charge is approximately equivalent to 5 basis points of the current mortgage rate. It is worth noting that a higher upfront fee would sharply decrease homeowners' incentive to refinance, according to Keys, Pope, and Pope (2016) and Agarwal, Driscoll, and Liabson (2013), and a $1000 increase in upfront fees might reduce the refinancing threshold by 25 basis points.

In addition to a higher LLPA fee, higher LTV ratios due to declining house prices would have also forced many mortgagors to purchase private mortgage insurance (PMI) for their mortgages. We find that without intervention, 14.7% of all GSE-backed non-cashout refinance loans originated since April (about 260,000) would have needed to purchase private mortgage insurance (PMI). A PMI is required for mortgages with LTV ratios larger than 80%, and its annual cost is approximately between 0.4% and 0.8% of the borrowed amount according to Freddie Mac. Once the mortgages' LTV ratios fall below 80%, it is optional to maintain the insurance. For each non-cashout refinance loan originated since April, we first check whether it would have needed a PMI in the hypothetical scenario of no forbearance and declining house prices, and then calculate the number of insurance

[14]The Loan-Level Price Adjustment (LLPA) Matrix provides the LLPAs applicable to loans delivered to GSE. LLPAs are assessed based upon loan features, such as credit score, loan purpose, occupancy, number of units, product type, etc. See `https://singlefamily.fanniemae.com/media/9391/display` for details

[15]In their website, Freddie Mac claims that one point which costs 1% principal in the form of up-front fee can approximately reduce the mortgage rate by about 0.25%. In other words, our estimation is on the conservative side.

Covid Economics 68, 10 February 2021: 1-46

months and the approximate cost. We find that without foreclosure moratorium, the affected homeowners would have paid an extra amount ranging from $1,450 to $2,900 on average for a PMI.

There is a large degree of heterogeneity for the increase in refinancing cost caused by the variation in homeowners' credit scores. We break down the increases in refinancing cost into different groups based on the homeowners' credit score at the origination of the loan. The percentage of households that would have incurred a higher LLPA cost and the average increase in refinancing cost are much higher for groups with lower credit scores. Table 3 shows that in the hypothetical scenario of no forbearance, only 35% of the households with credit score higher than 740 would have had a higher LLPA cost, and the average refinancing cost increase for this group would have been only $1,264 up-front or $4,987 in terms of interest payment. However, for groups with credit scores lower than 680, more than 50% of the households would have incurred a higher LLPA cost, and the average refinancing cost increase would have been a staggering number—more than $3,700 as up-front fee or more than $15,000 as interest payments. This is almost three times the extra cost incurred by mortgagors with high credit score. In other words, without the forbearance stabilizing effect on house prices, households with low credit scores might have had a much larger refinancing cost, and forced to forego the refinancing option completely.

The high sensitivity of refinancing costs to declines in house prices for households with low credit score also provides an additional explanation for the phenomenon studied by Keys, Pope, and Pope (2016): households with low credit score have a higher failure-to-refinance rate.[16] That is, for those households, even though refinancing seems optimal at the prevailing interest rate, they still tend to not refinance their existing mortgage. Our analysis suggests that the refinancing cost for low credit households could be much larger than for households with high credit score after the decline in house prices. However, refinancing opportunities due to lower interest rates often emerge during or after economic downturns together with decline in house prices, which makes households with low credit score less likely to refinance.

Following Di Maggio et al. (2017) and Di Maggio, Kermani, and Palmer (2019), we assume that the average marginal propensity to consume out of a reduction in monthly mortgage payments is around 0.75. Our calculation shows that forbearance, by reducing households' refinancing cost, increases the consumption of mortgagors by another $4.3 billion during the first nine months

[16] Keys, Pope, and Pope (2016) analyze the failure-to-refinance rates for different households in 2010.

Covid Economics 68, 10 February 2021: 1-46

Table 3: This table reports estimates of the fraction of mortgage borrowers who could have been subject to a higher LLPA fee without intervention, and the average increase in refinance cost either in the form of up-front fee or interest payment, across a range of borrower credit score characteristics. The sample consists of 30-year, fixed-rate, refinance loans originated from April 2020 to October 2020. Following the numbers provided in the Guarantee Fee Review report of Fannie Mae and Freddie Mac released in 2015, we assume that a 25 basis point upfront charge is approximately equivalent to a 5 basis points raise in current mortgage rate.

FICO Score	Item Count	Share with increase in LLPA fees	Average LLPA cost in the form of up-front fee	Average LLPA cost in the form of interest payment
$x \geq 740$	2,030,037	35.1%	\$1,264	\$4,987
$740 > x \geq 720$	198,047	44.3%	\$1,452	\$5,773
$720 > x \geq 700$	128,834	46.3%	\$1,827	\$7,289
$700 > x \geq 680$	58,824	47.0%	\$2,529	\$10,194
$680 > x \geq 660$	28804	50.8%	\$3,610	\$14,663
$660 > x \geq 640$	14085	49.0%	\$4,214	\$17,241
$640 > x \geq 620$	845	45.0%	\$4,501	\$18,443
$620 > x$	91	35.2%	\$4,124	\$16,517
Total	2,459,567	37.0%	\$1,419	\$5,626

subsequent to the intervention.

To reiterate the core message of this section, the impact of foreclosures is not limited to households who default on their mortgages. Importantly, foreclosure has negative effects on the availability of funding (i.e., refinancing eligibility and equity extraction) and raises the cost of funding of a wider range of households. This further impedes the recovery of the real economy by reducing aggregate consumption and increasing homeowners' probability of default in the near future. In this way, mortgage forbearance, a policy preventing foreclosures, can amplify the stimulative effect of monetary policies through the refinancing channel.

It is worth emphasizing that our estimates are conservative. Even though GSE-backed refinance loans make up a large of portion of the total refinance loans, there is still a significant number of refinance mortgages securitized by VA/FHA or held as portfolio loans. The forbearance also benefits borrowers holding those mortgages, because they are vulnerable to loss in home equity and decline in house prices. Moreover, there are still millions of households who would benefit from refinance but have not yet done so for various reasons, and they are likely to benefit from the foreclosure moratorium in the future.[17] Additionally, because the mortgage rate continues to decrease, there will be more households that become refinance candidates and can potentially benefit from forbearance.

4.4 Sensitivity of Refinancing to House Price Declines

In this section, we explore how sensitive the refinancing activity is to various magnitudes of declines in house prices. We calculate how the refinancing cost and eligibility of households who refinanced from April 2020 to October 2020 would be affected, under a variety of house price decline assumptions.

As shown in Table 4, the impact of the counterfactual house price decline estimated in Section 4.2 is most comparable to a naïve house price decline of 3% relative to the realized house price. The first three rows in Table 4 indicate that the share of non-cashout loans which requires extra insurance and the corresponding cost are very sensitive to the magnitude of house price declines. A

[17]There are still 18.5 million households who meet the refinancing criteria and can reduce their annual interest payment by at least 0.75% by refinancing at the mortgage rate of September (see the Black Night September Mortgage Report). This can, on average, save them $304 per month.

Covid Economics 68, 10 February 2021: 1-46

10% decline in house prices would lead to around 30% of loans requiring additional insurance at an average cost of \$6,712, while a 1% decline in house prices would lead to only around 7.3% of loans requiring additional insurance with an average cost of \$1,186. This is because the larger the house price decline, and the higher the LTV ratios of these loans. As a result, it will be more likely for the loans to require extra insurance, and the insurance period until the LTV ratios fall below 80% will also be longer. For the same reason, a larger house price decline will make a larger share of the loans ineligible to refinance through the standard conventional program. A 10% decline in house prices could have made 11.09% of the refinance loans from April 2020 to October 2020 ineligible for standard refinancing. If this were to happen, special programs that allow a large number of high LTV refinancing would be necessary, such as the Home Affordable Refinance Program (HARP) put forward during the 2008 financial crisis.

For cash-out refinancing, the amount of home equity that could be cashed out is also highly sensitive to the magnitude of house price declines. As it can be seen from Table 4, if house price declines were to raise from 1% to 10%, the share of cash-out refinance loans that would have been forced to lower their cashed-out equity increases from 16.5% to 35.6%, and the corresponding decrease in equity extraction also climbs from \$3,262 to \$33,364. This is because a lower house price reduces home equity which can be cashed out.

The share of homeowners that would be subject to a higher LLPA fee is sensitive to the magnitude of house price declines. This means that as the hypothetical drop in house prices gets larger, more homeowners would incur a higher refinancing cost, which would then lead to lower refinancing thresholds for homeowners.

This sensitivity analysis indicates that even a slight decline in house prices has a large impact on refinancing activities. More specifically, it suggests that it is crucial to stabilize house prices to increase refinancing activity, reduce refinancing cost, and increase equity extraction. It is only in this way that households can benefit from the stimulative effect of lower interest rate through mortgage refinancing during a crisis period.

Covid Economics 68, 10 February 2021: 1-46

CEPR PRESS

Table 4: This table reports the estimates of the impact of house price declines on refinancing eligibility, refinancing costs, and equity extraction. We consider a variety of assumptions on house price declines in the months from April 2020 through October 2020. The naïve house price declines considered are 1%, 2%, 3%, 5%, 8%, 10%, and they are relative to realized house prices in the origination month of each loan. To perform a comparison with naïve house price declines, in the first column we also include a baseline estimation of how much refinancing activities would have been impacted without intervention. The sample consists of 30-year, fixed-rate, GSE-backed refinance loans originated from April 2020 till October 2020. To put the numbers in context, the average monthly house price growth rate in 2020 is around 1%.

	House Price Decline:						
	Model Predicted Decline	1%	2%	3%	5%	8%	10%
Share of Non-cashout Loans Needs Additional Insurance	14.8%	7.32%	10.40%	12.77%	14.65%	27.08%	29.67%
Item Counts	267,956	132,670	188,556	231,666	265,761	491,164	538,062
Average Extra Cost	$3,619	$1,186	$1,934	$2,482	$4,217	$4,927	$6,712
Share of Non-cashout Loans Ineligible for Refinancing	3.34%	0.66%	0.84%	2.57%	4.12%	8.20%	11.09%
Item Counts	60,632	12,029	15,196	46,578	74,729	148,657	201,191
Share of cash-out Loans with Less Equiqty Extraction	22.47%	16.54%	18.48%	20.19%	21.67%	33.51%	35.56%
Item Counts	145,148	106,853	119,404	130,454	139,989	216,476	229,742
Average Decrease in Equity Extraction	$15,371	$3,262	$6,565	$9,879	$16,501	$26,594	$33,364
Share of Loans with Higher LLPA Cost	37.01%	19.29%	28.08%	32.10%	38.59%	53.78%	71.16%
Item Counts	910,384	474,441	690,556	789,549	949,237	1,322,802	1,750,306
Avg Extra LLPA Cost as up-front fee	$1,419	$1,553	$1,485	$1,403	$1,348	$1,365	$1,326
Avg Extra LLPA Cost as Interest Payment	$5,626	$6,201	$5,918	$5,588	$5,368	$5,438	$5,270

Covid Economics 68, 10 February 2021: 1-46

5 Robustness Analysis

In this section, we analyze the robustness of our predictions. First, we consider alternative foreclosure measurements and estimate how much foreclosure would be prevented by forbearance. Second, we analyze alternative econometric models and their produced estimates for the spillover effect of foreclosures on house prices. We examine the predictions of these models on the impact of forbearance on house prices (the first arrow of the graph in Fig. 1).

Our sensitivity analysis of the refinancing activity to hypothetical house price declines (the second arrow in Fig. 1) should be robust as long as GSEs do not alter their eligibility requirements and standards for charging LLPA fees. This is because (1) our calculation is based on GSE eligibility requirements and LLPA matrices, and (2) our sample covers all GSE refinance loans originated from April 2020 to October 2020.

5.1 Alternative Measures of Foreclosures and Delinquencies

Our model uses US Home Foreclosure Filings Total to measure foreclosures in each month. The Mortgage Bankers' Association (MBA) also provides data on foreclosure starts, foreclosure inventory, and delinquencies. However, the frequency of their data is at most quarterly, which would be too coarse for our counterfactual analysis on refinancing activities using loan-level data. For example, if we know that the counterfactual cumulative house price declines in July, August, and September are $1\%, 2\%, 4\%$ respectively, we can correspondingly calculate the counterfactual LTV ratios for all loans originated in each of those months, and analyze the effect on their refinancing cost and eligibility. However, if we only know that the counterfactual average house price decline in Q3 is around 2.3%, it would be hard to accurately calculate the counterfactual LTV ratios in each month. We could overestimate the impact on loans originated in July, and underestimate the impact on loans originated in September.

Even though the MBA data are inappropriate for our analysis on refinancing activity, we can still use them to calculate the counterfactual estimate of foreclosures in the absence of intervention. We use the foreclosure starts rate and serious delinquency rate as our measurement of foreclosure and delinquency. Similar to what we have done in Section 4.1, we calculate the transition rate from serious delinquency to foreclosure start in each quarter from aggregate data, and then estimate

Covid Economics 68, 10 February 2021: 1-46

what the foreclosure starts rate would have been without intervention. We find that the foreclosure starts rate might have been 0.32% and 0.64%, respectively in Q2 and Q3 2020, in the absence of intervention, whereas the actual foreclosure starts rate is 0.03% in Q2 and Q3 2020 (See Figure 7). This supports our conclusion that, in the absence of intervention, there would have been many more foreclosures.

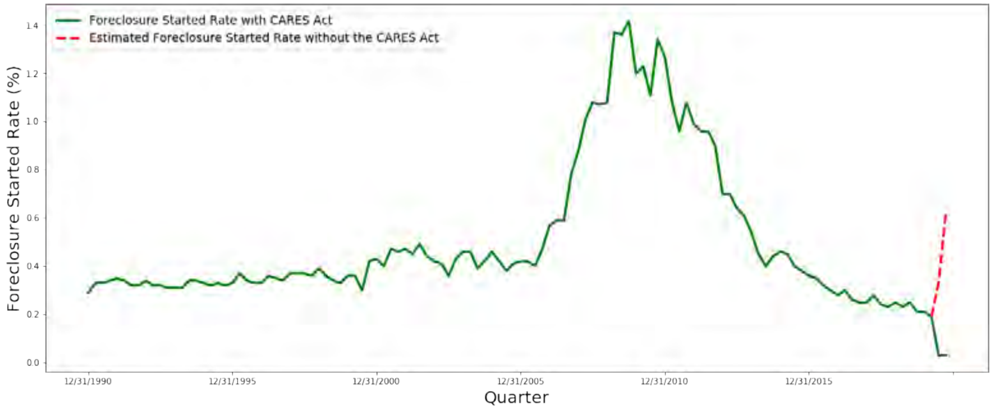

Figure 7: US Foreclosures Started Rate with the foreclosure moratorium (green solid line) and estimated US Foreclosures Started Rate in the hypothetical scenario of no moratorium (red dotted line).

5.2 Alternative Estimates of House Prices

In our study, we build a SVAR model to estimate how foreclosure would have impacted house prices in the absence of forbearance. We then use the resulting counterfactual house price declines and our loan-level data to estimate the impact of the foreclosure moratorium on refinancing activities. There are alternative econometric models developed in the literature, which estimate the spillover effect of foreclosures on house prices (see, for instance, Calomiris, Longhofer, and Miles (2013) and Mian, Sufi, and Trebbi (2015)). As discussed earlier, we did not directly use the model in Calomiris, Longhofer, and Miles (2013), but rather constructed a similar model fed with data at a monthly, instead of quarterly, frequency for our refinancing analysis. As a robustness check, we compare the magnitude of house price declines implied by their model with ours.[18]

[18] The difference in the estimates may result from the different data size, use of state-level and low-frequency data in their model, different measurement of house prices and foreclosures, and different identification methods.

Covid Economics 68, 10 February 2021: 1-46

Our model estimates that spillover effects from a large wave of foreclosures prevented by the moratorium in Q2 and Q3 2020 could have led to an approximately 5% house price decline in Q3 2020 relative to the realized house price in the same quarter. In Section 5.1, we also estimate that without the moratorium, the foreclosure starts rate could have been 0.32% and 0.62% instead of 0.03%, respectively in Q2 and Q3 2020. In Calomiris, Longhofer, and Miles (2013), they use the log foreclosure starts rate from MBA to measure foreclosures, and their impulse response function indicates that a comparable shock (that is, a quarterly foreclosure shock that increases the log foreclosure starts rate from $ln(0.03)$ to $ln(0.32)$ in Q2 2020, and from $ln(0.03)$ to $ln(0.62)$ in Q3 2020) in Q2 2020 results in a quarterly house price decline of roughly 2% in Q3 2020. The contemporaneous foreclosure shock[19] in Q3 2020 could lead to an additional 1% house price decline in Q3 2020. All this suggests that the procedure in Calomiris, Longhofer, and Miles (2013) would yield similar, but slightly smaller estimates of house price drops without the moratorium.

It is difficult to generate a shock comparable to the foreclosures shocks in Mian, Sufi, and Trebbi (2015) with our data and compare our estimates with theirs directly. This is because their model considers state-level, cross-section data between 2007 and 2009 whereas ours uses aggregate, time-series data. Therefore, we can only compare the results indirectly and resort to approximations. Calomiris, Longhofer, and Miles (2013) compare their estimates with that of Mian, Sufi, and Trebbi (2015), and show that theirs are more conservative: relative to Mian, Sufi, and Trebbi (2015), a similar magnitude of foreclosure shock would only result in about a third of the house price decline (2.7% compared to 8%) for a nine-quarter horizon in Calomiris, Longhofer, and Miles (2013). A three-time stronger response of the house price log growth rate to foreclosure shocks relative to Calomiris, Longhofer, and Miles (2013) would then suggest a roughly 9% house price decline in Q3 2020, after a comparable foreclosure shock in Q2 and Q3 2020. Compared to ours, their estimates of house price drops in the absence of a moratorium are higher.

[19] A contemporaneous foreclosure shock only affects house prices in their identification method that orders foreclosures before house prices.

Covid Economics 68, 10 February 2021: 1-46

CEPR PRESS

6 Conclusion

This paper investigates the impact of mortgage forbearance on the real economy through the housing market and the refinancing channel. We show that mortgage forbearance embedded in the CARES Act not only prevents a large wave of foreclosures that might otherwise have occurred, but also supports refinancing activities by stabilizing house prices. Our analysis implies that the foreclosure moratorium has prevented approximately 900,000 foreclosures in the first seven months of its enactment and a house price decline up to 8% in the period from April to October 2020.

We have quantified how forbearance supports household borrowing through the refinancing channel along three dimensions: relaxing eligibility constraints of mortgagors, increasing their equity extraction, and lowering their refinancing cost. We find that during the first seven months, the foreclosure moratorium allowed more than 60,000 mortgagors (about 3.3% of non-cashout refinancing) to become eligible for refinancing, increased home equity cashed out for around 145,000 households (about 22% of cash-out refinancing) by around \$15,000 on average, and lowered the refinancing cost for at least 900,000 households (about 37% of total refinancing lenders) by around \$5,600 per loan in terms of interest payments. Our estimation shows that forbearance can increase aggregate consumption by \$6.4 billion.

Unlike policies which disproportionally benefit the least hard-hit households during crisis times, mortgage forbearance greatly helps households who are subject to stricter credit standards and declining home equity. Households with low credit score benefit from a saving in terms of refinancing cost by roughly \$3,700 as up-front fee, or equivalently, \$15,500 in the form of interest payment.

Our results have important implications for the design of effective policy interventions. In crisis times like the COVID-19 pandemic, where a large number of households are simultaneously hit by income shocks, decisive and comprehensive foreclosure-prevention interventions are necessary from the very beginning. Lessons from the 2008 financial crisis indicate that the occurrence of a large wave of foreclosure has a negative long-term impact on economy, which is extremely difficult to counteract even with the aid of massive stimulus and debt relief (see Piskorski and Seru (2020)). Our analysis demonstrates that, by preventing a foreclosure shock from occurring, early intervention through mortgage forbearance avoids its potential amplification through the refinancing channel and raises aggregate consumption. This highlights the importance of implementing

Covid Economics 68, 10 February 2021: 1-46

foreclosure prevention policies early during a crisis instead of trying to limit the negative consequences afterward, especially given that ex-post relief policies often come with severe frictions (e.g., Piskorski and Seru (2020), Agarwal et al. (2015), Agarwal et al. (2017)). Moreover, our analysis also demonstrates the extent to which forbearance stabilizes house prices and prevents households' leverage from increasing. In this way, income shocks do not spillover to households who are not directly hit by those shocks, and these households benefit from the provision of liquidity and credit through the refinancing channel. Hence, household debt forbearance policies aimed at reducing foreclosures can have a sizeable impact on the pass-through of monetary policy to households. All this suggests that a combined regulatory effort is needed to support the economy during these crisis periods.

It is also important to stress a few limitations of our findings. First, we focus on the effect of the foreclosure moratorium on household borrowing through the house price channel. We ignore other channels, such as quantitative easing which lowers borrowers' mortgage rates, and unemployment benefits which may help borrowers avoid delinquencies. Our objective is to study the efficacy of foreclosure moratorium conditional on the existence of other stimulus, and we show that mortgage forbearance remains crucial to households even if other intervention actions have been taken. Second, we cannot address the long-term welfare implications of mortgage forbearance. Doing so would require us to properly assess whether the policy makers will further extend forbearance, whether the economy will have a fast recovery, and how fast the delinquent mortgagors will cure their mortgages after they exit forbearance. Finally, it is worth observing that our analysis, like other counterfactual analyses, relies on models and assumptions. For example, we use a Markov transition model to estimate the counterfactual level of foreclosures without intervention and a SVAR model similar to Calomiris, Longhofer, and Miles (2013) to assess the response of house prices to counterfactual foreclosure shocks. Even though we intend to give a lower bound on the efficacy of the forbearance program by carefully identifying strategic delinquencies and making conservative parameter estimates, our model specifications may still have drawbacks. Nonetheless, our sensitive analysis, which is derived directly from GSE requirements and thus not subject to model misspecification, still indicates that household borrowing through the refinancing channel relies heavily on policies that stabilize house prices, such as the foreclosure moratorium.

Covid Economics 68, 10 February 2021: 1-46

References

Agarwal, S., B. Ambrose, A. Bandyopadhyay, and Y. Yildirim. 2020a. Communications between borrowers and servicers: Evidence from the covid-19 mortgage forbearance program. Working paper.

Agarwal, S., G. Amromin, I. Ben-David, S. Chomsisengphet, T. Piskorski, and A. Seru. 2017. Policy intervention in debt renegotiation: Evidence from the home affordable modification program. *Journal of Political Economy* 125:654–712.

Agarwal, S., G. Amromin, S. Chomsisengphet, T. Landvoigt, T. Piskorski, A. Seru, and V. Yao. 2015. Mortgage refinancing, consumer spending, and competition: Evidence from the home affordable refinancing program. Working Paper 21512, National Bureau of Economic Research.

Agarwal, S., S. Chomsisengphet, H. Kiefer, L. Kiefer, and P. Medina. 2020b. Inequality during the COVID-19 pandemic: The case of savings from mortgage refinancing. Working paper.

Agarwal, S., J. C. Driscoll, and D. I. Liabson. 2013. Optimal mortgage refinancing: A closed-form solution. *Journal of Money, Credit and Banking* 45:591–622.

Alston, L. J. 1984. Farm foreclosure moratorium legislation: A lesson from the past. *American Economic Review* 74:445–57.

Amromin, G., N. Bhutta, and B. J. Keys. 2020. Refinancing, monetary policy, and the credit cycle. *Annual Review of Financial Economics* 12:67–93.

Anenberg, E., and E. Kung. 2014. Estimates of the size and source of price declines due to nearby foreclosures. *American Economic Review* 104:2527–51.

Arslan, Y., B. Guler, and T. Taskin. 2015. Joint dynamics of house prices and foreclosures. *Journal of Money, Credit and Banking* 47:133–69.

Beraja, M., A. Fuster, E. Hurst, and J. Vavra. 2018. Regional heterogeneity and the refinancing channel of monetary policy*. *The Quarterly Journal of Economics* 134:109–83.

Bhutta, N., and B. J. Keys. 2016. Interest rates and equity extraction during the housing boom. *American Economic Review* 106:1742–74.

CEPR PRESS

Calomiris, C. W., S. D. Longhofer, and W. R. Miles. 2013. The foreclosure–house price nexus: A panel var model for U.S. states, 1981–2009. *Real Estate Economics* 41:709–46.

Campbell, J. Y. 2012. Mortgage market design*. *Review of Finance* 17:1–33.

Campbell, J. Y., S. Giglio, and P. Pathak. 2011. Forced sales and house prices. *American Economic Review* 101:2108–31.

Canner, G., K. Dynan, and W. Passmore. 2002. Mortgage refinancing in 2001 and early 2002. *Federal Reserve Bulletin* 88:469–81.

Chatterjee, S., and B. Eyigungor. 2009. Foreclosures and house price dynamics: A quantitative analysis of the mortgage crisis and the foreclosure prevention policy. *FRB of Philadelphia Working Paper No. 09-22* .

Cherry, S. F., E. X. Jiang, G. Matvos, T. Piskorski, and A. Seru. 2021. Government and private household debt relief during covid-19. Working Paper 28357, National Bureau of Economic Research.

Clauretie, T. M., and T. Herzog. 1990. The effect of state foreclosure laws on loan losses: Evidence from the mortgage insurance industry. *Journal of Money, Credit and Banking* 22:221–33.

Cloyne, J., C. Ferreira, and P. Surico. 2019. Monetary policy when households have debt: New evidence on the transmission mechanism. *The Review of Economic Studies* 87:102–29.

Cloyne, J., K. Huber, E. Ilzetzki, and H. Kleven. 2019. The effect of house prices on household borrowing: A new approach. *American Economic Review* 109:2104–36.

Collins, J. M., and C. Urban. 2018. The effects of a foreclosure moratorium on loan repayment behaviors. *Regional Science and Urban Economics* 68:73 – 83.

Consumer Financial Protection Bureau. 2020. CARES Act mortgage forbearance: What you need to know .

Di Maggio, M., A. Kermani, B. J. Keys, T. Piskorski, R. Ramcharan, A. Seru, and V. Yao. 2017. Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review* 107:3550–88.

Covid Economics 68, 10 February 2021: 1-46

Di Maggio, M., A. Kermani, and C. J. Palmer. 2019. How quantitative easing works: Evidence on the refinancing channel. *The Review of Economic Studies* 87:1498–528.

Diamond, R., A. M. Guren, and R. S. Tan. 2020. The effect of foreclosures on homeowners, tenants, and landlords. *National Bureau of Economic Research* .

Eberly, J., and A. Krishnamurthy. 2014. Efficient credit policies in a housing debt crisis. *Brookings Papers on Economic Activity* 45:73–136.

Emmons, W. R. 2020. How many mortgage foreclosures is forbearance preventing? *St.Louis Fed's Housing Market Perspectives* .

Fuster, A., and P. S. Willen. 2017. Payment size, negative equity, and mortgage default. *American Economic Journal: Economic Policy* 9:167–91.

Gerardi, K., L. Lambie-Hanson, and P. S. Willen. 2013. Do borrower rights improve borrower outcomes? evidence from the foreclosure process. *Journal of Urban Economics* 73:1 – 17.

Guren, A. M., and T. J. McQuade. 2020. How do foreclosures exacerbate housing downturns? *The Review of Economic Studies* 87:1331–64.

Harding, J. P., E. Rosenblatt, and V. W. Yao. 2009. The contagion effect of foreclosed properties. *Journal of Urban Economics* 66:164 – 178.

Hartley, D. 2010. The impact of foreclosures on the housing market. *Economic Commentary* .

Hurst, E., and F. Stafford. 2004. Home is where the equity is: Mortgage refinancing and household consumption. *Journal of Money, Credit and Banking* 36:985–1014.

Iacoviello, M. 2011. Housing wealth and consumption. *FRB International Finance Discussion Paper No. 1027* .

Jarrow, R., D. Lando, and S. M. Turnbull. 1997. A Markov model for the term structure of credit risk spreads. *Review of Financial Studies* 10:481–523.

Keys, B. J., T. Piskorski, A. Seru, and V. Yao. 2014. Mortgage rates, household balance sheets, and the real economy. Working Paper 20561, National Bureau of Economic Research.

Keys, B. J., D. G. Pope, and J. C. Pope. 2016. Failure to refinance. *Journal of Financial Economics* 122:482 – 499.

Mayer, C., E. Morrison, T. Piskorski, and A. Gupta. 2014. Mortgage modification and strategic behavior: Evidence from a legal settlement with countrywide. *American Economic Review* 104:2830–57.

Mayer, C., K. Pence, and S. Sherlund. 2009. The rise in mortgage defaults. *Journal of Economic Perspectives* 23:27–50.

Mian, A., and A. Sufi. 2011. House prices, home equity-based borrowing, and the US household leverage crisis. *American Economic Review* 101:2132–56.

———. 2014. House price gains and U.S. household spending from 2002 to 2006. Working Paper 20152, National Bureau of Economic Research.

Mian, A., A. Sufi, and F. Trebbi. 2015. Foreclosures, house prices, and the real economy. *The Journal of Finance* 70:2587–634.

Mitman, K. 2016. Macroeconomic effects of bankruptcy and foreclosure policies. *American Economic Review* 106:2219–55.

Piskorski, T., and A. Seru. 2018. Mortgage market design: Lessons from the Great Recession. *Brookings Papers on Economic Activity* 49:429–513.

———. 2020. Debt relief and slow recovery: A decade after Lehman. *Journal of Financial Economics* .

Covid Economics 68, 10 February 2021: 1-46

Appendix A Specification of the Markov Transition Model

The transitions of the mortgage delinquency status are modeled via a discrete-time, time-homogeneous Markov chain with a finite state space $S = \{0, D30, D60, D90, F, R, P, L\}$. S represents the set of all possible states, with "0" being performing loans, "D30, D60, D90" being respectively 30-day, 60-day, 90-day delinquencies, "F" being foreclosures, "R" being repurchased, "P" being paid-off, and "L" being loss mitigation. In particular, foreclosures, repurchased, paid-off, and loss mitigation will result in the removal of the loans from the pools, i.e., those states are absorbing states in the Markov chain.

To control for loan-level characteristics, we divide loans in each month into different groups, $g = 1, 2, 3...$, according to their current LTV ratios (30, 60, 80, 100, 120), credit scores (600, 660, 740), Ginnie Mae programs (FHA, VA, PIH, RD), and states (judicial/non-judicial).

For each group g, the discrete-time, finite state, Markov chain is specified by a 8×8 transition matrix $Q(g)$:

$$Q(g) = \begin{pmatrix} p_{(0,0)}(g) & p_{(0,D30)}(g) & p_{(0,D60)}(g) & p_{(0,D90)}(g) & p_{(0,F)}(g) & p_{(0,R)}(g) & p_{(0,P)}(g) & p_{(0,L)}(g) \\ p_{(D30,0)}(g) & p_{(D30,D30)}(g) & p_{(D30,D60)}(g) & p_{(D30,D90)}(g) & p_{(D30,F)}(g) & p_{(D30,R)}(g) & p_{(D30,P)}(g) & p_{(D30,L)}(g) \\ p_{(D60,0)}(g) & p_{(D60,D30)}(g) & p_{(D60,D60)}(g) & p_{(D60,D90)}(g) & p_{(D60,F)}(g) & p_{(D60,R)}(g) & p_{(D60,P)}(g) & p_{(D60,L)}(g) \\ p_{(D90,0)}(g) & p_{(D90,D30)}(g) & p_{(D90,D60)}(g) & p_{(D90,D90)}(g) & p_{(D90,F)}(g) & p_{(D90,R)}(g) & p_{(D90,P)}(g) & p_{(D90,L)}(g) \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

where $p_{(i,j)}(g) \geq 0$ for $i, j \in S$, and $\sum_{j \in S} p_{(i,j)}(g) = 1$. The (i, j) entry, $p_{(i,j)}(g)$, represents the probability of a loan in group g going from state i to state j in one month. Since states F, R, P, L are absorbing states, $p_{(i,j)}(g) = 0$ for $i \neq j, i \in \{F, R, P, L\}$, and the diagonal entries, $p_{(i,i)}(g) = 1$ for $i \in \{F, R, P, L\}$.

For any period $t \in [0, T]$ and any group g, following Jarrow, Lando, and Turnbull (1997), we estimate the transition probability $p_{(i,j)}(T, g)$ using Ginnie Mae loan-level data as follows:

$$p_{(i,j)}(T, g) = \frac{\sum_{t=0}^{T} N_{(i,j)}(t, g)}{\sum_{t=0}^{T} N_i(t, g)},$$

Covid Economics 68, 10 February 2021: 1-46

where $\sum_{t=0}^{T} N_{(i,j)}(t,g)$ is the total number of transitions from state i to state j in the period of $t \in [0,T]$ for group g, and $\sum_{t=0}^{T} N_i(t,g)$ is the total number of loans in state i in the same period of time for group g. In our study, we choose the period $[0,T]$ as the time prior to the implementation of forbearance, since we want to estimate the transition probability in the absence of forbearance.

The expected number of foreclosures for month t, given the number of loans in each state at month $t-1$, is calculated as:

$$\mathbb{E}_{t-1}[N_F(t)] = \sum_g \sum_i p_{(i,F)}(g) N_i(t-1,g)$$

where $N_i(t,g)$ is the total number of loans in state i at $t-1$ for group g, $p_{(i,F)}(g)$ is the probability that a loan in state i transits to state F. In particular, $p_{(i,F)}(g) > 0$ only if $i = D90$ because it is only loans that are more than 90-day delinquent to be foreclosed.

The underlying rationale for the above calculation is: for each month in forbearance and for each group of loans, we first multiply our estimated transition rate from 90-day delinquency to foreclosure with the number of delinquencies identified as non strategic in the previous step. This yields the expected number of foreclosures next month for each group, in the absence of forbearance. After summing up the expected numbers for all groups, for each month we obtain the estimated numbers of Ginnie Mae backed mortgages which would have been foreclosed without forbearance.

Finally, we estimate the number of total US foreclosure filings that would have occurred without intervention as follows. We first calculate the average proportion of Ginnie Mae backed foreclosed loans in the US total foreclosure filings using the pre-COVID sample from March 2019 till March 2020. We then divide the number of Ginnie Mae backed, foreclosed mortgages estimated above by this proportion.

Appendix B Specification of the SVAR Model

In this section, we give the detailed specification of our SVAR model. The vector of the five endogenous variables is given by

$$\boldsymbol{x_t} = (\Delta une_t, \Delta y_t, f_t, \Delta s_t, \Delta hp_t)',$$

where Δune is the log growth in unemployment rate, Δy denotes the log growth rate of GDP, f denotes the log of the number of foreclosures, Δs represents the log growth rate of new houses for sale, and Δhp is the log growth rate of the house price index. We use monthly data from January 2005 to March 2020 to estimate the SVAR model. A detailed description of the data used in our VAR model and the corresponding plots are given in Appendix D. The structural VAR representation is:

$$\boldsymbol{A_0 x_t} = \boldsymbol{\alpha} + \sum_{i=1}^{T} \boldsymbol{A_i x_{t-i}} + \boldsymbol{\epsilon_t},$$

where $\boldsymbol{\epsilon_t}$ denotes the vector of mutually and serially uncorrelated structural innovations.

Similar to Calomiris, Longhofer, and Miles (2013), we assume that $\boldsymbol{A_0^{-1}}$ has a recursive structure. The ordering is indicated by the definition of $\boldsymbol{x_t}$ above. In particular, house prices depend contemporaneously on all other variables. The model also imposes the exclusion restriction that structural shocks specific to the housing market will not immediately affect the real economy, but rather with a delay of at least a month. This restriction is consistent with the fact that information on house prices is not available instantaneously. This is because the data on house prices are released with lags, and there exists large heterogeneity in agent real estate valuations, allowing us to rule out instantaneous feedback. The number of foreclosures depends contemporaneously on unemployment and GDP, since the growth of the real economy and unemployment rate immediately affects how many people will default on their mortgage loans. This ordering emphasizes that (i) real economic activities affect contemporaneously both house prices and the number of foreclosures, (ii) foreclosures can immediately affect house prices, and (iii) changes in house prices affect future foreclosures.

The reduced-form VAR model is estimated using the least-squares method, and the resulting estimates are used to construct the structural VAR representation of the model. These impulse response is based on wild bootstrap with 1000 replications. Figure B.1 lays out the two estimated impulse response functions (IRFs) of our interest.

The impulse responses between foreclosures and house prices are supportive of the existence a self-reinforcing feedback loop between changes in foreclosures and changes in house prices found by Guren and McQuade (2020), Calomiris, Longhofer, and Miles (2013), Arslan, Guler, and Taskin

Covid Economics 68, 10 February 2021: 1-46

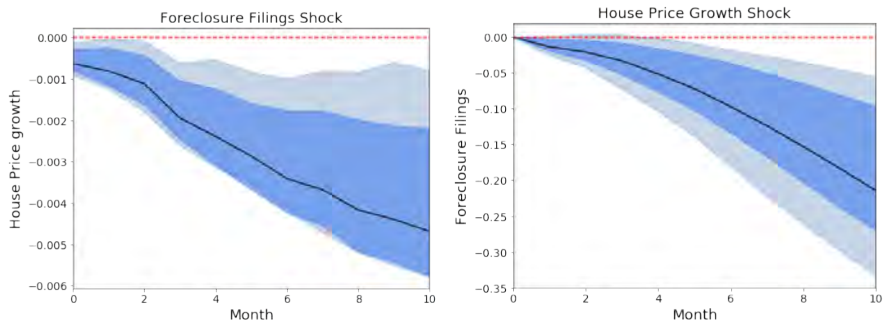

(a) (b)

Figure B.1: Impulse responses to one-standard-deviation structural shock over 10 months, with 68% (dark blue) and 90% (light blue) confidence interval. Panel (a) illustrates the responses of the growth rate of house prices to one-standard structural innovation of foreclosure. Panel (b) illustrates the response of foreclosure to one-standard structural innovation of house prices growth.

(2015), and Chatterjee and Eyigungor (2009). It can be seen from Figure B.1 that an unexpected surge of foreclosures then causes a very persistent and highly statistically significant decrease in house price growth. Conversely, an unexpected increase in house prices growth triggers a highly persistent decrease in foreclosures. The impulse responses can be explained by the fact that a surge in foreclosures tends to flood the housing market with foreclosed real estate that banks are eager to unload promptly, resulting in lower prices. Moreover, when house prices increase, a homeowner can sell a house at a profit even if she is unable to make monthly mortgage payments.

Appendix C Calculation of House Price Declines

In this section, we introduce the methodology used to calculate house price declines in the absence of forbearance.

Let $\boldsymbol{x_t} = (\Delta une_t, \Delta y_t, f_t, \Delta s_t, \Delta hp_t)'$ be the time series of our five endogenous variables. Denote $\boldsymbol{x_t^*} = (\Delta une_t^*, \Delta y_t^*, f_t^*, \Delta s_t^*, \Delta hp_t^*)'$ to be the counterfactual of the same five variables in the hypothetical scenario without foreclosure forbearance. The two time series should be identical until April 2020 when the foreclosure moratorium is introduced. The structural VAR representation is:

$$\boldsymbol{A_0 x_t} = \boldsymbol{\alpha} + \sum_{i=1}^{T} \boldsymbol{A_i x_{t-i}} + \boldsymbol{\epsilon_t},$$

Covid Economics 68, 10 February 2021: 1-46

where $\boldsymbol{\epsilon_t}$ denotes the vector of mutually and serially uncorrelated structural innovations.

Using a moving average representation of the SVAR, we obtain $\boldsymbol{x_t} = \boldsymbol{x_0} + \sum_0^t \Phi_i \boldsymbol{\epsilon_{t-i}}$, and $\boldsymbol{x_t^*} = \boldsymbol{x_0} + \sum_0^t \Phi_i \boldsymbol{\epsilon_{t-i}^*}$. It then follows that

$$\boldsymbol{x_t} - \boldsymbol{x_t^*} = \sum_{i=0}^{t} \Phi_i (\boldsymbol{\epsilon_{t-i}} - \boldsymbol{\epsilon_{t-i}^*}). \tag{1}$$

Assume that the mortgage forbearance only exerts structural shocks on foreclosures, and the structural shocks of other variables will not be affected by mortgage forbearance. We then have

$$\boldsymbol{\epsilon_t} - \boldsymbol{\epsilon_t^*} = [0, 0, \epsilon_{t,f} - \epsilon_{t,f}^*, 0, 0]',$$

where $\epsilon_{t,f} - \epsilon_{t,f}^*$ is the difference of the structural shocks on foreclosures. In particular, we can write the third component of equation (1) as:

$$f_t - f_t^* = \sum_{i=0}^{t} [\Phi_i]_{(3,3)} (\epsilon_{t-i,f} - \epsilon_{t-i,f}^*),$$

where $f_t - f_t^*$ is the difference of foreclosures between the two scenarios, and $[\Phi_i]_{(3,3)}$ is the $(3,3)$ component of the matrix Φ_i. Since we know the actual foreclosure filings f_t and the counterfactual foreclosures estimates f_t^* from Section 4.1, we can calculate the difference $\epsilon_{t,f} - \epsilon_{t,f}^*$ recursively through the above equation. After obtaining $\epsilon_{t,f} - \epsilon_{t,f}^*$, we can use equation (1) to calculate $\Delta hp_t - \Delta hp_t^*$, i.e., the difference of log house price growth in the two scenarios. The ratio between the actual house prices and the counterfactual can then be calculated as

$$\frac{HPI_t}{HPI_t^*} = e^{\sum_{i=0}^{t} \Delta hp_t - \Delta hp_t^*}$$

.

Appendix D Data for The SVAR Model

In this section, we describe and provide the source of data used to estimate our SVAR model. We also plot all the time series in Figure D.2.

Covid Economics 68, 10 February 2021: 1-46

Real GDP and Unemployment Rate: We use monthly data of GDP and unemployment rate. The monthly Real GDP Index is from HIS Markit, and the unemployment rates are from the U.S. Bureau of Labor Statistics. The unemployment rate is defined as the number of unemployed people divided by the size of the labor force.

House Prices: We use the seasonally adjusted House Price Index (HPI) from FHFA. The FHFA HPI is a weighted, repeat-sales index which measures average price changes in repeat sales on the same properties. It covers all single-family properties whose mortgages have been purchased or securitized by Fannie Mae or Freddie Mac since January 1975. This index is published with a two-month lag. At the time our paper has been released, the last published date was October 2020.

Foreclosures: We use the U.S. Home Foreclosure Filings Total from Realty Trac Inc. to measure foreclosures. This index counts properties at different stages of the foreclosure process, such as receiving a notice, auction, and actual foreclosure, in each month from 2005-01 till 2020-10.

New Home for Sale: The data for United States New Home Sales are provided monthly by the US Census Bureau. A new home is considered to be offered for sale when it is being built to be sold. In permit-issuing areas of the United States, this is recorded when the permit to build is issued. In non-permit-issuing areas, this is recorded when work has begun on the footings or foundation, and a sales contract has not been signed nor a deposit accepted. The seasonally adjusted monthly data is jointly released by the U.S. Census Bureau and the U.S. Department of Housing and Urban Development (HUD).

Covid Economics 68, 10 February 2021: 1-46

Covid Economics 68, 10 February 2021: 1-46

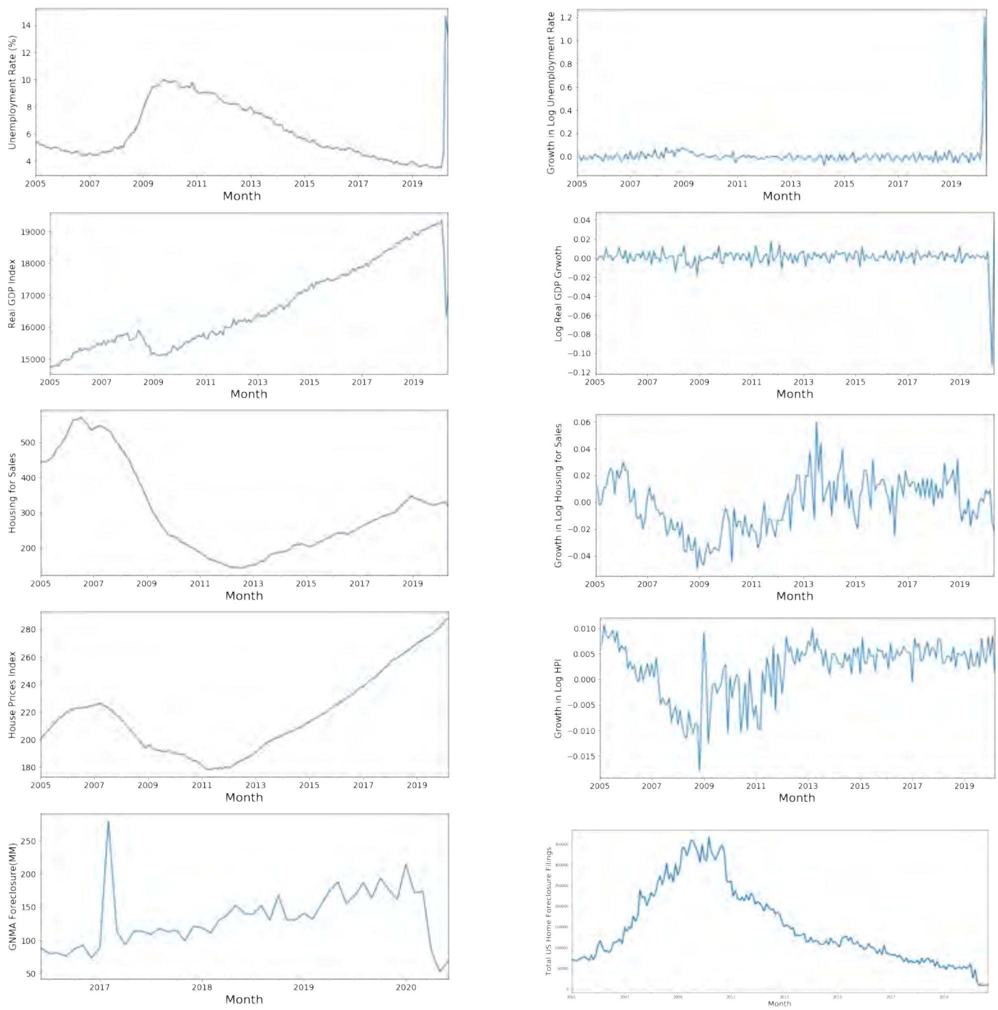

Figure D.2: Data: The figure displays the time series of unemployment rate, real GDP index, new home for sale, House Price Index, foreclosures. It also displays the time series of the log growth rate of unemployment rate, real GDP index, new home for sale, and House Price Index.

Covid Economics Issue 68, 10 February 2021

Social repercussions of pandemics[1]

Philip Barrett[2] and Sophia Chen[3]

Date submitted: 3 February 2021; Date accepted: 7 February 2021

Epidemics may have social scarring effects, increasing the likelihood of social unrest. They may also have mitigating effect, suppressing unrest by dissuading social activities. Using a new monthly panel on social unrest in 130 countries, we find a positive cross-sectional relationship between social unrest and epidemics. But the relationship reverses in the short run, implying that the mitigating effect dominates in the short run. Recent trends in social unrest immediately before and after the COVID-19 outbreak are consistent with this historic evidence. It is reasonable to expect that, as the pandemic fades, unrest may reemerge in locations where it previously existed.

Covid Economics 68, 10 February 2021: 47-68

1 The views expressed in this working paper are those of the authors and do not necessarily represent those of the IMF, its Executive Board, or its management. Working papers describe research in progress by the authors and are published to elicit comments and to encourage debate. We thank Deniz Igan, Gian Maria Milesi-Ferretti, Malhar Nabar, and Antonio Spilimbergo for helpful comments. Luisa Calixto provided excellent research assistance.

2 Economist, International Monetary Fund.
3 Economist, International Monetary Fund.

I. INTRODUCTION

In 1832, the great cholera pandemic hit Paris. In just a few months, the disease killed 20,000 of the city's 650,000 population. Most fatalities occurred in the heart of the city, where many poor workers lived in squalid conditions, drawn to Paris by the Industrial Revolution. The spread of the disease heightened class tensions, as the rich blamed the poor for spreading the disease and the poor thought they were being poisoned. Animosity and anger were soon directed at the unpopular King. The funeral of General Lamarque— pandemic victim and defender of popular causes—spurred large anti-government demonstration on the barricaded streets: scenes immortalized in Victor Hugo's novel *Les Misérables*. Historians have argued that the epidemic's interaction with pre-existing tensions was a principal cause of what came to be known as the Paris Uprising of 1832, which may in turn explain subsequent government repression and public revolt in the French capital in the 19th century (Snowden 2019).

From the Plague of Justinian and the Black Death to the 1918 Influenza Epidemic, history is replete with examples of disease outbreaks casting long shadows of social repercussions: shaping politics, subverting the social order, and some ultimately causing social unrest (North and Paul 1973, Bristow 2017, Elledge 2020). Despite ample examples, quantitative evidence on the link between epidemics and social unrest is scant and limited to specific episodes. This paper fills the gap. We examine whether epidemics may lead to higher likelihood of social unrest using global evidence in recent decades. Understanding the implications of epidemics on social unrest is crucial for preparing for potential social repercussions caused by the COVID-19 pandemic.

There are several ways in which epidemics could affect the likelihood of social unrest. On the one hand, epidemics, like other threats to human health such as natural disasters, can subvert the social order. Mishandling of epidemics may reveal deeper problems such as insufficient social safety nets, incompetent government, or the public's lack of trust in institutions. Outbreaks of contagious diseases have historically caused "fear of the other" and backlash against certain groups (Deverell 2004; Hogarth 2017; Randall 2019). And containment and mitigation efforts could be seen as excessive and unnecessarily costly ex

Covid Economics 68, 10 February 2021: 47-68

post—ironically, this may happen if these efforts are successful in stopping the spread of a disease. Furthermore, possibly severe economic damage from epidemics, especially if affecting disproportionately the poor, could exacerbate inequality and sow the seeds of future social unrest. These are the *scarring* factors of epidemics that may give rise to social unrest.

On the other hand, epidemics are humanitarian crises that bring abrupt disruptions to lives. Such disruptions may impede the communication and transportation needed to organize major protests. Similarly, public opinion might favor cohesion and solidarity in times of duress. In some cases, incumbent regimes may also take advantage of an emergency to consolidate power and suppress dissent. As a result of these *mitigating* factors, the social scarring in the form of unrest may not show up quickly. For these reasons, the overall effect of epidemics on social unrest may depend on the horizon. Its quantitative impact is a priori ambiguous, depending on the offsetting effects of the scaring and mitigating factors at different horizons

Our analysis relies on a new cross-country dataset on social unrest—the Reported Social Unrest Index (RSUI). The RSUI is an index constructed based on press coverage of social unrest. Social unrest events are identified using spikes of the index. Barrett et al. (2020) show that these identified events line up very closely with narrative descriptions of unrest in a variety of case studies, suggesting that the index captures real events rather than shifts in media sentiment or attention. This approach provides a consistent, monthly measure of social unrest for 130 countries from 1985 to the present. It addresses a key challenge for research on social unrest on identifying when events of unrest have occurred, with extant sources of information mostly available at low frequency, with a significant time lag, or have incomplete coverage.

Using this social unrest dataset merged with a comprehensive dataset of epidemics around the world, we find that countries with more frequent and severe epidemics also experienced greater unrest on average. However, this cross-sectional relationship is likely not causal. For example, common factors, such as geography or income level may lead to more unrest and more or more serious epidemics. We thus also estimate a dynamic panel model. The causal

Covid Economics 68, 10 February 2021: 47-68

interpretation of this model rests on the assumption that that the exact timing of epidemics is random. In other words, the likelihood of an epidemic outbreak at a given time does not vary systematically with other factors that may lead to unrest. This seems to be a reasonable assumption within a relatively short timeframe, as in our empirical setting.

The dynamic panel model shows that the likelihood of unrest following epidemics goes down. This suppressive effect is sufficiently large as to drive protest almost to zero in the average country. The COVID-19 experience is not included in our standardized epidemics dataset, and so serves as an out-of-sample test of our quantitative findings. It is so far consistent with this historical pattern, with an almost complete cessation of recorded unrest during the peak months of the crisis in spring 2020, followed by a slow recovery since.
Our paper is related to a large literature on social and political instability. One strand of the literature examines the impact of social and political instability on growth, output, investment (Alesina and Perotti 1996; Alesina et al. 1996; Jong-A-Pin 2009; Aisen and Veiga 2013; Bernal-Verdugo et al. 2013;), and stock market performance (Barrett et al, forthcoming). Miguel et al. (2004) find that economic growth is strongly negatively related to civil conflict: a negative growth shock of five percentage points increases the likelihood of conflict by one-half the following year. A separate strand of the literature examines the determinants of social unrest. Ponticelli and Voth (2020) find a positive correlation between fiscal austerity and social unrest in Europe in the period between World War I and the Global Financial Crisis. Fearon and Laitin (2003) and Collier, Hoeffler, and Rohner (2009) find that poor policies and institutions are important determinants of social unrest among low-income countries (LICs) or other emerging and developing economies (EMDEs).

Empirical evidence on the relationship between epidemics and social unrest are scant. One exception is Cervellati et al (2014), who find that exposure to multi-host vector pathogen (e.g. Malaria, zika, yellow fever) affects the likelihood of civil wars, a particularly extreme form of social unrest. In another study, Cervellati et al (2018) find that a higher exposure to malaria increases the incidence of civil violence in African countries. Both papers focus on specific forms of unrest in Africa and explore cross-section variation in the exposure to a specific type of contagious diseases. In comparison, our paper offers broad-based evidence

Covid Economics 68, 10 February 2021: 47-68

and explores both cross-sectional and time-series variations from the short-run to long-run horizons.

The rest of the paper is organized as follows. Section II describes the data and measures and present evidence on the recent trend of social unrest. Section III discusses the empirical methodology and results. Section IV concludes.

II. Data and measurements

A. Measuring social unrest

We use a newly constructed data on social unrest events based on the Reported Social Unrest Index (RSUI) (Barrett et al. 2020). The authors use articles from major international news sources to create the country specific RSUI. For each country, the RSUI is constructed using the number of articles on social unrest as a fraction of total articles. The authors use text-based criteria the to identify articles on social unrest. For example, relevant articles must include words such as "protest", "riot" or "revolution". They must also exclude certain terms to avoid counting reports about previous events or revolution-themed movies. The selected articles must be at least 100 words long and must mention the name of the country in question. At the country level, the RSUI exhibits very large spikes that are associated with major episodes of social unrest. The authors develop quantitative criteria to formalize these spikes and identify a list of social unrest events at the country-month frequency. [4] They compare these events against consensus narratives for a number of case studies, showing that they align closely, and conclude that this method captures actual major social unrest events. The final dataset consists of 569 events in 130 countries from mid-1980s to early 2020.

We merge the social unrest data with EM-DAT—a comprehensive database of international epidemics and natural disasters, with information on the timing and location of more than 11,000 events since 1990. Although our focus is on epidemics, we also collect data on five

[4] See Barrett et al. (2020) for multiple tests of this event coding against alternative narrative sources. The event dataset has several advantages over the index alone. It addresses the limitation in the RSUI construction that media coverage may not be fully comparable across countries. It also permits better cross-checks for false positives and permits a labelling of each event, thus facilitating event narratives.

Covid Economics 68, 10 February 2021: 47-68

other types of disasters. Four types of disasters—floods, storms, earthquakes, and landslides—are, along with epidemics, the most common events in the sample.(online appendix Table A1). Droughts, although the only the seventh most common type of event appear have particularly extensive impacts, on average affecting ten times as many people as the next most far-reaching disaster, storms. The comparison to other disaster types is useful because, as discussed earlier, epidemics and natural disasters are both humanitarian crises that present similar challenges to social orders.

B. Recent social unrest events

Social unrest pre-COVID-19

Between January 2019 and January 2020, the events dataset identifies 59 unrest events in forty countries. A number of major protests occurred late 2019 to early 2020, most notably in the Middle East and South America but also elsewhere. None of these events appear to be directly linked to major natural disasters or epidemics.[5] Instead, most of the unrest events were motivated by political factors. This recent wave of social unrest events was the continuation of a longer trend since 2016 (Figure 1), which itself reversed a gradual decline in unrest following a peak after the Arab Spring of 2011.

Social unrest during COVID-19

Since the start of the COVID-19 outbreak, the number of major unrest events worldwide has fallen sharply and in March reached its lowest level in almost five years (Figure 1). The decline in social unrest corresponds closely with a generalized decline in mobility driven by regulations such as shelter-in-place orders, and voluntary social distancing. To illustrate the remarkably close association between the timing of the decline in protest and the abrupt cessation of social activities, Figure 1 also includes global averages of two series from

[5] It may be possible that disasters (or handling thereof) are the true primitive cause of social unrest causing political disagreements to which the events are erroneously attributed. Anecdotal evidence suggests that this is not the case for 2019.

Covid Economics 68, 10 February 2021: 47-68

Covid Economics 68, 10 February 2021: 47-68

Google Community Mobility Reports. The mobility series use cell phone location history to measure activity in specific categories, of which we include two (activities in retail and recreation spaces and transit stations), although others look very similar. We interpret this time series correlation as suggestive evidence that mitigating effects of the latest epidemic have likely outweighed any scarring that might have incentivized unrest. We take up this question again more formally in the next section .

Figure 1: Global unrest and mobility since 2016

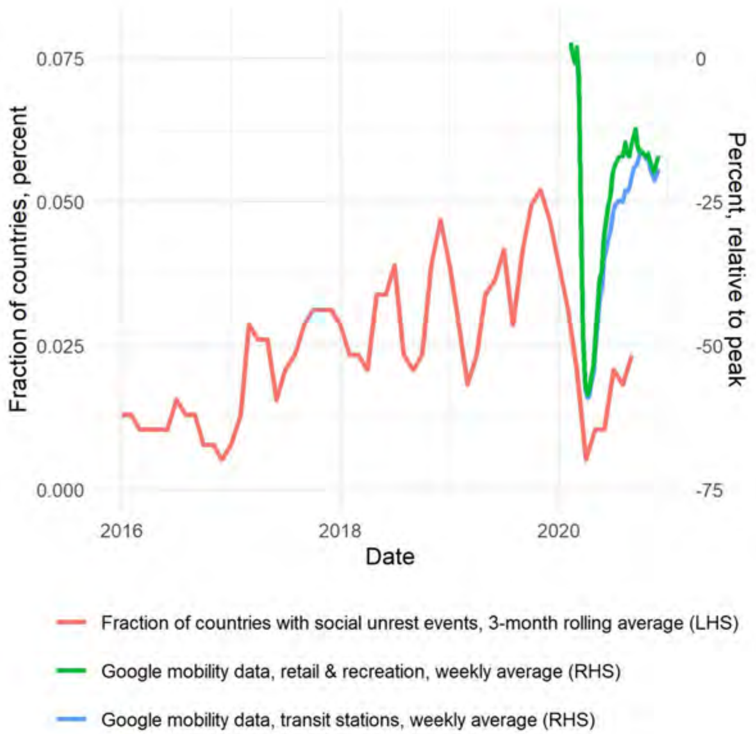

Source: Barrett et al. (2020), Google Community Mobility Reports and authors' calculations.

Notable exceptions to this pause in social unrest include the United States and Lebanon. Yet even in these cases, the largest protests were related to issues which long preceded the COVID-19 epidemic: racial injustice in the United States and governance in Lebanon.

Covid Economics 68, 10 February 2021: 47-68

Figure 2: Daily media mentions of unrest in the United States, early 2020

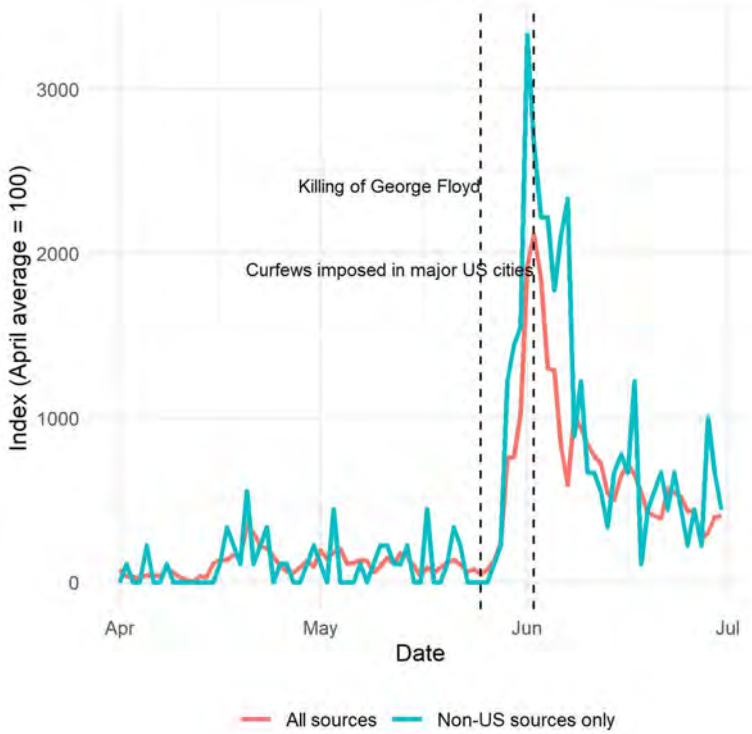

The protests in the United States are a good test of the media-based approach to measuring unrest. An important challenge to our method is that media coverage of other high-profile issues, such as an ongoing pandemic, may "crowd out" the coverage of unrest. In Figure 2 we replicate the search criteria of Barrett et al. (2020) at a daily frequency for the United States.[6] It shows that press articles related to unrest increase sharply at almost exactly the same time as major street protests broke out across the United States, despite a severe and continuing pandemic, suggesting that media coverage remains a good indicator of unrest even during a pandemic.

[6] Barrett et al. (forthcoming) similarly construct a daily social unrest index for a large sample of countries.

CEPR PRESS

Covid Economics 68, 10 February 2021: 47-68

III. EMPIRICAL METHODOLOGY AND RESULTS

Here we analyze more formally the relationship between disasters and unrest. We report two main findings. First, there is a strong cross-sectional relationship in the data: countries with more disasters also have more unrest, even after accounting for variations in region, income, and exposure to waves of unrest. Second, there is a negative within-country relationship between epidemics and unrest, largest at around 4-6 months after the epidemic starts, consistent with the experience under COVID-19, but that other disasters show no obvious within-country intertemporal relationship with unrest.

We interpret this evidence as supporting two conclusions. The cross-country results are consistent with the notion that the permanent and pervasive risk of disasters could lead to more unrest on average. However, the within-country findings suggest that the realization of that risk—that is, the occurrence of disasters at a particular time—does not cause social unrest. If anything, epidemics are more likely to mitigate unrest than to spur it.

A. Cross-sectional evidence

We start by analyzing the cross-sectional relationship between disasters and unrest, asking whether countries have more epidemics have on average more social unrest. This relationship usefully captures long run and persistent variations across countries. This long-run perspective is important because epidemics may leave decades-long shadows in the society, as we alluded to in our earlier account of history of epidemics. However, the limitation of a cross-sectional analysis is its identification: there could be many factors causing a correlation between unrest and disasters.

Figure 3 plots the cross-sectional relationship between per capita disasters and unrest events for the six disaster types in our sample.[7] In almost all cases, there is a positive relationship between unrest and disasters that is stable across income groups. Interestingly, the magnitude

[7] The per capita measure accounts for the fact that large countries (typically) have more area in which natural disasters can occur, and more people who might want to engage in unrest. That said, our findings are almost identical without the per capita scaling.

CEPR PRESS

Covid Economics 68, 10 February 2021: 47-68

of the relationships is also relatively stable across the different types of disasters, despite considerable variation in the countries exposed to different disasters.[8]

Figure 3: Cross-sectional relationship between disasters and unrest

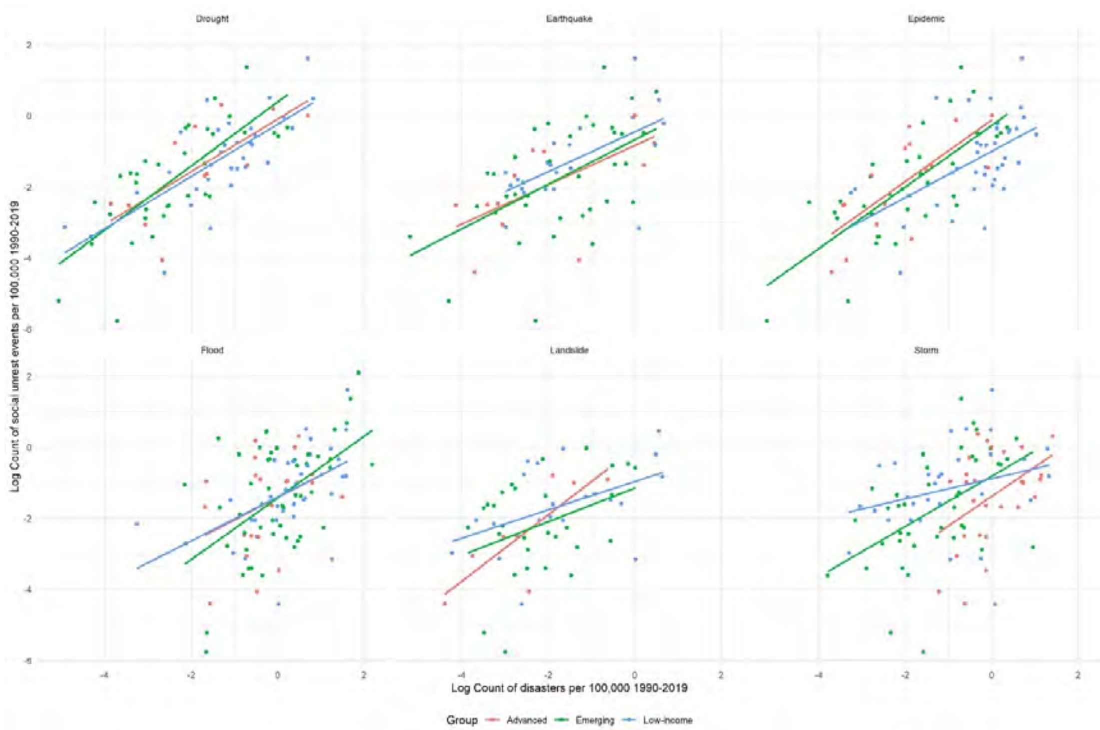

Source: Barrett et al. (2020), EM-DAT, and authors' calculations.

The graphical evidence seems convincing. In the Appendix we check that this relationship is robust, reporting the results of the following regression:

$$y_i = \alpha + \textstyle\sum_d \beta x_{i,d} + \gamma_{r(i)} + \nu_{y(i)} + e_i \qquad (1)$$

where i and r index country and region respectively, y_i is the log number of social unrest events per capita since 1990 in country i, $x_{i,d}$ is the log number of disasters, e_{ir} is the error

[8] Of course, some countries are exposed to multiple disasters, but many are not. For example, earthquakes are relatively common in Japan, but droughts are rare there. The reverse is true in Niger.

term, and $\gamma_{r(i)}$ and $\nu_{y(i)}$ are fixed effects for the region, $r(i)$, and income group, y(i), for country i. With these controls, the coefficients of interest β_d capture the relationship between disasters and unrest across countries within the same region and income group.
The results are reported in Table A2 and confirm the relationships in Figure 3. We find a positive and statistically significant cross-country relationship between disasters (all types combined) and social unrest (column (1)). The relationships for other types of disasters are statistically indistinguishable from each other (columns (2)-(4)) and are robust across regions and country income group.

Social unrest often occurs in regional waves, such as Arab Spring in 2011, or Latin American protests in 2019. Indeed, Barrett et al. (2020) show that recent social unrest both domestically and in neighboring countries is correlated with higher future social unrest. Likewise, disasters—particularly epidemics—are frequently regionally clustered. It thus seems plausible that the cross-sectional relationship is merely the result of coordinated waves of disasters and protest. To investigate this issue, we also estimate equation (1) using five-year averages within each country, including time fixed effects. The inclusion of time fixed effects accounts for global factors that may give rises to waves of unrest and disasters.

These results, reported in Table A2, columns (5) to (8), similarly show a robust and positive relationship between disasters and unrest. In fact, in many cases, this relationship is slightly stronger. The comparison now is across countries within a given five-year period, and so these results say that even in the medium run, countries with more disasters also have more unrest.

Overall, the cross-country evidence is relatively clear: there is a positive and strong relationship between natural disasters of all types and unrest. However, this does not necessarily mean that unrest is higher *because* of disasters. It could be that other factors, such as geography or institutions, induce a positive correlation between the two. To investigate the consequences of a disaster, we turn to the dynamic within-country behavior of unrest following a disaster next.

Covid Economics 68, 10 February 2021: 47-68

B. Evidence from a dynamic panel

Turn to the dynamic relationship between social unrest and disasters, we estimate the following panel regression, run separately for each type of disasters:

$$y_{it} = \alpha_i + \eta_t + \sum_{j=1}^{n} \beta_j x_{it}^j + \gamma' z_{it} + e_{it}, \qquad (2)$$

where y_{it} is an indicator for a social unrest event in country i in year t, x_{it}^j is an indicator variable that takes a value of 1 if the latest disaster event occurred j months prior, z_{it} is a vector of controls, and α_i and η_t are country and time fixed effects respectively. We use a linear probability so that we can include admit a wide battery of fixed effects—something much harder in a nonlinear framework.

Note that this is slightly different from a local projection framework (Jordà 2005). There, one estimates separate regressions of $y_{i,t+h}$ on x_{it} and the set of controls over horizons $h = 0, 1, \dots, H$. The local projection approach is somewhat problematic in the current setting, simply because both unrest and disasters are rare, with unrest and epidemics occurring in around 1.2 and 1.9 percent of country-months respectively. Therefore, despite having a large sample, at any given horizon h the probability of encountering many observations with both a disaster at time t and an epidemic at time h is very low. As a result, the local projection estimates have very low power, producing point estimates with very large standard errors.

A natural solution to this problem is to try to combine horizons to increase power. In the standard local projection framework this is impossible, as regressions at horizons h and h' are complete separate objects, so there is no common variance-covariance matrix that can be used to compute standard errors. In the specification in equation (2), we have no such problem. We can easily estimate the coefficient on a horizon of, say, 4-6 months simply by replacing x_{it}^j with an indicator equals to 1 if the last disaster occurred 4, 5, or 6 months prior. This approach allows us to increase the power of our estimates while preserving the monthly frequency of the data.

Alternative ways of using the local projection framework exist, but they often require changing the interpretation of the coefficients. For example, using a rolling average for

Covid Economics 68, 10 February 2021: 47-68

unrest would capture the probability of unrest up to horizon h but inferring the probability of unrest in any given interval would be challenging. In comparison, the panel specification in equation (2) is simple and more transparent.

The dynamic panel regression approach is not without drawbacks. In particular, our point estimates are relative to some particular horizon after the disaster rather than the moment before it. We use more than 60 months after a disaster as the comparison horizon. This means that the coefficients measure the conditional average unrest likelihood in the short run (less than 5 years) relative to the period of 5 years or longer after a disaster occurred in the country. If there are permanent effects of disasters on unrest, this would bias our estimates downwards. However, Appendix Figure A1 suggests that such potential bias is quantitatively insignificant. There, we compare our results to the local projection equivalents. There is no evidence of a systemic difference on the point estimates between the two sets of results. Yet, the local projection estimates are much noisier than the panel estimates. Overall, this comparison supports the argument that the dynamic panel approach improves power without introducing noticeable biases.

We report the results for epidemics in Table 1. Column (1) is simply a baseline, reporting the average frequency of unrest events in the sample. Column (2) reports results of a simple dynamic specification. It shows a clear decline in the probability of unrest following an epidemic, with the effect peaks at 4-6 months. At its peak, the decline is sufficiently large as to almost entirely offset the average probability of unrest (a fall of 1.2 percent versus an average of 1.4 percent). This specification and all subsequent ones include country fixed effects and so abstract from the cross-sectional relationship discussed in the previous section, yielding a true dynamic response.

Other specifications add extra controls, including time since the last social unrest events in the same and neighboring countries, the intensity of the last epidemic, and time fixed effect. Our preferred specification is (3), which strikes a balance between fitting the data and preserving the sample. It also performs well on measures of fit–measured by Akaike Information Criterion and R^2–of those without time fixed effects. Consistent with (2), it

Covid Economics 68, 10 February 2021: 47-68

CEPR PRESS

Covid Economics 68, 10 February 2021: 47-68

shows probability of unrest is lower during epidemic, with the effect peaks at 4-6 months. With the additional control, the effect on longer-horizons becomes smaller and statistically insignificant.[9] Point estimates from the richer specifications (4-6) are broadly stable although the reduction in sample size (as some backward-looking variables are not available for all periods) leads to lower power. The inclusion of time fixed effects in column (6) likely absorb useful variation in epidemics, which can be correlated across countries even if exogenous.

We perform the same exercise for other types of disasters and find much weaker evidence of a dynamic pattern. As an illustration, Table 2 presents the results for floods, the largest category of natural disasters. There is no statistically significant pattern, in contrast to epidemics. Results on other natural disasters are similar and omitted in the interest of space. What could explain this difference between natural disasters and epidemics? It could be that there is something unique about the contagious nature of epidemics and its suppressive force on protest and other forms of unrest, which are inherently social activities. Although natural disasters of all kinds likely impede protest movements by hampering transport, communications and the like, these challenges are likely less severe, more short-lived, and limited to the activity itself. In an epidemic, however, attending a large scale and possibly contagious event likely have serious and long-lasting consequence for the health of the participants (and those they meet),which may be a much higher deterrent to unrest.

Finally, a word on identification. While we do not advocate strongly for a causal interpretation, natural disasters are likely very close to randomly assigned (especially earthquakes and storms). For epidemics, its occurrence may be related to country-specific conditions. For example, country with worse health infrastructure may be less likely to prevent an epidemic. However, the specific timing of epidemics is likely random, particularly at the monthly level. This consideration on the timing of disasters informs our empirical design, which combines outcome horizons to increase power without losing the monthly

[9] Although the long-run effect is small, is natural to expect that epidemics may indeed have some long-run effects, perhaps either because higher death rates can affect demographics for years to come, or because they are sufficiently traumatic that they have long-lasting psychological effects on the population.

frequency. Under this assumption, the effect of estimated in a tight window reflects a causal link.

What should we make of the empirical evidence overall? As discussed earlier, epidemics may have both scarring and mitigating effects on social unrest. Whether an epidemic will increase or decrease the overall likelihood of social unrest may depend on whether the scarring effect or the mitigating effect dominates over time. Our empirical results confirm this ambiguous relationship. On the one hand, the finding of a robust and positive cross-sectional relationship between epidemics and social unrest is consistent with a long-run scarring effect. On the other hand, evidence is weak in the medium run. In the short run, the likelihood of social unrest is slightly lower following epidemics, consistent with the notion that the mitigating factors dominates in the short run.

IV. Conclusion

This paper examines the implications of epidemics on social unrest using global evidence in recent decades. Drawing on a new cross-country dataset of social unrest, we find a positive cross-sectional relationship between social unrest and epidemics. While this result may reflect a positive long-run effect, we find that the relationship reverses in the short run. This difference between the long-run and short-run result suggest are consistent with the theoretical prediction on the scarring and mitigating effects of epidemics.

Recent trends in social unrest immediately before and after the COVID-19 outbreak are also consistent with this historic evidence. Unrest was elevated before the COVID-19 crisis began but has declined as the pandemic has continued. If history is a guide, it is reasonable to expect that, as the pandemic fades, unrest may reemerge in locations where it previously existed, not because of the COVID-19 crisis per se, but simply because underlying social and political issues have not been tackled.

Covid Economics 68, 10 February 2021: 47-68

Table 1: Impact of epidemics on unrest, linear probability model

	Dependent variable: Indictor for social unrest event					
	(1)	(2)	(3)	(4)	(5)	(6)
Epidemic, current month		-0.011***	-0.009*	-0.008	-0.011*	-0.006
		(0.003)	(0.005)	(0.006)	(0.006)	(0.006)
Epidemic, last 1-3 months		-0.009***	-0.007	-0.008	-0.009	-0.006
		(0.003)	(0.005)	(0.006)	(0.006)	(0.007)
Epidemic, last 4-6 months		-0.012***	-0.011**	-0.010*	-0.010*	-0.006
		(0.003)	(0.005)	(0.006)	(0.006)	(0.006)
Epidemic, last 7-12 months		-0.006**	-0.001	-0.001	-0.001	0.003
		(0.003)	(0.004)	(0.005)	(0.005)	(0.005)
Epidemic, last 13-24 months		-0.010***	-0.008**	-0.008*	-0.010**	-0.007
		(0.003)	(0.004)	(0.004)	(0.005)	(0.005)
Epidemic, last 25-60 months		-0.006***	-0.002	-0.002	-0.003	-0.001
		(0.002)	(0.003)	(0.003)	(0.003)	(0.003)
Months since last social unrest event			-0.0004***	-0.0004***	-0.0004***	-0.0004***
			(0.00003)	(0.00004)	(0.00004)	(0.00004)
Months since last social unrest event, neighboring country				0.00002	0.00004	0.00003
				(0.00005)	(0.00004)	(0.00004)
Deaths per capita in last epidemic					-0.0001***	-0.0001**
					(0.00004)	(0.0001)
Constant	0.014***					
	(0.001)					
Country FEs	No	Yes	Yes	Yes	Yes	Yes
Time FEs	No	No	No	No	No	Yes
Normalized AIC	-1.437	-1.442	-1.056	-1.006	-1.127	-1.155
R^2	0.014	0.019	0.05	0.051	0.049	0.075
Observations	27,505	27,505	18,123	15,137	12,953	12,953

Note: This table reports results of dynamic panel regressions. The dependent variable is a dummy variable that takes a value of 1 for social unrest events. Standard errors clustered at the country-month level are shown in parenthesis.

Covid Economics 68, 10 February 2021: 47-68

Table 2: Impact of flood on unrest, linear probability model

	Dependent variable: Social unrest event					
	(1)	(2)	(3)	(4)	(5)	(6)
Flood, current month		-0.007*	-0.003	-0.004	0.001	0.003
		(0.004)	(0.007)	(0.008)	(0.009)	(0.009)
Flood, last 1-3 months		-0.007*	-0.002	-0.006	-0.0004	-0.0001
		(0.004)	(0.006)	(0.007)	(0.007)	(0.008)
Flood, last 4-6 months		-0.005	0.001	-0.003	0.002	0.002
		(0.003)	(0.006)	(0.007)	(0.007)	(0.008)
Flood, last 7-12 months		-0.004	0.001	-0.002	0.004	0.004
		(0.003)	(0.006)	(0.007)	(0.008)	(0.008)
Flood, last 13-24 months		-0.003	-0.001	-0.004	0.003	0.004
		(0.003)	(0.006)	(0.007)	(0.007)	(0.007)
Flood, last 25-60 months		-0.006**	-0.006	-0.009	-0.005	-0.003
		(0.003)	(0.006)	(0.006)	(0.007)	(0.007)
Months since last social unrest event			-0.0004***	-0.0004***	-0.0004***	-0.0004***
			(0.00004)	(0.00004)	(0.00004)	(0.00004)
Months since last social unrest event, neighbor				0.00001	0.00003	0.00001
				(0.00004)	(0.00004)	(0.00003)
Deaths per capita in last epidemic					-0.0001**	-0.00005
					(0.00003)	(0.0001)
Constant	0.013***					
	(0.001)					
Country FEs	No	Yes	Yes	Yes	Yes	Yes
Time FEs	No	No	No	No	No	Yes
Normalized AIC	-1.487	-1.49	-1.077	-1.013	-1.125	-1.15
R^2	0.013	0.017	0.047	0.051	0.048	0.072
Observations	36,143	36,143	23,175	19,033	16,242	16,242

Note: This table reports results of dynamic panel regressions. The dependent variable is a dummy variable that takes a value of 1 for social unrest events. Standard errors clustered at the country-month level are shown in parenthesis.

Covid Economics 68, 10 February 2021: 47-68

References

Aisen, Air and Francisco Veiga, 2013, "How Does Political Instability Affect Economic Growth?" *European Journal of Political Economy*, 29:151-167.

Alesina, Alberto and Perotti, Roberto, 1996, "Income Distribution, Political Instability, and Investment". *European Economic Review*, 40(6):1203-1228.

Alesina, Alberto, Sule Özler, Nouriel Roubini, and Philip Swagel, 1996, "Political Instability and Economic Growth", *Journal of Economic Growth*, 1:189-211.

Barrett, Philip, Maximiliano Appendino, Kate Nguyen, and Jorge de Leon Miranda, 2020, Measuring Social Unrest Using Media Reports, IMF Working Paper 20/129.

Barrett, Philip, Mariia Bondar, Sophia Chen, Mali Chivakul, Deniz Igan, forthcoming, Social Unrest and Financial Markets, IMF Working Paper.

Bernal-Verdugo, Lorenzo E., Davide Furceri, and Dominique Guillaume, 2013, The Dynamic E_ect of Social and Political Instability on Output: The Role of Reforms. IMF Working Paper No. 13/91.

Bristow, Nancy, 2017, American *Pandemic: The Lost Worlds of the 1918 Influenza Epidemic*, Oxford University Press.

Cervellati, Matteo, Uwe Sunde, and Simona Valmori, 2017, "Pathogens, Weather Shocks and Civil Conflicts", *The Economic Journal*, 127(607), Pages 2581–2616.

Cervellati, Matteo, Elena Esposito, Uwe Sunde, and Simona Valmori, 2018, "Long-term Exposure to Malaria and Violence in Africa", *Economic Policy*, 33(95), Pages 403–446.

Collier, P., A. Hoeffler and D. Rohner, 2009, "Beyond Greed and Grievance: Feasibility and Civil War", *Oxford Economic Papers*, No. 61, Pages 1–27.

Deverell, William F., 2004, *Whitewashed Adobe the Rise of Los Angeles and the Remaking of Its Mexican Past.* University of California Press.Elledge, Jonn, 2020, "Revolts and Revolutions: How History Reveals the Ways Coronavirus Could Change Our World Forever", *Prospect*, April 17, 2020.

Fearon, J.D. and D.D. Laitin, 2003, "Ethnicity, Insurgency, and Civil War", *American Political Science Review*, No. 97, Pages 75–90.

Gonzales-Torres, Ada and Elena Esposito, 2017, Epidemics and Conflict: Evidence from the Ebola outbreak in Western Africa, mimeo.

Hogarth, Rana A., 2017, *Medicalizing Blackness Making Racial Difference in the Atlantic World, 1780-1840*, The University of North Carolina Press.International Monetary Fund (IMF), 2020a, World Economic Outlook, Chapter 1: Global Prospects and Policies Global Prospects and Policies, Box 1.4.

International Monetary Fund (IMF), 2020b, World Economic Outlook, Chapter 2: Dissecting the Economic Effects.

Jong-A-Pin, Richard, 2009, "On the Measurement of Political Instability and Its Impact on Economic Growth", *European Journal of Political Economy*, 25:15-29.

Jordà, Òscar, 2005, "Estimation and Inference of Impulse Responses by Local Projections",

American Economic Review, 95(1):161-182.

Miguel, Edward, Shanke Satyanath, and Ernest Sergenti, 2004, "Economic Shocks and Civil Conflict: An Instrumental Variables Approach", *Journal of Political Economy*, 112(4), Pages 725-753.

North, Douglass C. and Robert Paul, 1973, *The Rise of the Western World: A New Economic History*, Cambridge University Press.

Ponticelli, Jacopo and Hans-Joachim Voth, 2020, "Austerity and Anarchy: Budget Cuts and Social Unrest in Europe, 1919-2008", *Journal of Comparative Economics* 48 (1), Pages 1-19.

Randall, David K., 2010, *Black Death at the Golden Gate: The Race to Save America from the Bubonic Plague*, W. W. Norton & Company.

Snowden, Frank M., 2019, *Epidemics and Society: From the Black Death to the Present*, Yale University Press.

Covid Economics 68, 10 February 2021: 47-68

Appendix

Table A1: EM-DAT disasters since 1990 (with at least 50 observations)

Type	Number	Avg Deaths	Avg affected	Avg Damage (USD)	Avg Mortality (%)
Flood	4096	48	751467	187812	0.01
Storm	2942	139	321082	483299	0.04
Epidemic	1235	163	18993	0	0.86
Earthquake	819	1007	173415	903272	0.58
Landslide	523	50	13025	15187	0.39
Extreme temperature	522	335	198391	105692	0.17
Drought	472	51	3671603	322349	0.00
Wildfire	340	6	19438	366170	0.04
Volcanic activity	155	16	48054	14949	0.03

Covid Economics 68, 10 February 2021: 47-68

Table A2: Cross-section regressions

	Dependent variable: Log number of social unrest events 1990-2019							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log number of disasters, per capita	0.515***				0.793***			
	(0.034)				(0.028)			
Log number of droughts, per capita		0.434***	0.411***	0.418***		0.626***	0.626***	0.624***
		(0.063)	(0.057)	(0.058)		(0.037)	(0.037)	(0.037)
Log number of earthquakes, per capita		0.466***	0.454***	0.457***		0.745***	0.745***	0.748***
		(0.063)	(0.058)	(0.058)		(0.041)	(0.041)	(0.041)
Log number of epidemics, per capita		0.590***	0.573***	0.573***		0.714***	0.714***	0.715***
		(0.067)	(0.061)	(0.062)		(0.039)	(0.039)	(0.038)
Log number of floods, per capita		0.723***	0.758***	0.743***		0.968***	0.968***	0.967***
		(0.147)	(0.137)	(0.136)		(0.048)	(0.048)	(0.048)
Log number of landslides, per capita		0.499***	0.480***	0.481***		0.771***	0.771***	0.773***
		(0.071)	(0.065)	(0.065)		(0.040)	(0.040)	(0.040)
Log number of storms, per capita		0.490***	0.505***	0.526***		0.784***	0.784***	0.788***
		(0.083)	(0.077)	(0.079)		(0.039)	(0.039)	(0.039)
Constant	-0.974***	-1.003***	-0.983***	-1.291***	-0.942***	-0.962***	-0.962***	-0.935***
	(0.066)	(0.070)	(0.087)	(0.202)	(0.067)	(0.084)	(0.084)	(0.136)
Regression type	x-sect	x-sect	x-sect	x-sect	5-yr panel	5-yr panel	5-yr panel	5-yr panel
Region FEs	No	No	Yes	Yes	No	Yes	Yes	Yes
Time FEs	No	No	No	No	No	No	Yes	Yes
Income group FEs	No	No	No	Yes	No	No	No	Yes
Observations	504	504	504	504	795	795	795	795
R^2	0.281	0.291	0.363	0.368	0.503	0.585	0.585	0.592
Adjusted R^2	0.280	0.282	0.350	0.354	0.503	0.580	0.580	0.584

Note: This table reports results of cross-sectional (columns 1-4) and 5-year panel (columns 5-8) regressions. The dependent variable is the log number of social unrest events over the 1990-2019 period. Robust standard errors shown in parenthesis.

Figure A1: Dynamic regressions versus local projections

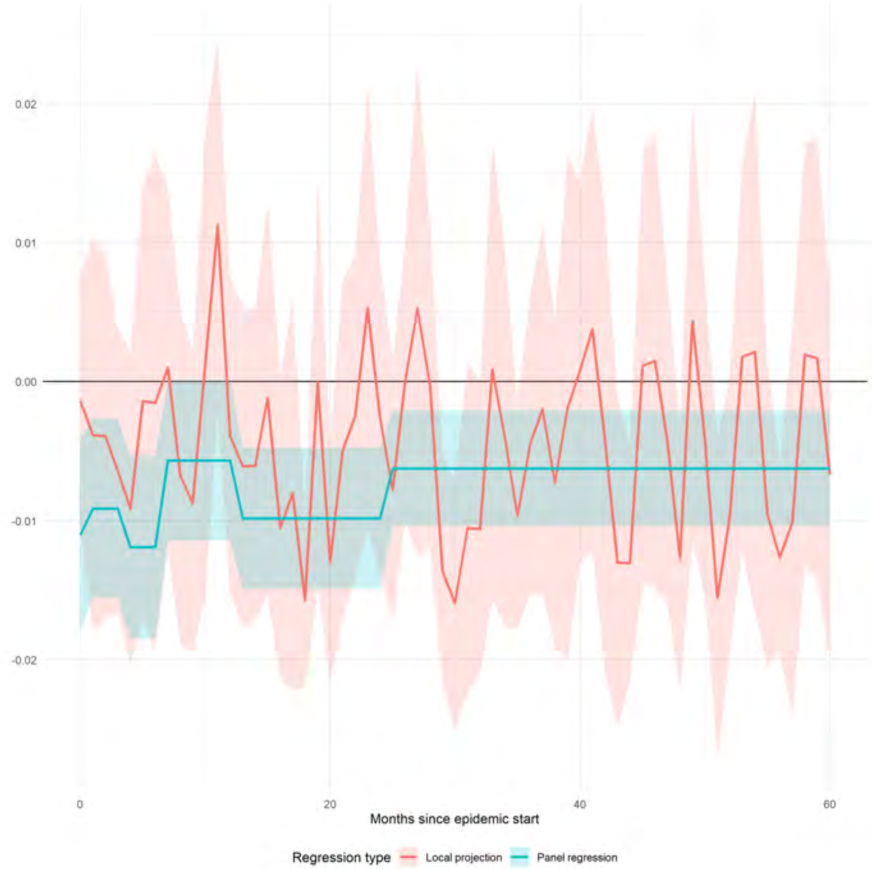

Note: This figure plots the point estimates and 90 percent confidence intervals of local projection and panel regression models. Source: Barrett et al. (2020), EM-DAT, and authors' calculations.

Covid Economics 68, 10 February 2021: 47-68

Covid Economics Issue 68, 10 February 2021

Learning at home: Distance learning solutions and child development during the COVID-19 lockdown[1]

Hugues Champeaux,[2] Lucia Mangiavacchi,[3]
Francesca Marchetta[4] and Luca Piccoli[5]

Date submitted: 22 January 2021; Date accepted: 27 January 2021

Covid Economics 68, 10 February 2021: 69-99

School closures, forcibly brought about by the COVID-19 crisis in many countries, have impacted children's lives and their learning processes. There will likely be substantial and persistent disparities between families in terms of educational outcomes. Distance learning solutions adopted by schools have been heterogeneous across countries, within countries and between school levels. As a consequence, most of the burden of children's learning has fallen on their parents, with likely uneven results depending on the socio-economic characteristics of the family. Using a real time survey data, collected in April and early May 2020 in France and Italy, we estimate child fixed effects models to analyze how the lockdown has affected parents' evaluations of their children's emotional wellbeing and of their home learning process. The analysis also focuses on the role played by online classes, or other interactive methods, on children's home learning and emotional status. We find that the Spring 2020 lockdown had a stronger negative effect on boys, on children attending kindergarten (in Italy) or secondary school (in France), and on children whose parents have a lower education level. We also find that the increase in the time spent in front of screen is correlated to worse learning progresses and emotional status, while the opposite is true for the time spent reading. The use of interactive distance learning methodologies, which has been

1 The authors thank the participant to the 35th Annual Conference of the Italian Association of Labour Economists for their useful comments and suggestions. Hugues Champeaux and Francesca Marchetta acknowledge the support received from the Agence Nationale de la Recherche of the French government through the "Investissements d'avenir" program (ANR-10-LABX-14-01.) All errors remain ours.

2 CERDI, Université Clermont Auvergne, CNRS.
3 University of Perugia and IZA.
4 CERDI, Université Clermont Auvergne, CNRS.
5 University of Trento and IZA.

much more common in Italy than in France, appears to significantly attenuate the parents' negative perception of the impact of lockdown on the learning progress of their children.

Covid Economics 68, 10 February 2021: 69-99

Covid Economics 68, 10 February 2021: 69-99

1 Introduction

The COVID-19 crisis in Spring 2020 forced many countries around the world to close schools for a prolonged period of time, and teaching has been moved online on an unprecedented scale.[1] How much parents can help their children in dealing with education at home varies widely across families, as do the resources given to parents by their children's schools, since, even within the same countries or regions, schools have adopted different learning solutions. Consequently, the COVID-19 outbreak will lead to an increase in the inequality of human capital development for the affected cohorts of children.

This paper aims to analyze, at an early stage, the perceptions of parents in regard to how the COVID-19 lockdown has affected the use of time, the learning process and the emotional status of pupils aged 3 to 16 in France and Italy. We also explore the effectiveness of different distance learning solutions adopted by schools to mitigate the perceived negative effects of the lockdown. This preliminary evaluation is especially relevant both for short run and long run policy objectives. In the short run, as the actual health situation is still critical and schools are still being closed from time to time for certain periods in countries across the world, it is important to understand how distance education worked during the Spring 2020 lockdown and how it can be adjusted so that students do not lag behind. For instance investment in educational technology and teacher training may be a short term priority. More generally, it would be important to help policymakers identify those children who have suffered more during lockdown so that adapted educational program could be offered to them in the medium and long run.

To address these research questions and to identify policy priorities, we explore the time use of Italian and French children during the Spring 2020 lockdown through a descriptive analysis and we estimate child fixed effect models for parents' evaluations of children's learning and emotional status, using original data collected on a sample of families from April 7 to May 10, 2020.

Thanks to the availability of the individual data that were collected right in the middle of the lockdown concerning children's time-use, home schooling and emotional status, we have a large sample of children whose parents were interviewed via a real-time online survey. This allows us to offer the first comprehensive evaluation of the effects of the lockdown on children's learning and contribute to the few emerging studies that have already analyzed the heterogeneity of home schooling experienced by families during the lockdown, in different countries. Burgess and Sievertsen (2020) for instance describe the possible effects of the outbreak on children's education. Andrew et al. (2020) collected data on children aged 4-15 between April 29 and May 12, 2020 in the UK and find large variations in home learning resources provided by schools and in parents' ability to support home learning. They find that private schools are much more likely to offer online classes and, even in state schools, online classes are more likely to be offered to children living in richest families. Mangiavacchi et al. (2020) show that the emotional status and the quality of time-use of Italian children improved when fathers were more involved in childcare.

The cross-country focus on France and Italy is noteworthy since both countries are hugely affected by COVID-19 and their school systems are mostly public. This implies that the analysis would not be severely confounded by children's enrollment in private schools that are more likely to have better

[1] According to UNESCO, up to 192 simultaneous country-wide closures had affected 91.2% of the world's student population at the beginning of April (source: UNESCO Institute for Statistics Database, COVID-19 Impact on Education).

educational technologies, as shown for the UK (Andrew et al., 2020). At the same time, the comparison between France and Italy is interesting because their educational systems differ in terms of both policy priority and results (Woessmann, 2016), and in the way distance learning resources has been provided during the crisis. They also differ in terms of the duration of school closure during the 2019/20 academic year: Italy started on March 4, 2020, keeping schools closed until the end of the academic year; French schools closed on March 17, 2020 and gradually reopened starting from May 10 on a voluntary basis. Moreover, even if these two countries share similar normative determinants of the time cost of children, like Catholic values, Latin cultural heritages and asymmetric gender roles, they present remarkable differences in terms of public spending for families, family policies and childcare services (Anxo et al., 2011; Pailhé et al., 2019).[2]

Our paper also contributes to the literature that analyzes the relevance of time at school for children's cognitive development. Lavy (2015), for instance, estimates the impact on academic achievement of differences in instructional time across countries and finds these differences to cause significant variation in test score outcomes: one more hour per week over the school year in the main subjects increases test scores by around 6% of a standard deviation. Other studies focused on the increase in educational inequalities when schools remain closed for a long period due to different exogenous shocks. For instance, Jaume and Willén (2019) found that being exposed to the average incidence of strikes during primary school reduces the labor earnings of males and females by 3.2% and 1.9% respectively. Even if we cannot measure children's test scores or use other objective metric for cognitive development, we can rely on the parental judgment of their educational progress with home learning and we can identify the potential detrimental effects of the lockdown on the way children allocated their time between productive and unproductive activities.

Related to this, we believe that our study provides a contribution to the growing literature on the allocation of children's time out of school as one of the determinants of cognitive and socio-emotional skills (Fiorini and Keane, 2014; Del Boca et al., 2017). During the Spring 2020 outbreak, children stayed at home for three months in Italy and two months in France and they had to completely reorganize their time. It is thus meaningful to study how this reorganization is related to their learning ability and their emotional status. In particular, if children's available time has been spent in productive activities, this could at least partially compensate for the detrimental effect of school closures on their cognitive development. Following the evidence on the importance of child's reading and parental reading to children at age 0-5 (Kalb and Van Ours, 2014), as well as the evidence of the detrimental effect on cognitive development of time spent in front of a screen (Walsh et al., 2018), we asked parents about their children's time spent reading (or listening stories) and time spent in front of the screens, both before and during the lockdown. This allows us to evaluate how children reacted to the lockdown in terms of time-use and test the mediating effects of time re-allocation on learning and emotional well-being.

Finally, the current paper contributes to the literature on *Edutech* and distance learning, evaluating parents' perceptions about the effectiveness of different distance learning approaches. To the best of our knowledge, existing economic literature focuses on college students, who were the subjects of a number of experiments (Coates et al., 2004; Xu and Jaggars, 2013; Bettinger et al., 2017; Pellizzari et al., 2019),

[2]According to the OECD family database for 2015, France spends 3.7% of GDP on family policies, while Italy only spends 2.5%.

Covid Economics 68, 10 February 2021: 69-99

CEPR PRESS

which showed mixed evidence on the effects of online classes on achievement compared to traditional lectures. In regard to the differences between alternative online learning solutions, Figlio et al. (2013) analyzed the difference between live classes and watching videos with the same lectures on the internet in a experimental settings and found that live-only instruction is slightly better than internet instruction. School closures during the lockdown obliged all teachers to suddenly adopt distance learning strategies, but often without receiving clear guidelines from their superiors. Schools and teachers were thus free to choose from a large typology of methods, which differ in the degree of interaction. This offers an ideal experimental setup to study the impact of distance learning on younger students. In the survey, we asked parents to report which distance learning methods were offered to their children. This allows us to test the difference in the parents' evaluation of their children's home learning and emotional status when live classes or chats have been implemented compared with less interactive methods, like sharing materials or videos.

We find that, on average, Italian parents are more worried about their children's home learning process with respect to their French counterparts, and this is particularly true when French parents are highly educated or they were at home during the lockdown. As to children's emotional wellbeing, the negative effect of the lockdown that we estimate on the basis of parents' perceptions is twice as large for Italian children. We estimate that both French and Italian children increased the time spent reading by 0.3 hours on average, and the time spent in front of screens (out of classes) by 1.3 hours on average during lockdown. According to the parents, the negative effects of lockdown on both their children's learning and emotional status is attenuated when children spend more time reading, while it is amplified when they spend more hours watching TV or in passive screen activities (YouTube, social media, and similar). Our regression results suggest that younger children (aged 3-6) suffered more from the lockdown, both in terms of learning progresses and emotional status, in particular in Italy, where 40 per cent of them did not receive any type of distance learning support from their teachers. Children attending secondary schools also experienced significant losses in terms of learning progress when they could not attend online classes, and this is particularly evident in France, where almost 30 per cent of children did not benefit from interactive distance learning methods. In general, the use of interactive methods seem to attenuate the negative evaluation parents give on the lockdown effects. It is interesting to observe that important differences emerge in the share of students that could attend online lectures in the two countries: it was substantially larger in Italy for all school levels. Within countries, substantial heterogeneity exists, always favoring older students.

The rest of the paper is organized as follows. Section 2 describes the institutional settings, focusing on education systems and the management of distance learning in the two countries. Section 3 describes data and presents a descriptive analysis on the time use of children before and after the lockdown. Section 4 presents the estimation strategy and the results of the analysis of the effect of lockdown on the learning process and the emotional status of pupils. Section 5 concludes.

2 Education systems and distance learning in France and Italy

The organization and governance of the educational system explain the large international differences in student achievement combined with family background (Woessmann, 2016). Family background and institutions are quite likely to also shape the educational penalty that children of different countries

Covid Economics 68, 10 February 2021: 69-99

may suffer from the school closure period that was undertaken to limit the spread of the COVID-19 virus during the Spring 2020 lockdown. Andrew et al. (2020) have already shown the importance of families' economic situation as a determinant of children's time use during the lockdown in the UK, but pre-existing educational institutions may also matter. It is not merely that differently organized schools may have offered distance learning solutions that are likely heterogeneous in quality, but also that pupils who have been trained to be self-directed in their academic work may have experienced lower losses. In addition, as the burden of children's education substantially fell on the shoulders of their parents during the lockdown, family support policies may also have played a significant (although indirect) part: when families feel that they are supported by the state (and society at large) in their task of raising their children and that the quality of this process really matters for the society, their involvement and effort may be larger, even in such an emergency context. The opposite may also be true: when not supported by the state, parents may feel that they need to compensate for state or school absenteeism.

French and Italian school systems, as well as their family welfare policies, share some similarities but also have significant institutional differences. The French and Italian education systems are similar at first glance. Table A1 shows that they are both largely public systems (6.9% of pupils attend private schools in Italy, in France this is around 21.5%, although almost entirely publicly funded) characterized by compulsory education until 16 years of age. Both countries have four levels of education, with lower secondary education lasting three years in Italy and four in France, and higher secondary education lasting five years in Italy and three in France. Teachers have about the same starting salary (about 30K dollars PPP for kindergarten and primary education, and about 32.5K for secondary education). Despite being apparently similar, the French system achieves better results. According to the 2018 OECD PISA report French scores are higher than Italian in all subjects: reading, mathematics, and science. French schools achieve higher attendance rates at all levels, but particularly at early ages.[3] Italian adolescents also have lower expectations for academic achievements: less than 30% expect to complete tertiary education, with respect to almost 40% of French students.

Italian students go to school more days during the year (200 vs 162, about 23.5% more) for primary and lower secondary levels, but school days are much more concentrated, as summer holidays last 4/5 weeks more. French classes are larger by more than 4 students and in general French teachers have more pupils at all levels. Other characteristics of the school organization are likely to be relevant for achievement: for instance, in Italian schools most children in primary and lower secondary school maintain the same teachers for the entire duration of the school level, while in France this typically does not happen, with most teachers changing every year. In addition, classmates and classrooms change in France from one year to the next, and, for older children, even during the day. Higher secondary schools in France tend to have dedicated counselors (more than 50% versus about 10% of Italian schools) to help children in their transition towards tertiary education.

Public expenditure per student is larger in France (except for primary education), especially for higher secondary schools, and overall public expenditure on education is almost 50% greater in France in terms of percentage of GDP. Finally, French schools have much younger teachers: primary school teachers under 30 make up 12% of the total versus 1% of Italy, while the share of teachers aged 50 or more are 22% of the total versus 56% of Italy. Finally, in France many more teachers are men.

Another important difference across the two countries is the cost of raising children, which is higher,

[3]Since 2019, education has been compulsory from the age of 3 in France.

Covid Economics 68, 10 February 2021: 69-99

in terms of childcare, for Italian parents. This is particularly true for large families and for families with preschool children. Italian mothers adjust for this burden by substituting housework with childcare and reducing their leisure time more than man (Pailhé et al., 2019). The lower cost of raising children in France is clearly related to the stronger social support provided to families that has a long tradition and is effective in keeping the fertility rate quite high in the country.[4] Family policy spending is one of the largest among the OECD countries, accounting for about 3.7% of GDP in 2015, which is around 50% larger than the Italian expenditure (2.5% of the GDP). Child benefit is generous in France, especially from the third child, while in Italy at the time of writing such benefit is absent.[5] In addition, there are widespread subsidized day-care centers for children aged 2 months to 3 years old with long opening hours (up to 11 hours per day for day-care centers), as well as school recreation centers in all pre-schools and primary schools where children can stay before and after school time for a low price. On the other hand, in Italy day care centers are scarcely present in the territory, mostly relegated to the private sphere and, in some regions, they are quite expensive, implying that Italian parents are often forced to rely on grandparents or other informal childcare solutions, or sacrifice their job to take care of their children.[6]

In regard to children, according to Cardoso et al. (2010), Italians tend to study more at home (about 154 minutes per day vs 93), watch less TV (99 minutes vs 118) and socialize less (38 minutes vs 52) with respect to their French counterparts. This is confirmed by more recent OECD data: Italian children socialize less (22.9% of them do not invite friends to their homes to play or eat, etc., versus 13.8% of French children) and do less regular leisure activities, such as holidays, swimming, riding their bikes, football, and so on: a stunning 55% of Italian children do not do any of these activities regularly, versus 39.3% of French children. These figures highlight the fact that for many Italian children, school represents the only place where they can develop social skills and the closure periods may have a relevant negative impact also in this field of social capital development.

2.1 Education during the pandemic

The COVID-19 pandemic hit early both in Italy and France, with the first confirmed cases occurring in the last days of January. The contagion evolution forced both governments to act early with nationwide restrictive measures. In Italy, all schools closed on March 4 (some regions closed schools a couple of weeks earlier), while the French government followed early on, closing schools on March 16. By March 17, both countries had already implemented home confinement measures and by March 23 both countries had already issued travel limitations to citizens. These measures stayed in place until May 11, when both counties started removing limitations. France gradually reopened schools at the end of the lockdown, with full re-opening set on June 22. In Italy schools started straight after the summer holidays, that is on September 14 for most regions.[7]

Even with the closure of school buildings, educational activities were maintained by the French and

[4] The fertility rate in France in 2019 stood at 1.8 children per woman, above the OECD countries average of 1.6 and well above the 1.3 children per woman recorded in Italy.

[5] The recently approved financial law introduced a new unconditional child benefit that is scheduled to start on July 2021.

[6] According to the European Platform for Investing in Children (EPIC), in Italy more than man only 25.7% of children aged 0-3 attended formal ECEC, while the figure almost doubles, reaching 50.0%, in France. As to kindergarten age, the enrollment rate is much larger in both countries, 91.0% in Italy and 94% in France.

[7] Only daycare services were allowed to work, under strict restrictions, starting from July 1.

Covid Economics 68, 10 February 2021: 69-99

Italian governments. As the pandemic was not anticipated, schools and teachers from both countries benefited from some degree of freedom regarding the implementation of distance learning methods. In Italy, schools were left to their own initiative by the government, which provided some guidance through Ministry of Education guidelines and its website, which indicated the software platforms that could be used. Nevertheless, schools had almost total freedom in deciding if and how to implement distance learning solutions. In France, the Ministry of Education decreed "pedagogical continuity" for the pupils early on, providing official chatrooms and educational platforms, but, as with Italy, teachers were not obliged to use them, and instead were free to decide what type of learning methods to offer to their students.[8]

In both countries, children differ in terms of IT equipment availability, in terms of parental investment (which may depend on the parents' level of education and working status during lockdown) and on the types of distance learning they benefited from during lockdown.[9] All these factors likely generated high heterogeneous impacts of schools closures on children's learning achievements and emotional status, as well as different behavioral reactions.

3 Data and Descriptive Analysis

We use original data, specifically designed to study the effects of the lockdown on families and their components, which we collected through a real-time online questionnaire.[10] We started to disseminate the surveys on April 7 in Italy and on April 21 in France. Both surveys were available until the end of the outbreak, on May 10. The final assembled data provide information on 3,352 families with children in Italy and 2,154 in France. As participation in the surveys was voluntary with no sampling strategy, we cannot claim representativity of the populations of reference at national levels. For Italy, thanks to the relevant sample size and the ability to reach all the regions and different socio-economic groups, the geographical and family type distributions are in line with the national statistics reported by ISTAT (see Table A2, Panel A). The only notable exceptions are for the South of Italy, which is slightly under-represented, and the share of mono-parental households, which is strongly under-represented. The situation is similar for France: the sample is relatively well balanced at the geographical level (excepted for the Paris area), while single parents are still under-represented (see Table A2, Panel B).

The survey includes basic information on the respondents' and their partners' personal characteristics including gender, age, location of residence, highest level of education, marital status, and parental status. It also collects detailed current and retrospective information on the respondents' and their partners' labor market participation (including employment, sector of employment, labor supply evolution and hours of teleworking) and on the division of household tasks.[11] The surveys included a

[8] For instance, the CNED platform 'Ma classe à la Maison' was used by about 24 per cent of lower secondary students (DEPP, 2020).

[9] About 9 per cent of school principals declared that all or most of their students had outdated, defective or unsuitable equipment (DEPP, 2020)

[10] Both surveys were jointly developed with an European team of researchers. Similar surveys were also disseminated in Spain (Lidia Farré and Libertad Gonzales), Germany (Christiane Schwieren) and Austria (Doris Weichselbaum). The French and Italian surveys added a specific section on children. The anonymous questionnaires were disseminated through targeted campaigns on the main social networks, such as Facebook and Twitter, participation was on a voluntary basis and no rewards were offered upon completion of the questionnaire. The English translations of the French and Italian questionnaires are available upon request.

[11] For many items we asked the respondent to recall which was the situation just before the lockdown and which is the

CEPR PRESS

Covid Economics 68, 10 February 2021: 69-99

Table 1: Children's samples in France and Italy - descriptive statistics

Variable	France mean	France sd	Italy mean	Italy sd
Girls	0.49	0.50	0.49	0.50
Age	9.60	3.78	7.95	3.75
Children living in two-parents households	0.87	0.33	0.92	0.26
Children living in one-parents households	0.13	0.33	0.07	0.26
Children living in one-child households	0.20	0.40	0.27	0.44
Children living in two-child households	0.52	0.50	0.56	0.50
Children living in three-child households	0.23	0.42	0.15	0.36
Children living in four-child households	0.04	0.19	0.01	0.12
Children living in households with five or more children	0.01	0.12	0.00	0.07
Children attending kindergarten	0.25	0.43	0.32	0.47
Children attending primary school	0.42	0.49	0.44	0.50
Children attending lower secondary school	0.25	0.43	0.16	0.37
Children attending upper secondary school	0.08	0.27	0.08	0.27
Age of mother	39.79	5.81	41.85	5.25
Age of father	41.98	6.56	44.47	5.90
Children whose mother has a university degree	0.57	0.49	0.58	0.49
Children whose father has a university degree	0.40	0.49	0.37	0.48
Children whose mother was at home during lockdown	0.79	0.41	0.77	0.42
Children whose father was at home during lockdown	0.63	0.48	0.52	0.50
Children whose mother is working before lockdown	0.85	0.35	0.81	0.40
Children whose father is working before lockdown	0.93	0.25	0.96	0.19
Children whose mother is working during lockdown	0.70	0.46	0.55	0.50
Children whose father is working during lockdown	0.78	0.41	0.74	0.44
Observations	3272		4477	

specific section on children living in the household. We asked about the parental time use in terms of number of hours spent on active childcare and home-schooling, their subjective opinions on the child's educational improvement during lockdown, as well as on the emotional status of children and on the relationship between the parents and children. We also asked questions about children's time use before and after the closure of the schools. In particular, we collected information about hours spent studying, performing extracurricular activities, reading and watching tv (and other passive screens). Finally, we collected data about the distance learning methods offered to each child and about the availability of IT equipment, such as computers, tablets, and smartphones, in order to identify situations of digital divide.[12]

The two samples are composed of 3,272 children in France and 4,477 children in Italy, totaling to a general sample of 7,749 children (see Table 1). Children are balanced on gender in both countries and are a little older in France (the average age is 9.6) than in Italy (average age 7.9). 87% of children in France and 92% in Italy are living in two-parents households. More than half of the children live in two-children households (52% in France and 56% in Italy). The incidence of children living in families

current situation.

[12]For France, just 0.5% of primary and secondary school children do not have access to IT equipment. While in Italy the figure rises to about 9%, we do not observe significant differences in the impact of the lockdown with respect to the rest of the sample.

Figure 1: Children's daily activities before and during lockdown

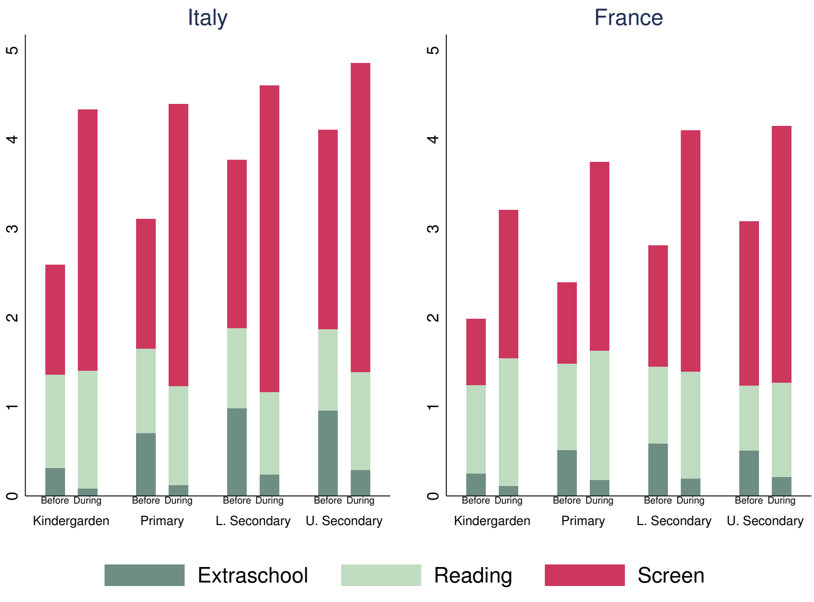

with three children is higher in France (23% against 15%), reflecting the difference in fertility rates between the two countries. Italian parents are a bit older, reflecting the older age at first child of Italian parents, while the incidence of mothers and fathers with university degrees and their work status before the lockdown are similar in the two countries. Our data confirm that mothers' labor supply was highly affected by lockdown, as already shown by Del Boca et al. (2020) and Mangiavacchi et al. (2020) for Italy, by Farré et al. (2020) for Spain and by Andrew et al. (2020) for the UK. Table 1 shows that Italian mothers were more affected than French ones: the incidence of children whose mothers are working moved from 81% before the lockdown to 55% during the lockdown in Italy, and from 85% to 70% in France.

3.1 Children's use of time during lockdown

Figure 1 shows the evolution of time devoted to three relevant daily activities of children: extraschool activities, reading and screen time. The color green is assigned to "productive" activities from a human capital accumulation perspective.[13] Time investments made by children in productive activities has shown to be particularly important during adolescence (Del Boca et al., 2017; Giménez-Nadal et al., 2019). During lockdown, we observe the expected strong reduction in extraschool time, that reduced to less than 10 minutes per day for both countries. This reduction seems to impact Italian children

[13]The figure does not report the reduction in school hours, which is self-evident. According to the previous literature (Lavy, 2015), this huge reduction in school time is likely to have a negative effect on future test scores, which is likely to be larger in Italy where schools remained closed for a longer period of time.

Figure 2: Change in daily time devoted to reading and screen

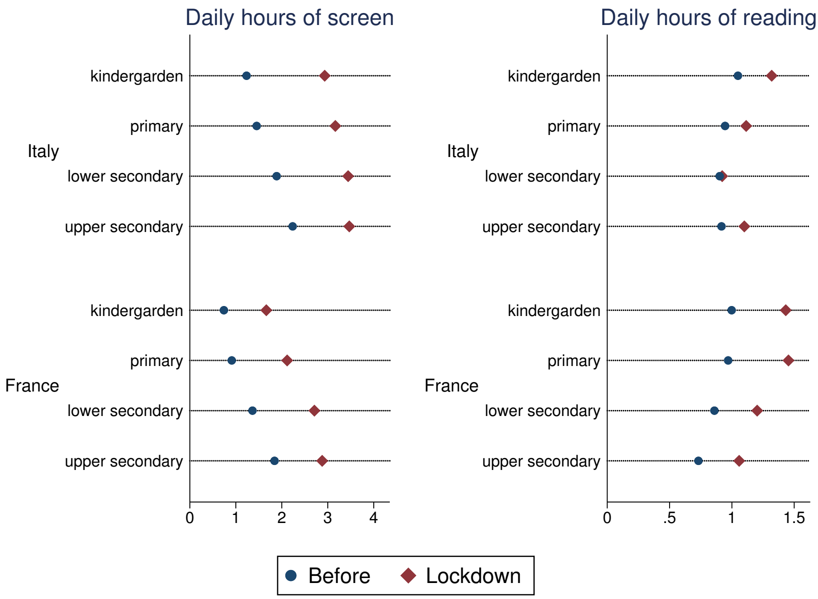

to a greater degree, since they spent more time on those activities before the lockdown (36 minutes per day vs 25 minutes per day in France). The reduction in the time spent in extraschool activities is compensated for by an increase in reading time in France. In Italy this compensation is observed only for pre-school children, where reading is performed by parents, which implies a reduction on overall productive time for all other children.[14]

During the lockdown period, children in both countries have allocated a significant part of their time that was previously devoted to school to passive screen time. Time spent watching TV or time spent on the internet (videos, socials) doubled in both countries, increasing from 1 to 2 hours on average for French children and from 1.5 to 3 hours on average for the Italian ones. Figure 2 shows that, although important heterogeneity exists across school levels with respect to the initial amount of screen exposition, the increase was almost the same across all levels (see also Andrew et al., 2020, for similar results in the UK).

3.2 Distance learning methods and children's educational progress

Upon closure in March 2020, in both countries teachers had to put in place distance learning activities, even if they were not prepared in the slightest for such a task. The Ministries of education provided some guidance and offered some software platforms that could be used, but schools and teachers had

[14]Reading time was almost one hour in both countries before schools closure, in France it increased up to 1 hour and 20 minutes, 10 minutes more than in Italy. See Kalb and Van Ours (2014) on the importance of reading for children's cognitive development.

Covid Economics 68, 10 February 2021: 69-99

Figure 3: Distribution of different home learning resources provided by schools - by country and level

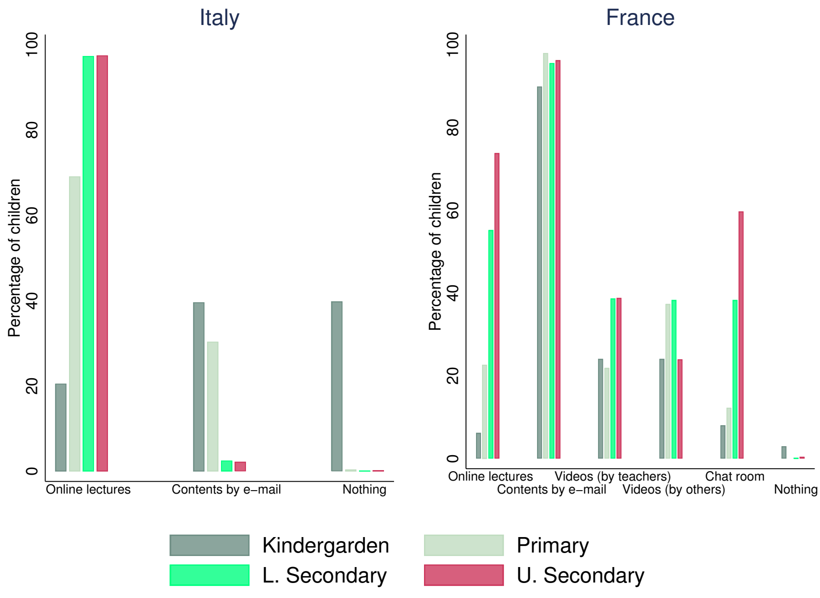

Covid Economics 68, 10 February 2021: 69-99

almost total freedom in deciding if and how to implement distance learning activities. This, of course, caused an extremely heterogeneous response (DEPP, 2020).

In the survey we asked parents to report which distance learning activities were implemented by their children's teachers and the parents' perceptions about their children's learning progress during the lockdown. In the Italian questionnaire, we asked parents if the teachers (i) shared only educational material, by mail or other digital platforms; (ii) also offered live online lectures; (iii) did not offer any distance learning activity. We also asked about the number of hours of online classes offered to the child. In the French questionnaire we asked more details about the activities offered by teachers: parents had to indicate if their children were involved in any of the following activities: attended online classes, participated in chats with teachers and classmates, received videos created by their teachers or by others, or received educational content via e-mail or via a platform.

Figure 3 shows substantial differences between the two countries and across school levels. In both countries, almost all secondary school children received content and assignments by e-mail or through a platform, yet while in Italy almost all of them attended online classes as well, the percentage of online lectures for France stands at 70%. In primary school, almost all children received content by email or via a platform, while on-line classes were offered to 65% of Italian students and to 20% of French ones. For kindergarten, it is interesting to observe that almost 42.1 per cent of Italian children were not involved in any activity, while only 3.6 per cent of French children were in the same situation. Online classes are unsurprisingly less common for young children (6 per cent for France, 19 per cent for Italy), with teachers preferring less interactive distance learning methodologies.

Figure 4: Differences in learning evaluations when interactive learning resources were offered by schools

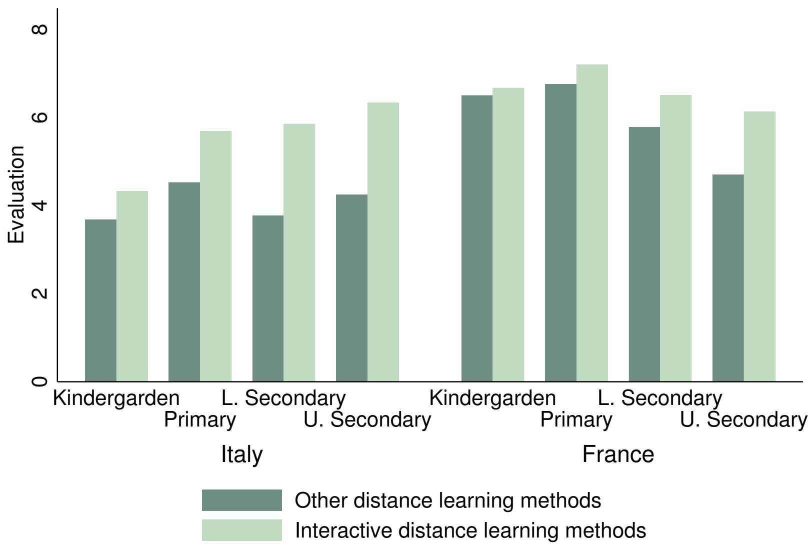

Note: Pupils without any educative materials provided by their teachers are dropped in this figure.

The type of distance learning activities proposed by the teachers seems to drive the parents' evaluation of their children's learning during the lockdown, especially for older children. Figure 4 describes a variable indicating parental judgment (from 1 to 10) on the quality of children's learning progress. When this variable takes the value 10, it means that the parents evaluated the children's learning with distance learning activities to be as good as during a normal school period. For both countries and at every level, the parents' judgment was better when children were able to follow interactive lectures (on-line classes for Italy and online classes and live chat for France) and the difference grows larger with school level. In Italy, the overall judgment of parents was particularly low at all levels when no interactive classes were offered and the implementation of interactive classes, as well as their intensity, substantially improved parents' evaluations.[15] Interestingly, the judgment of French parents' was better for all school levels, except for children in upper secondary school, with the difference among the two country being particularly important for children in kindergarten or primary school. This significant difference may be due to two different, but not alternative, explanations. On the one hand, it may depend on the type of inputs children received before the lockdown: the French system in kindergarten and primary school prepares children to be more independent and more flexible to changes (see Section 2), so French children may have adapted better to homeschooling. On the other hand, the level of parental stress was likely to be higher in Italy at the time of the survey.[16] This could have biased

[15]In kindergarten, children who attend online classes do it for very few hours per week, while secondary school children almost maintained the same teaching schedule they had before the lockdown.

[16]The situation in Italy was more severe in terms of the number of COVID-19 cases and deaths. According to the WHO Coronavirus Disease (COVID-19) Dashboard, by May 11, France had experienced 137,073 cases and 26,338 deaths and

Covid Economics 68, 10 February 2021: 69-99

Figure 5: The evolution of emotional status and relationship with parents by school level and country

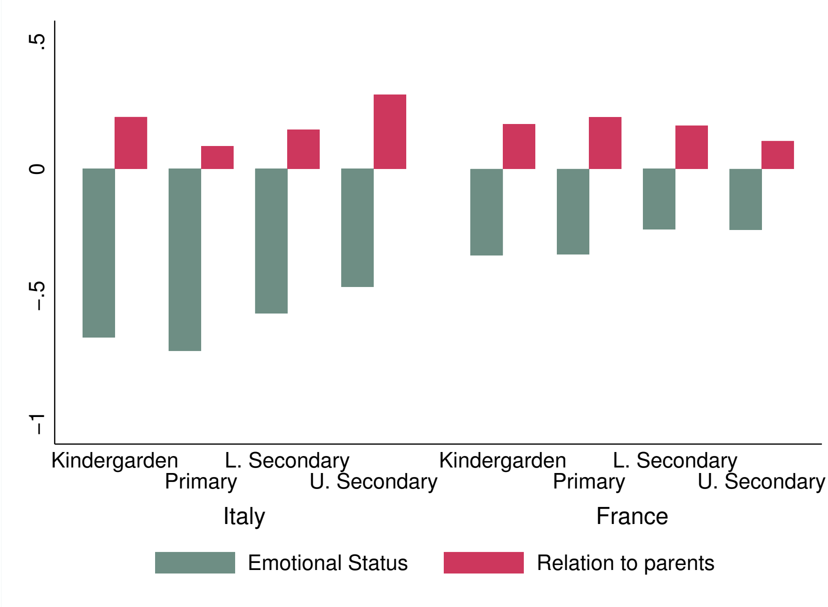

Covid Economics 68, 10 February 2021: 69-99

parental perceptions towards a negative evaluation of their children's learning in Italy. Moreover, on April 13, President Macron of France announced that schools would be reopened starting from May 11. This may have reassured French parents about the temporary nature of school closures, while no statements was given by the Italian Prime Minister, Giuseppe Conte, in regards to school reopening, clearly indicating a more than likely reopening after the summer vacation (which actually happened).

3.3 Parents' evaluation of children's emotional status

The COVID-19 outbreak increased the stress and burden on parents and the social isolation of children from their peers and teachers. This situation may also affect the socio-emotional skills of children, such as their mental health, wellbeing, and behavior. The risk of an increase in socio-emotional problems may be higher for those living in low educated and the poorest households, who have lower socio-emotional skills in normal periods also (Attanasio et al., 2020). Boys are also more at risk since they are more likely to experience behavioral issues than girls (Bertrand and Pan, 2013) as well as all adolescents. On the other hand, positive interactions between parents and children can improve socio-emotional skills (Moroni et al., 2019). For these reasons, in the survey we asked parents to report the evolution of their children's emotional status and the evolution of the parent-child relationship. For both questions the response items were: "it is much worse", "it is slightly worse", "remains stable", "it is slightly better", "it is much better". We recoded the variables in order to have zero when the emotional status and the relationship with parents were judged stable, and values -1 and 1 for the largest variations.

Italy 219,070 cases and 30,560 deaths.

Figure 5 plots these two variables by country. In general, parents report, on the one hand, an overall reduction in children's emotional status, on the other hand a slight increase in the quality of the parent-child relationship. Parents in both countries are slightly more worried for younger children (those in kindergarten and primary school) compared to those in secondary school. Italian parents appear, again, more worried about their children's emotional status when compared to French parents.

4 Estimation Method and Results

In this section, we analyze first how the lockdown has affected children's learning process and emotional wellbeing, according to their parents' perceptions. We then move on to analyze the role played by interactive learning methods on the same outcomes.

4.1 Children's learning and emotional status during the lockdown

For the empirical analysis, we use two indicators as dependent variable in the regressions: i) parental evaluation of the child's educational progress in a 1 to 10 scale (1 for "not progressing at all" and 10 for "progressing at the same pace as when she/he was attending classes at school");[17] ii) parental evaluation of the child's emotional status in a -1 to 1 scale, as explained in section 3.2 above.[18] Both variables are interpreted as a variation with the lockdown, which allows us to perform fixed effect regressions of the form:

$$Y_{it} = \theta LD + \beta FR \cdot LD + \gamma X_{it} + u_i + e_{it}, \tag{1}$$

where Y_{it} is the selected outcome for individual i at the time t, LD is the temporal dummy equal to one for the period during the lockdown, which is interacted with FR, a dummy equal to one for French children. θ thus measures the impact of the lockdown in Italy on the dependent variable, while β shows the differential impact of the lockdown in France.[19] X_{it} is a set of child-specific time-varying regressors which include the time spent in front of a screen and reading, in hours per day, before and during the lockdown, and two dummies indicating whether the mother and father were actually working for any amount of time (including smartworking or teleworking) before and during the lockdown.[20] u_i represents child fixed effects and e_{it} is the idiosyncratic error. Standard errors are clustered at regional level.

To analyze the different impacts that the lockdown may have had on different population groups, and to avoid an excessive set of interactions, we prefer to run the same model on different sub-samples. We thus split the sample by gender, by education level attended, by the education level of both parents, and by the work status of both parents during the lockdown. We also separately look at children with siblings or those without, and at children living with a single parent.

[17] We fixed educational progress equal to 10 for the period before the lockdown. This means that, by definition, the coefficient can be either negative or equal to zero. The structure of our questionnaire does not allow us to capture an improvement in the learning process during lockdown.

[18] We fixed the emotional status equal to 0 for the period before the lockdown

[19] Similar results are obtained when we estimate a fully interacted version of this model, where all covariates are interacted with the dummy FR.

[20] This variable is thus not referring to a person who has a job contract or not, but rather if s/he had actually worked in the reference period. Several types of workers, especially in the public sector, may have been receiving a salary without working any hours during the lockdown.

Covid Economics 68, 10 February 2021: 69-99

Covid Economics 68, 10 February 2021: 69-99

Table 2: The impact of the lockdown on education and emotional status

	Learning	Learning	Emotional status	Emotional status
	(1)	(2)	(3)	(4)
Lockdown	-5.135***	-4.954***	-0.655***	-0.568***
	(0.091)	(0.152)	(0.037)	(0.050)
France · Lockdown	1.722***	1.697***	0.349***	0.303***
	(0.114)	(0.125)	(0.045)	(0.038)
Mother is working		0.077		0.044
		(0.127)		(0.045)
Father is working		0.026		-0.000
		(0.140)		(0.056)
Screen (time)		-0.123***		-0.070***
		(0.047)		(0.021)
Reading (time)		0.209**		0.114***
		(0.093)		(0.028)
Child fixed effects	Yes	Yes	Yes	Yes
Observations	15,412	13,439	15,412	13,439
Within R-squared	0.620	0.628	0.126	0.169

All results were estimated using the fixed effects model on panel data from Italian and French 2020 COVID-19 online surveys. "Lockdown" is a dummy equal to one for the observations during the pandemic. "France" is a dummy equal to one if the child is French. Here it is interacted with "Lockdown". "Screen" is a continuous variable reporting time spent by the child in front of screens before and during the lockdown. "Reading" is the time spent reading. "Mother (Father) is working" is a time-variant dummy equal to one if the child's mother (father) is working during the period.
Each estimation controls for child individual fixed effects. Standard Errors in parentheses are clustered at region level. ***, **, * indicate significance at the 1%, 5% and 10% level respectively.

Table 2 presents the results of the regressions for the parental evaluation of the child's learning process (columns 1 and 2) and for her/his emotional status (columns 3 and 4), with and without covariates for the whole sample.

According to parents, the negative impact of the COVID-19 lockdown on their children's learning process has been substantially stronger in Italy, by 1.7 points on a 1 to 10 scale. The results are almost unchanged when accounting for the control variables. The increase in screen time is strongly associated with a worse evaluation of the learning progress, while an increase in the time spent reading improved parents' evaluation, with this being in line with previous literature on the impact of reading on human capital development (Kalb and Van Ours, 2014). Parents' work status does not significantly relate to their children's learning process. It is worth noting the particularly large value of the within R^2 even without covariates. It indicates that the lockdown impacted (within both countries) how parents perceived the learning progress of their children in a relatively homogeneous way.

A similar pattern is observed when analyzing children's emotional status: the impact is clearly negative in both countries, but it is almost twice as large in Italy. In a -1 to 1 scale, according to parents, Italian children worsened their emotional status by almost 0.6 points, while French ones only by 0.3 points. The increase in screen time has a negative association with emotional status, while

Figure 6: The impact of the lockdown on education - different subsamples and by country

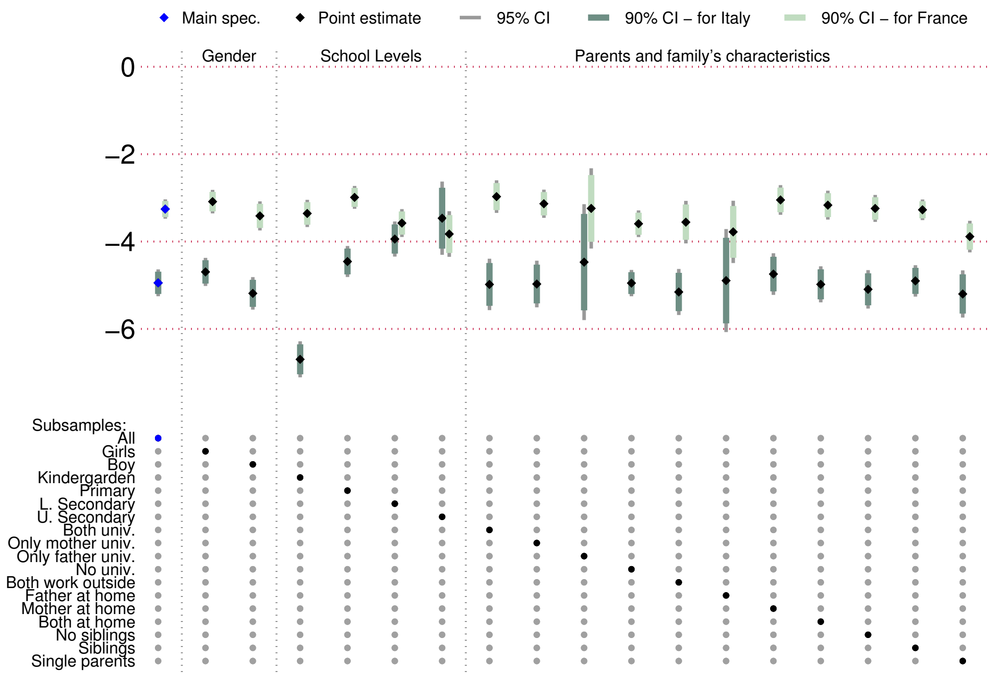

the increase in reading has a positive effect. Again, we find no evidence of a direct role of parents' work status. In this case the within R^2 is much smaller, indicating a more heterogeneous response of children's emotional status to the lockdown.

In what follows, we explore the heterogeneity in the response to the lockdown in different sub-populations. Figures 6 and 7 plot the lockdown coefficient values, as well as the 95 and 90 percent confidence intervals for learning evaluation and emotional status respectively. The coefficient is reported separately for France and Italy and for each sub-sample of the population that we examined.[21]

Looking at the parents' judgments, the lockdown has been more detrimental to learning achievement for boys than for girls in both countries. Notable differences across school levels emerge. In Italy preschoolers seem to have particularly suffered in terms of learning achievement with respect to older children, and, more in general, Italian parents give better evaluations of the learning progress of their older children. For France, we do not observe the same dynamics. Parents seem relatively satisfied with the learning progress of their primary school children, while they give worse evaluations for children in kindergarten and particularly for the ones in secondary school. Tentative explanations for these

[21] Full estimation tables are reported in Tables A3 and A4 in the Appendix.

Figure 7: The impact of the lockdown on emotional status - different sub-samples and by country

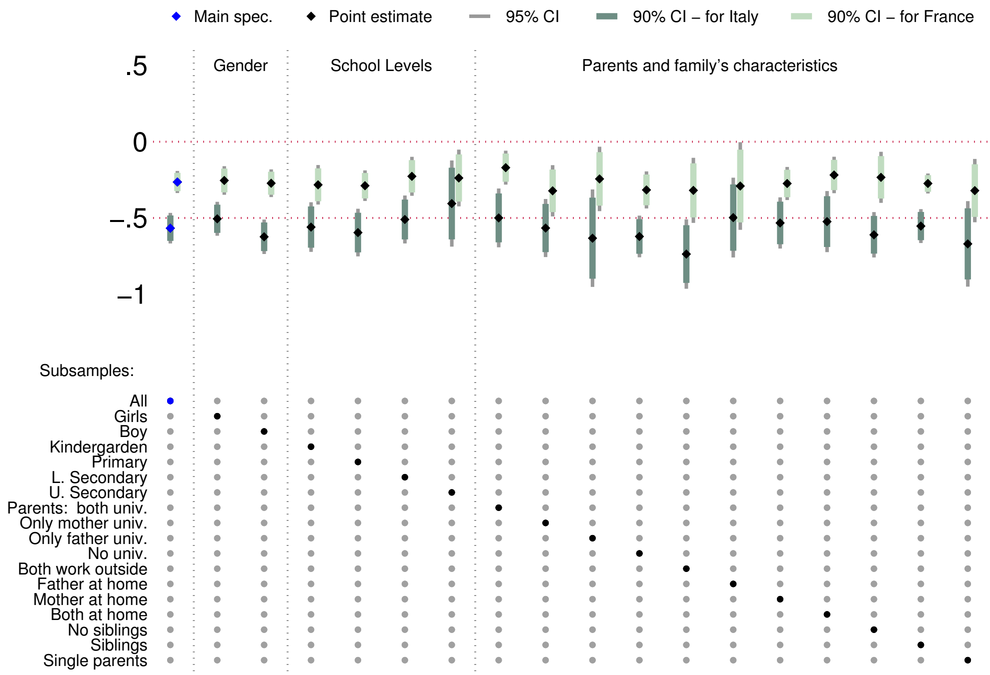

different results may be related to the different types of distance learning methods used by teachers in the two countries. Italian secondary school pupils attended more online classes, which seem to be quite appreciated by Italian parents (see below), and this may have improved their parents' judgment. As pre-school pupils, the extreme negative evaluation of Italian parents is probably explained by the fact that about 40 per cent of them did not receive any learning material from their teachers during the lockdown, as presented in section 4.2 below. Looking at parents' characteristics, in France we see that when the mother is not at home, as well as when both parents do not have an university degree, they seem to be more worried about their children's education. More educated parents, as well as parents who were at home during lockdown, are likely to be more comfortable in taking care of their children's education. Finally, French single parents seem to be more worried about their children's learning progress, while this is not the case for Italian ones. We do not see any differences among children with or without siblings.

Looking at the impact of lockdown on the emotional status, we see again that boys seem to suffer more than girls, according to their parents, but only in Italy. Parents report a worse emotional status

Covid Economics 68, 10 February 2021: 69-99

Covid Economics 68, 10 February 2021: 69-99

for younger children.[22] Much of the negative effect on emotional status may be due to their very limited interactions with peers. For older children, this reduction in person interaction could have been partially compensated by virtual interaction, which may have mitigated the negative effect of lockdown on their emotional status. As for learning, we observe that university educated parents are less worried about their children's emotional status. Children with siblings appear to have suffered less from an emotional point of view in Italy, but not in France. Children living with single parents present a larger decrease in their emotional status in both countries.

4.2 Distance learning methods

In France and in Italy, several distance learning methods were quickly implemented just after the closure of the schools. As there was no major guidance from both governments related to the educational materials or education monitoring, teachers were relatively free to decide how to teach. For Italy, starting from the replies to the question on distance learning described above in section 3.2, we are able to build three dummy variables that take the value of 1 respectively when: (i) no materials was provided by the teachers, (ii) only material without interactive content was provided by the teachers (Homework), (iii) full or partial interactive content was provided by the teachers (OnlineL). In the first situation, pupils had no relationship with their teachers or school and did not follow any exercises. In the second, they were connected with the teachers using emails or internet platforms and did homework. In the last one, they both followed online lectures and received materials by email.

In France, parents can record the information on their children's education by selecting several options for distance learning. They can click on six choices: no material provided by the teachers (1), chat room with other pupils and the teachers (2), pedagogical videos from other teachers (3), pedagogical videos from their teachers (4), material provided by emails without interactive content (5), online lectures (6). Putting aside the first option, the other choices are not independent and individuals can select the "chat room" option and "pedagogical videos" for example. In order to compare our results between the two countries in the regressions, we gather together all individuals with only emails and videos from their teachers in a sole modality and we consider them as receiving only material without interactive content: the dummy Homework takes value 1 for them. All individuals with at least one form of interactive content (online lectures or chat room) were considered to be receiving interactive content and take the value 1 to the dummy OnlineL. We then integrate the baseline specification in equation (1), adding interactions between the lockdown dummy and the dummies related to the distance learning methods and we follow the basic estimated fixed effects model on education level sub-samples as:

$$Y_{it} = \theta LD_t + \beta_1 OnlineL_i \cdot LD_{it} + \beta_2 Homework_i \cdot LD_{it} + \gamma X_{it} + u_i + e_{it} \tag{2}$$

for children at kindergarten education level; and:

$$Y_{it} = \theta LD_t + \beta OnlineL_i \cdot LD_{it} + \gamma X_{it} + u_i + e_{it} \tag{3}$$

for children at primary and secondary levels. We prefer to split the sample into different school levels for this part of the analysis since we believe that interactive distance learning methods are not

[22] Estimates are less precise for children in upper secondary school because of the smaller sample size

Table 3: Effects of distance learning methods on evaluation

	France			Italy		
	Kindergarten	Primary	Secondary	Kindergarten	Primary	Secondary
	(1)	(2)	(3)	(4)	(5)	(6)
Lockdown	-5.035***	-3.156***	-4.189***	-8.079***	-5.212***	-5.563***
	(0.936)	(0.139)	(0.182)	(0.307)	(0.296)	(0.907)
Online Lectures	1.742	0.411*	0.633***	2.466***	1.107***	1.830**
· Lockdown	(1.130)	(0.212)	(0.207)	(0.348)	(0.250)	(0.844)
Homework	1.784*			1.878***		
· Lockdown	(0.938)			(0.342)		
Child fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,449	2,395	1,763	2,156	3,414	1,756
Within R-squared	0.687	0.666	0.750	0.910	0.804	0.782

All results were estimated using the fixed effects model on panel data from Italian and French 2020 COVID-19 online surveys. "Lockdown" is a dummy equal to one for the observations during the pandemic. "Online Lectures" is a dummy equal to one if the child had interactive distance lectures during the pandemic. "Homework" equals to one when the child had no interactive lectures (for instance, pedagogical content sent by emails). Column 1 is based on a sample of children at kindergarten level. As many of the pupils at this level did not receive any educational material from their teachers, we keep all individuals of our sample. The coefficient in front of the "Lockdown" variable in Column 1 is consequently the effect of lockdown on the parents' judgments of their child's improvement when the child had no lectures. "Online Lectures · Lockdown " is the differential effect when the child had interactive lectures. "Homework · Lockdown" is the differential effect when the child had no-interactive lectures. For other estimates presented in columns 2 to 4, we exclude pupils without lectures from our sample because of the slight size of this sub-sample. The coefficients in front of "Lockdown" variable show the effects of having no-interactive lectures as compared to the differential effect of having online lectures.
Individual controls are the time characteristics of the child, the time spent in front of passive screen, time spent reading or listening to stories. We also control for the mother and father's participation in the labor market.
Each estimation controls for child individual fixed effects. Standard Errors in parentheses are clustered at region level. ***, **, * indicate significance at the 1%, 5% and 10% level respectively.

Covid Economics 68, 10 February 2021: 69-99

exactly the same and may have very different outcomes for younger and older children. We estimate separately the model for Italy and for France because of the differences in the original questions on distance learning methodologies, which do not allow a perfect comparison across countries.

In Equation (2), θ captures the effect of having no educational material provided by the teachers during the pandemic, β_1 is the differential effect from θ of having interactive learning, β_2 is the differential effect from θ of having educational content without interaction. As we can observe in Figure 3, there were very few pupils at primary or upper levels who received no educational material during the school closure, we thus decided to drop these children from our sample. Therefore, in Equation (3) the coefficient θ in front of the variable LD_t captures the effect of having educational content without interaction during the lockdown and, β in front of the interactive term between LD_t and $OnlineL_i$, captures the differential effect from θ of having interactive lectures.

X_{it} is a vector of time varying controls, which include the work force participation of the parents and the time spent by children in front of screens or reading. All of the regressions present standard errors clustered at the regional level.

Table 3 reports estimates of children's learning processes for France and Italy, respectively. In order to make the interpretation of our results as simple as possible, we present these results in a graph in Figure 8 (a) and (b). For pupils enrolled in French kindergartens, we do not find significant differences between interactive and no interactive learning. At 95% confidence intervals, we also do not find any differences with no education continuity. However, we cannot exclude that this result may be driven

Covid Economics 68, 10 February 2021: 69-99

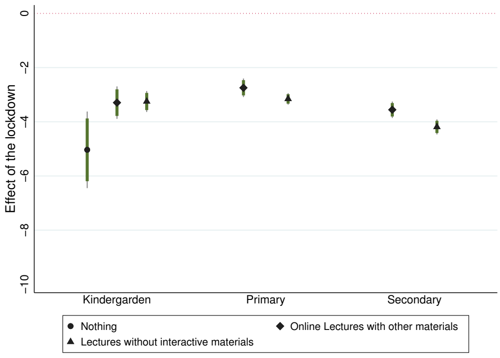

(a) Effects on educational progress [France]

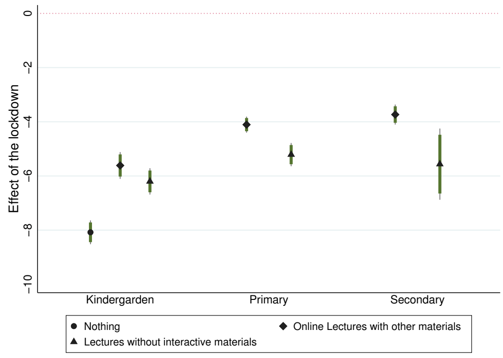

(b) Effects on educational progress [Italy]

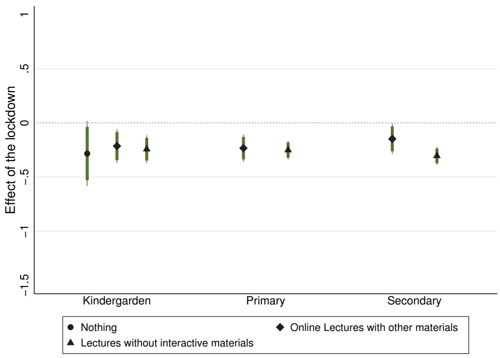

(c) Effects on emotional status [France]

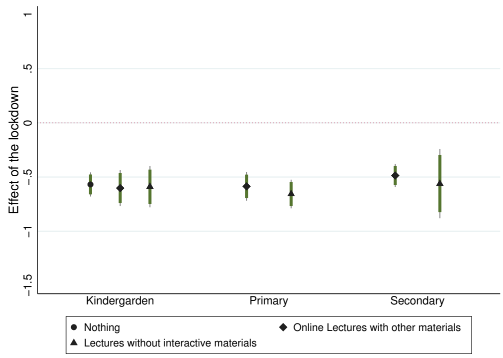

(d) Effects on emotional status [Italy]

Figure 8: Effects of the distance learning methods

by the weak size of our subsample of children with no education continuity. For Italian children at the same education level, overall we note a worse evaluation of children's progress than in France, as evidenced in the previous section. We recall that 40 per cent of Italian pre-school children did not receive any pedagogical continuity. This is reflected in an extremely low evaluation of education progress for those children, which is significantly lower with respect to those receiving interactive and non-interactive lectures. No significant differences are observed in pre-school children attending online classes in addition to receiving material.

For both primary and secondary levels, our estimates indicate that, according to parents' perceptions, interactive lectures are more advantageous for educational progress than non-interactive methods. The effect is again stronger for Italian children.[23] We can also remark that, although Italian parents are more worried about the learning progress of their primary school children with respect to French parents on average, the differences across the two countries become far smaller when Italian students attend online classes. This may suggest that young French students are likely more independent than the Italian ones and that Italian parents are reassured when their children have a closer contact with their

[23] Estimates for Italian secondary school children may not be very accurate because very few students in that group have only no-interactive lectures, and this is reflected in the large confidence interval, as we can observe in Figure 8 (c).

Table 4: Effects of distance learning methods on emotional status

	France			Italy		
	Kindergarten	Primary	Secondary	Kindergarten	Primary	Secondary
	(1)	(2)	(3)	(4)	(5)	(6)
Lockdown	-0.284	-0.252***	-0.304***	-0.569***	-0.657***	-0.562**
	(0.199)	(0.0568)	(0.0542)	(0.0779)	(0.0915)	(0.220)
Online Lectures	0.0694	0.0183	0.156**	-0.0335	0.0705	0.0757
· Lockdown	(0.217)	(0.0790)	(0.0669)	(0.0909)	(0.0587)	(0.201)
Homework	0.0405			-0.0205		
· Lockdown	(0.158)			(0.111)		
Child fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,449	2,395	1,763	2,156	3,414	1,756
Within R-squared	0.160	0.158	0.101	0.353	0.408	0.250

All results were estimated using the fixed effects model on panel data from Italian and French 2020 COVID-19 online surveys. "Lockdown" is a dummy equal to one for the observations during the pandemic. "Online Lectures" is a dummy equal to one if the child had interactive distance lectures during the pandemic. "Homework" equals to one when the child had no interactive lectures (for instance, pedagogical contents sent by emails). Column 1 is based on a sample of children at kindergarten level. As many of pupils at this level did not receive any educational material from their teachers, we keep all individuals of our sample. The coefficient in front of the "Lockdown" variable in Column 1 is consequently the effect of lockdown on the parents' judgment of their child's emotional status when the child had no lectures. "Online Lectures · Lockdown " is the differential effect when the child had interactive lectures. "Homework · Lockdown" is the differential effect when the child had no-interactive lectures. For other estimates presented in columns 2 to 4, we exclude pupils without lectures from our sample because of the slight size of this sub-sample. The coefficients in front of "Lockdown" variable show the effects of having no-interactive lectures as compared to the differential effect of having online lectures.
Individual controls are the time characteristics of the child, the time spent in front of a passive screen, time spent reading or listening to stories. We also control for the mother and father's participation in the labor market.
Each estimation controls for child individual fixed effects. Standard Errors in parentheses are clustered at region level. ***, **, * indicate significance at the 1%, 5% and 10% level respectively.

teachers. This difference could also be explained by the perspectives of re-opening the French schools in late May, which may have reassured French parents with respect to their children's learning progress even in the absence of interactive distance learning methods: while in France President Macron declared that schools would reopen in June, in early May, for Italian parents – in the absence of any statement on it by the Italian Prime Minister – it became self-evident that schools would not reopen before the summer vacation, with negative consequences on the expectations for their children's learning progress.

Regarding children's emotional status, results reported in Table 4 and Figure 8 (c) and (d) show that, on average, distance learning methods do not seem to play a crucial role in the psychological health of young children, whereas they seem relevant for secondary school students, at least in France. As for the Italian secondary school children, since only a few of them did not attend online lectures, estimations are not very accurate.[24]

5 Conclusions

School closures, forcibly caused by the COVID-19 crisis in many countries, impacted children's lives and their learning processes. There will likely be substantial and persistent disparities between families

[24]We also explored the possible heterogeneity of distance learning methods according to the gender of the pupils. Results show no differences between boys and girls in the learning evaluation outcome for kindergarten and primary levels. In secondary schools interactive distance learning seems to benefit girls more. As to the emotional status outcome, we do not observe any significant difference between girls and boys.

Covid Economics 68, 10 February 2021: 69-99

in terms of educational outcomes. This situation may also affect the socio-emotional skills of children, such as their mental health, well-being, and behavior. Distance learning solutions adopted by schools are heterogeneous across countries, within countries and between school levels. As a consequence, most of the burden of children's learning falls on their parents, with likely uneven results depending on the socio-economic characteristics of the family. Using a real-time survey data collected in April and early May 2020 in France and Italy on a large sample of families, we analyze how the Spring 2020 lockdown has affected children's emotional well-being and home learning processes according to their parents.

We show that both French and Italian parents were particularly worried by their children's home learning processes, with Italian parents giving lower evaluations than French parents to pre-primary and primary school levels. Using child fixed effect estimates, we find that that learning progress during the lockdown was particularly difficult for very young children (aged 3-6), especially for the ones who did not receive any distance learning support from their teachers, i.e. 40% of them in Italy vs only 2% in France. Children attending secondary schools also experienced important losses in terms of learning achievements when they could not attend online classes, and this is particularly evident in France, where it was the case for almost 30%. More generally, our regression results suggest that, for parents, attending online classes played a role in reducing the negative impact of the lockdown on the home learning process. This is true for all school levels except for kindergarten. As to the implementation of distance learning technologies, our data show important differences in the share of students that could attend online lectures both across countries and across school levels, with observed heterogeneity favoring Italian students and higher grades students.

For children's emotional wellbeing, our data indicate that, according to their parents, Italian children suffered more than French ones. In both countries, parents reported a worse emotional status for younger children. Much of the negative effect on emotional status may be due to their very limited interactions with peers. For older children, this reduction in personal interaction may have been partially compensated by virtual interaction, which could have mitigated the negative effect of the lockdown on their emotional status. On-line classes seem to have attenuated the social capital losses of secondary-school pupils during the lockdown. As for learning, we observe that university educated parents were less worried about their children's emotional status.

Finally, our results show that both French and Italian children increased the time spent reading by 0.3 hours on average, and the time spent in front of a screen (out of classes) by 1.3 hours on average during lockdown. The negative effects of the lockdown on both children's learning and emotional status is attenuated when the children spend more time reading, while it is amplified when they spend more hours watching TV or in passive screen activities (YouTube, social media, and similar).

All in all, our results indicate parents perceive that younger children to suffer more from the lockdown, both in terms of learning achievements and in terms of emotional stability. While older children could still interact with their friends using social networks, and could easily follow online classes, this was not possible, or at least very difficult, for children under 6. While interactive distance learning could help primary and secondary school children to learn the during school closures, it can be extremely hard for teachers to communicate remotely with very young children. Moreover, parents may find that it is more difficult to help their pre-school age children in their learning processes as the teaching methods for this age group are less standardized and demand more creative skills. These conclusions suggest that governments should be particularly concerned about keeping young pupils at school for as long as

Covid Economics 68, 10 February 2021: 69-99

possible during the health crisis, while at the same time providing kindergarten teachers with training that could help them to ensure the continuity of learning if schools were forced to close. As shown by our results, existing technologies for online classes were perceived by parents as being quite effective for the learning processes of older children. Governments should strongly recommend that teachers use such technologies and facilitate this task by providing clear guidelines and training materials to be used in the event of school closures.

A limitation of this work is that answers to the questionnaire were provided by parents and therefore may be substantially different from the responses on time use or wellbeing that the children would have given if surveyed directly. This may be especially true for parents of older children, who may find it difficult to distinguish the time spent by children doing homework from passive screen time. Also, parents' perceptions on learning progress may differ from the children's actual outcomes, which could only be assessed through cognitive tests. Nevertheless, we believe that being aware of parents' worry about their children's emotional status and learning progress may be useful for researchers and policy makers as a wake-up call for the closer monitoring of child development during and beyond the current pandemic crisis.

Covid Economics 68, 10 February 2021: 69-99

References

Andrew, A., S. Cattan, M. Costa-Dias, C. Farquharson, L. Kraftman, S. Krutikova, A. Phimister, and A. Sevilla (2020). Learning during the lockdown: real-time data on children's experiences during home learning. IFS Briefing Note BN288. 1, 2, 4, 8, 9

Anxo, D., L. Mencarini, A. Pailhé, A. Solaz, M. L. Tanturri, and L. Flood (2011). Gender differences in time use over the life course in France, Italy, Sweden, and the US. *Feminist Economics 17*(3), 159–195. 2

Attanasio, O., R. Blundell, G. Conti, and G. Mason (2020). Inequality in socio-emotional skills: A cross-cohort comparison. *Journal of Public Economics*. 104171, forthcoming. 12

Bertrand, M. and J. Pan (2013). The trouble with boys: Social influences and the gender gap in disruptive behavior. *American Economic Journal: Applied Economics 5*(1), 32–64. 12

Bettinger, E. P., L. Fox, S. Loeb, and E. S. Taylor (2017). Virtual classrooms: How online college courses affect student success. *American Economic Review 107*(9), 2855–75. 2

Burgess, S. and H. H. Sievertsen (2020). Schools, skills, and learning: The impact of COVID-19 on education. *VoxEu.org 1*. 1

Cardoso, A. R., E. Fontainha, and C. Monfardini (2010). Children's and parents' time use: empirical evidence on investment in human capital in France, Germany and Italy. *Review of Economics of the Household 8*(4), 479–504. 5

Coates, D., B. R. Humphreys, J. Kane, and M. A. Vachris (2004). "No significant distance" between face-to-face and online instruction: Evidence from principles of economics. *Economics of Education Review 23*(5), 533–546. 2

Del Boca, D., C. Monfardini, and C. Nicoletti (2017). Parental and child time investments and the cognitive development of adolescents. *Journal of Labor Economics 35*(2), 565–608. 2, 8

Del Boca, D., N. Oggero, P. Profeta, and M. Rossi (2020). Women's Work, Housework and Childcare, before and during COVID-19. *Review of Economics of the Household*. Forthcoming. 8

DEPP (2020). Continuité pédagogique période de mars à mai 2020. Document de travail n° 2020-E03. 6, 10

Farré, L., Y. Fawaz, L. González, and J. Graves (2020). How the COVID-19 Lockdown Affected Gender Inequality in Paid and Unpaid Work in Spain. IZA Discussion Paper No. 13434. 8

Figlio, D., M. Rush, and L. Yin (2013). Is it live or is it internet? Experimental estimates of the effects of online instruction on student learning. *Journal of Labor Economics 31*(4), 763–784. 3

Fiorini, M. and M. P. Keane (2014). How the allocation of children's time affects cognitive and noncognitive development. *Journal of Labor Economics 32*(4), 787–836. 2

Covid Economics 68, 10 February 2021: 69-99

Giménez-Nadal, J. I., L. Mangiavacchi, and L. Piccoli (2019). Keeping inequality at home: The genesis of gender roles in housework. *Labour Economics 58*, 52–68. 8

Jaume, D. and A. Willén (2019). The long-run effects of teacher strikes: evidence from Argentina. *Journal of Labor Economics 37*(4), 1097–1139. 2

Kalb, G. and J. C. Van Ours (2014). Reading to young children: A head-start in life? *Economics of Education Review 40*, 1–24. 2, 9, 14

Lavy, V. (2015). Do differences in schools' instruction time explain international achievement gaps? evidence from developed and developing countries. *The Economic Journal 125*(588), F397–F424. 2, 8

Mangiavacchi, L., L. Piccoli, and L. Pieroni (2020). Fathers matter: Intra-household responsibilities and Children's wellbeing during the COVID-19 lockdown in Italy. *IZA Discussion Paper No. 13519*. 1, 8

Moroni, G., C. Nicoletti, and E. Tominey (2019). Child socio-emotional skills: The role of parental inputs. IZA Discussion Paper No. 12432. 12

Pailhé, A., A. Solaz, and M. L. Tanturri (2019). The time cost of raising children in different fertility contexts: Evidence from France and Italy. *European Journal of Population 35*(2), 223–261. 2, 5

Pellizzari, M., M. P. Cacault, J. Laurent-Lucchetti, and C. Hildebrand (2019). Distance learning in higher education: Evidence from a randomized experiment. IZA Discussion Paper No. 12298. 2

Walsh, J. J., J. D. Barnes, J. D. Cameron, G. S. Goldfield, J.-P. Chaput, K. E. Gunnell, A.-A. Ledoux, R. L. Zemek, and M. S. Tremblay (2018). Associations between 24 hour movement behaviours and global cognition in US children: a cross-sectional observational study. *The Lancet Child & Adolescent Health 2*(11), 783–791. 2

Woessmann, L. (2016). The importance of school systems: Evidence from international differences in student achievement. *Journal of Economic Perspectives 30*(3), 3–32. 2, 3

Xu, D. and S. S. Jaggars (2013). The impact of online learning on students' course outcomes: Evidence from a large community and technical college system. *Economics of Education Review 37*, 46–57. 2

Covid Economics 68, 10 February 2021: 69-99

A Supplementary material

Covid Economics 68, 10 February 2021: 69-99

Table A1: Institutional comparison of the French and Italian educational systems.

	France	Italy
Age of attendance		
Kindergarten	3-5	3-5
Primary	6-10	6-10
Lower secondary	11-14	11-13
Higher secondary	15-18	14-18
School days per year		
Primary	162	200
Lower secondary	162	200
Higher secondary	180	200
Summer vacation		
weeks per year	8	12/13
Class size		
Primary	23.7	19.1
Secondary	25.2	21
Pupils per teacher		
Kindergarten	23.3	12.2
Primary	19.2	11.5
Lower secondary	14.4	11
Higher secondary	11.4	10.4
Attendance rate (% of the same age group)		
Nursery	56.3%	29.7%
Kindergarten	100.0%	93.9%
Primary	99.7%	97.4%
Secondary	86.4%	84.8%
Students enrolled in private institutions		
Kindergarten	13.3%	28.3%
Primary	14.9%	6.0%
Lower secondary	22.1%	3.6%
Higher secondary	29.0%	8.8%
Public expenditure per pupil (thousand US$ PPP)		
Kindergarten	8.2	7.4
Primary	7.6	8.0
Lower secondary	10.6	8.9
Higher secondary	14.1	9.4
Public expenditure		
Share of total public expenditure	10.8%	8.9%
Percentage of the GDP	3.7%	2.5%
Starting salary of teachers (thounsand US$ PPP)		
Kindergarten	30.9	30.4
Primary	30.9	30.4
Lower secondary	32.5	32.7
Higher secondary	32.5	32.7
Share of female teachers		
Kindergarten	89.4%	98.7%
Primary	83.5%	93.6%
Lower secondary	60.5%	76.7%
Higher secondary	59.8%	66.2%
Distribution of primary school teachers by age class		
Less than 30	12%	1%
30-39	33%	11%
40-49	34%	32%
50 or more	22%	56%
PISA scores		
Reading	493	476
Math	495	487
Science	493	468

Source: OECD.stat, Eurydice, PISA-OECD (last available year, most figures refer to 2017 or 2018)

Covid Economics 68, 10 February 2021: 69-99

Table A2: Representativeness of the sample.

A) ITALY			B) FRANCE		
	Our sample	ISTAT		Our sample	INSEE
Family type[1,2]					
Couples with 1 child	44.0%	47.9%	Couples with 1 child	31.2%	44.8%
Couples with 2 children	46.8%	41.7%	Couples with 2 children	50.7%	38.7%
Couples with 3 or more children	9.1%	10.4%	Couples with 3 children	15.4%	12.7%
			Couples with 4 or more children	2.6%	3.8%
Single parents	7.8%	24.5%	Single parents	14.4%	22.8%
Geographical distribution[3,4]					
Piemonte	9.6%	6.6%	AURA	12.4%	27.9%
Valle d'Aosta	0.4%	0.2%	Bourgogne-Franche-Comté	4.3%	6.8%
Liguria	2.6%	2.2%	Bretagne	5.2%	4.5%
Lombardia	20.0%	16.1%	Centre-Val-de-Loire	3.9%	5.0%
Trentino	2.2%	1.7%	Corse	0.5%	0.2%
Veneto	9.2%	8.0%	Grand Est	8.5%	8.0%
Friuli	2.3%	1.9%	Hauts-de-France	9.2%	6.3%
Emilia-Romagna	8.6%	6.9%	Île-de-France	18.9%	5.9%
Toscana	10.5%	6.0%	Normandie	5.1%	5.4%
Umbria	3.5%	1.4%	Nouvelle-Aquitaine	9.3%	7.6%
Marche	2.6%	2.6%	Occitanie	9.1%	9.9%
Lazio	9.1%	10.3%	Pays de la Loire	5.9%	4.4%
Abruzzo	1.5%	2.2%	PACA	7.8%	8.2%
Molise	0.4%	0.5%			
Campania	5.0%	10.6%			
Puglia	4.7%	7.1%			
Basilicata	0.5%	0.9%			
Calabria	1.3%	3.3%			
Sicilia	3.2%	8.8%			
Sardegna	2.8%	2.8%			

Notes: 1. ISTAT – Multipurpose Survey on Households: Aspects of Daily Life 2019. 2. INSEE – Census 2016 3. ISTAT – Resident Municipal Population on January 1 2019. 4. INSEE – Census 2016

Covid Economics 68, 10 February 2021: 69-99

CEPR PRESS

Covid Economics 68, 10 February 2021: 69-99

Table A3: Full estimates by sub-samples (FE regressions)- learning evaluation

	(1) Girls	(2) Boys	(3) Kindergarden	(4) Primary school	(5) Lower Secondary school	(6) Upper secondary school	(7) Both parents with university education	(8) Only mother with university education	(9) Only father with university education	(10) Both parents without university education	(11) No parents at home	(12) Only father at home	(13) Only mother at home	(14) Both parents at home	(15) Single child	(16) Child with siblings	(17) Single parent household
Lockdown	-4.705***	-5.193***	-6.700***	-4.456***	-3.942***	-3.465***	-4.994***	-4.980***	-4.471***	-4.958***	-5.156***	-4.893***	-4.756***	-4.990***	-5.102***	-4.909***	-5.208***
	(0.162)	(0.182)	(0.213)	(0.174)	(0.197)	(0.411)	(0.297)	(0.262)	(0.658)	(0.152)	(0.262)	(0.585)	(0.231)	(0.205)	(0.222)	(0.176)	(0.259)
FrancexLockdown	1.620***	1.779***	3.338***	1.466***	0.365	-0.362	2.022***	1.854***	1.229**	1.359***	1.589***	1.114**	1.702***	1.830***	1.861***	1.636***	1.319***
	(0.142)	(0.180)	(0.186)	(0.141)	(0.220)	(0.373)	(0.279)	(0.167)	(0.537)	(0.172)	(0.302)	(0.414)	(0.171)	(0.205)	(0.217)	(0.150)	(0.269)
Time spent in front of a screen	-0.178**	-0.071	-0.206**	-0.142**	-0.082	-0.050	-0.186***	-0.174*	-0.156	-0.055	-0.091	-0.112	-0.211**	-0.072	-0.092	-0.130**	-0.060
	(0.075)	(0.061)	(0.078)	(0.062)	(0.065)	(0.098)	(0.067)	(0.092)	(0.145)	(0.057)	(0.100)	(0.182)	(0.090)	(0.082)	(0.068)	(0.055)	(0.105)
Time spent reading	0.151	0.263**	0.273*	0.348***	0.178	0.204	0.205*	0.311**	-0.086	0.223	0.237	0.308	0.193	0.178	0.056	0.244**	0.313
	(0.113)	(0.098)	(0.138)	(0.098)	(0.122)	(0.190)	(0.112)	(0.148)	(0.197)	(0.140)	(0.261)	(0.287)	(0.116)	(0.108)	(0.136)	(0.100)	(0.227)
Mother is working	-0.028	0.075	-0.085	-0.086	-0.037	0.412	-0.003	-0.093	0.521	0.015	-0.153	-0.075	0.262	-0.049	-0.104	0.048	
	(0.187)	(0.184)	(0.279)	(0.188)	(0.203)	(0.342)	(0.230)	(0.261)	(0.651)	(0.167)	(0.725)	(0.658)	(0.257)	(0.209)	(0.321)	(0.157)	
Father is working	0.081	0.078	-0.022	0.104	0.099	0.407	0.093	0.138	-0.217	0.021	0.605	0.081	0.037	0.110	0.392	-0.043	
	(0.181)	(0.163)	(0.196)	(0.153)	(0.203)	(0.327)	(0.463)	(0.232)	(0.655)	(0.151)	(0.526)	(0.339)	(0.416)	(0.179)	(0.325)	(0.136)	
Parent is working																	-0.091
																	(0.410)
Observations	6,584	6,855	4,080	5,831	2,528	1,000	4,310	3,588	798	4,683	1,616	1,370	4,490	5,943	2,930	10,509	1,443
Within R-squared	0.753	0.777	0.851	0.758	0.766	0.760	0.757	0.777	0.737	0.773	0.800	0.769	0.765	0.756	0.788	0.759	0.772

Robust standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Covid Economics 68, 10 February 2021: 69-99

Table A4: Full estimates by sub-samples (FE regressions)- emotional status

VARIABLES	(1) Girls	(2) Boys	(3) Kindergarden	(4) Primary school	(5) Lower Secondary school	(6) Upper secondary school	(7) Both parents with university education	(8) Only mother with university education	(9) Only father with university education	(10) Both parents without university education	(11) No parents at home	(12) Only father at home	(13) Only mother at home	(14) Both parents at home	(15) Single child	(16) Child with siblings	(17) Single parent household
Lockdown	-0.505***	-0.625***	-0.565***	-0.596***	-0.510***	-0.405***	-0.500***	-0.567***	-0.632***	-0.623***	-0.737***	-0.497***	-0.534***	-0.524***	-0.612***	-0.553***	-0.668***
	(0.056)	(0.057)	(0.083)	(0.077)	(0.077)	(0.139)	(0.097)	(0.095)	(0.158)	(0.068)	(0.113)	(0.130)	(0.083)	(0.099)	(0.075)	(0.055)	(0.136)
FrancexLockdown	0.250***	0.352***	0.278***	0.307***	0.284***	0.168	0.329***	0.241***	0.389***	0.308***	0.423***	0.207	0.259***	0.304***	0.377***	0.279***	0.348***
	(0.046)	(0.047)	(0.071)	(0.057)	(0.065)	(0.136)	(0.072)	(0.077)	(0.141)	(0.068)	(0.129)	(0.146)	(0.064)	(0.082)	(0.094)	(0.040)	(0.125)
Time spent in front of a screen	-0.082***	-0.059**	-0.094***	-0.088***	-0.031	-0.022	-0.084**	-0.066	-0.027	-0.064**	-0.044	-0.089	-0.086***	-0.062*	-0.052	-0.075***	-0.017
	(0.024)	(0.024)	(0.031)	(0.026)	(0.032)	(0.057)	(0.036)	(0.047)	(0.050)	(0.027)	(0.034)	(0.059)	(0.025)	(0.036)	(0.031)	(0.021)	(0.045)
Time spent reading	0.108**	0.120***	0.094**	0.141***	0.110	0.103	0.075	0.148***	0.017	0.141***	0.119*	0.246***	0.096*	0.090**	0.080	0.122***	0.034
	(0.043)	(0.038)	(0.046)	(0.032)	(0.069)	(0.068)	(0.052)	(0.050)	(0.087)	(0.037)	(0.070)	(0.068)	(0.052)	(0.037)	(0.048)	(0.033)	(0.082)
Mother is working	0.078	-0.071	0.019	-0.002	-0.059	0.016	0.011	-0.133	0.087	0.046	-0.109	0.028	0.064	-0.026	0.017	-0.012	
	(0.068)	(0.082)	(0.120)	(0.065)	(0.098)	(0.131)	(0.101)	(0.100)	(0.136)	(0.068)	(0.400)	(0.332)	(0.077)	(0.085)	(0.096)	(0.059)	
Father is working	0.044	0.047	-0.018	0.024	0.119	0.214	-0.018	0.057	0.032	0.025	0.188	0.083	0.057	0.076	0.148*	0.004	
	(0.070)	(0.063)	(0.074)	(0.064)	(0.096)	(0.153)	(0.159)	(0.123)	(0.134)	(0.074)	(0.340)	(0.138)	(0.151)	(0.064)	(0.084)	(0.066)	
Parent is working																	0.015
																	(0.142)
Observations	6,584	6,855	4,080	5,831	2,528	1,000	4,310	3,588	798	4,683	1,616	1,370	4,490	5,943	2,930	10,509	1,443
Within R-squared	0.2795	0.285	0.298	0.336	0.196	0.167	0.226	0.280	0.291	0.328	0.367	0.285	0.311	0.240	0.329	0.268	0.212

Robust standard errors in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

A vaccine auction[1]

Romans Pancs[2]

Date submitted: 21 January 2021; Date accepted: 31 January 2021

This article describes an auction for selling vaccines in a pandemic. The environment borrows from the problem of allocating positions for sponsored links on pages with online search results but recognizes the externalities that one man's vaccination imposes on another. The auction is the pivot Vickrey-Clark-Groves mechanism and, so, inherits its properties: efficiency and strategy-proofness. The auction lets each bidder bid not only on his own behalf but also on behalf of others. The auction requires neither the bidders nor the auctioneer to forecast the efficacy of the vaccine or the evolution of the pandemic.

1 I thank Shaun McRae for many a fruitful conversation and credit him for the best idea in this article: that everyone be able to bid on everyone's vaccination, not just on his own.

2 ITAM.

CEPR PRESS

1 Introduction

Suppose the government has bought forward vaccines from pharmaceutical companies. How should the government allocate these vaccines?

When triaging, doctors ask: How to allocate scarce resources to maximize the number of quality adjusted life years saved? When deciding how much pollution to allow, the Environmental Protection Agency asks: How to balance the number of lives saved (each valued at about ten million dollars) against the abatement costs? Imitating either approach in order to allocate vaccines in a pandemic would be unwise for two reasons:

1. No one apart from the individual himself knows how much he values the life-style of greater freedom afforded by the vaccine relative to the alternative of the social isolation. (Steven Landsburg makes this point here.[1])

2. There are externalities. Prioritizing the protection of the most vulnerable may do little to slow down the spread of the virus. (Tyler Cowen makes this point here.[2]) When infected, individuals impose externalities on their colleagues, employers, and their health insurer, private or public (such as the NHS in the United Kingdom).

A superior approach, described in this note, is an auction with the following defining features:

- Each individual and each organization (e.g., a government agency, a charity, or a firm) can bid for a vaccine on behalf of any individual (including the bidder himself). While individuals (physical entities) get vaccinated, organizations do not.

- If getting vaccinated is more desirable earlier in the pandemic rather than later, then vaccines are allocated to the individuals in the descending order of the aggregate bids submitted on their behalf. That is, the individual with the greatest support as expressed by the sum of the bids submitted on his behalf gets the vaccine first; the individual with the second greatest support gets the vaccine second, and so on.

- Each bidder pays the externality that his bids impose on others by diverting the vaccine towards (or away from) those on whose behalf he bids. For instance, if a coffee shop's bid

[1]Steven Landsburg, "Priority Care," *The Big Questions* blog, 2 December 2020.
[2]Tyler Cowen, "Vaccine Distribution Shouldn't Be Fair," *Bloomberg Opinion*, 23 November 2020.

Covid Economics 68, 10 February 2021: 100-114

on behalf of its employee Alice causes Alice and Bob to switch their adjacent positions in the vaccination order, then the coffee shop pays the cumulative cost of delayed vaccination that this switch imposes on Bob (as revealed by Bob's bid on himself), on Bob's family (as revealed by their bids on Bob's behalf), on his health insurer, and on anyone else who may have bid on Bob's behalf. Thus, the auction is not pay-as-bid.

Because the described vaccine auction is a pivot VCG mechanism, it inherits the virtues shared by the auctions in the class:[3]

- The auction allocates the vaccine efficiently.
- The auction is strategy-proof: regardless of what others do, no bidder can do better than to bid his true valuations.

Here is how the auction mitigates the two major deficiencies of the triage-like vaccine distribution approach mentioned above:

1. The auction enables each bidder to express his private valuation for the vaccine. Examples:
   (a) A vaccine skeptic may express a negative valuation.
   (b) A nurse with virus-conferred immunity may express a zero valuation.
   (c) A youngster who looks after his ailing grandparents may express a large positive valuation.
2. The opportunity to bid on behalf of others enables individuals and organizations to internalize some of the externalities that vaccination entails. Examples:
   (a) A firm at the helm of a company town may subsidize the vaccination of (i.e., may bid on behalf of) its employees in order to accelerate herd immunity and avert costly lockdowns.
   (b) A susceptible individual who is allergic to the vaccine may subsidize the vaccination of his doorman, his housekeeper, and his hairdresser.

[3] Ausubel and Milgrom (2005) discuss the virtues and the weaknesses of VCG (Vickrey–Clarke–Groves) mechanisms. The two bullet points below is their Theorem 1. Their Theorem 2 says that the described vaccine auction (and, in particular, its feature that bidders be allowed to bid on behalf of others) is essentially necessary under rather weak conditions provided one insists on the two bullet points.

Covid Economics 68, 10 February 2021: 100-114

(c) A health insurance company (private or public) may subsidize the vaccination of the most vulnerable among the insured (e.g., the elderly, the obese, and those suffering from chronic conditions) in order to avoid paying their hospitalization bills.

(d) An airline or a coffee shop chain may subsidize the vaccination of its loyal customers—and potential super-spreaders—in order to mitigate the treat they pose to other customers and to the company's reputation.

The described vaccine auction generates revenue. In order to preserve bidders' incentives, it is important that this revenue not be promised to any bidder. In particular, the budget of the public healthcare provider (e.g., the NHS), who is a bidder in the auction, may not be topped up depending on the auction's realized revenue. Instead, the government can commit to channeling the auction revenue towards retiring the national debt.

Related Literature

The problem of vaccine allocation is closely related to the thoroughly studied and successfully solved problem of selling ad positions to advertisers in online search. The ads are sold via auctions, popularized by *Google* and *Yahoo!*. The problem of selling online ads consists in allocating positions of sponsored links on a search page. The vaccine allocation problem consists in allocating positions in the vaccination queue. The former problem has been formalized by Edelman, Ostrovsky and Schwartz (2007) and Varian (2007), on whose formalization I build.

As far as allocating vaccines is concerned, the only difference from the problem of selling sponsored search positions is that a variety of actors might care about uncle Bob's vaccination, not just uncle Bob himself. The proposed vaccine auction accommodates this feature.

2 The Model

The model extends the positions auction environment of Edelman, Ostrovsky and Schwartz (2007) and Varian (2007) by permitting each bidder to bid not only on his behalf but also on behalf of others. The proposed auction is the standard pivot VCG mechanism.

Covid Economics 68, 10 February 2021: 100-114

CEPR PRESS

Environment

In total I bidders participate in a vaccine auction. Typical bidders are indexed by i and j in $\mathcal{I} \equiv \{1, 2, \ldots, I\}$. A bidder is interpreted as an individual or an organization.

Once a vaccine has been authorized for use, T units of the vaccine become available sequentially over time. (In practice, a "unit" comprises the first dose of the vaccine and any necessary boosters.) We say that the t-th unit becomes available at time $t \in \mathcal{T} \equiv \{1, 2, \ldots, T\}$. Without loss of generality, let $T = I$. The (economic) time in the model need not correspond to the calendar time, as will be explained.

Bidders' common discount factor for getting vaccinated at time t is denoted by $\alpha_t \in \mathbb{R}_+$, which the auctioneer knows (more on which later). Each discount factor α_t is interpreted as the time-discounted reduction in the average mortality risk for those vaccinated at time t. Without loss of generality, $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_T$. Any variation in the speed with which vaccine units become available is folded into the sequence $(\alpha_t)_{t \in \mathcal{T}}$; as a result, the calendar time will generally differ from the (economic) time in the model. For instance, a batch of k vaccine units that appear from time t onwards in short succession are all discounted similarly; $\alpha_t, \alpha_{t+1}, \ldots, \alpha_{t+k-1}$ are all about the same. If the subsequent unit is shipped with much delay, then α_{t+k} is much lower than α_{t+k-1}, both because of impatience and because the pandemic continues to ravage. When herd immunity is reached, the sequence $(\alpha_t)_{t \in \mathcal{T}}$ falls precipitously. Moreover, if early vaccination is perceived as risky, then, all other things being equal, individuals may prefer to wait and see before getting vaccinated, in which case the convention $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_T$ would require that time $t = 2$ occur "before" time $t = 1$, where "before" is in terms of calendar time, while $t = 2$ and $t = 1$ refer to the (economic) time in the model.

Each bidder i values bidder j's vaccination at time t at $\alpha_t s_{ij}$, where $s_{ij} \in \mathbb{R}$. Here, s_{ii} pertains to bidder i's benefit from vaccinating himself and captures the bidder-specific adjustments to the mortality risk, the cost of mitigating exposure to the virus, and the value of life; and s_{ij} with $j \neq i$ pertains to the bidder's benefit from seeing someone else vaccinated (a relative, an employee, or an insuree). Each bidder i knows his valuations $\left(s_{ij}\right)_{j \in \mathcal{I}}$. A vaccine proponent has $s_{ii} \geq 0$, while a vaccine skeptic has $s_{ii} < 0$. A spiteful vaccine proponent has $s_{ij} < 0$ for some $j \neq i$.

Covid Economics 68, 10 February 2021: 100-114

Define an allocation $\mathbf{x} \equiv (x_{it})_{i \in \mathcal{I}, t \in \mathcal{T}}$ by letting $x_{it} = 1$ if bidder i is vaccinated at time t and letting $x_{it} = 0$ otherwise. An allocation $\mathbf{x}$ is feasible if no two bidders are vaccinated at the same time: $\sum_{i \in \mathcal{I}} x_{it} \leq 1$ for all $t \in \mathcal{T}$. Let X denote the set of all feasible allocations.

At an allocation $\mathbf{x}$ in X, the payoff of a bidder i who makes a payment p_i is

$$\sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{I}} x_{jt} \alpha_t s_{ij} - p_i.$$

Utility is transferable across bidders and the auctioneer, who collect $\sum_{i \in \mathcal{I}} p_i$.

The Vaccine Auction

Each bidder i submits a collection $\mathbf{b}_i \equiv (b_{ij})_{j \in \mathcal{I}}$ of bids: a bid b_{ii} on his own behalf and a bid b_{ij} on behalf of each bidder $j \in \mathcal{I} \backslash \{i\}$. The aggregate bid on bidder i's behalf is denoted by $B_i \equiv \sum_{j \in \mathcal{I}} b_{ji}$.

The auction's allocation rule $\mathbf{x}^*$ is efficient; that is, it associates with each bid profile $\mathbf{b} \equiv (\mathbf{b}_i)_{i \in \mathcal{I}}$ a feasible allocation $\mathbf{x}^* (\mathbf{b})$ that maximizes the total surplus:

$$\mathbf{x}^* (\mathbf{b}) \in \arg\max_{\mathbf{x} \in X} \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{I}} x_{it} \alpha_t B_i. \tag{1}$$

Because the sequence $(\alpha_t)_{t \in \mathcal{I}}$ is weakly decreasing by assumption, one solves (1) by assigning each time-t unit of the vaccine to the bidder with the t-th highest aggregate bid, with indifferences resolved arbitrarily. That is, for any bidder i, we have $x^*_{it} (\mathbf{b}) = 1$ if and only if B_i is the t-th largest component of $\mathbf{B} \equiv (B_i)_{i \in \mathcal{I}}$.

Each bidder's payment in the auction is the externality that he imposes on other bidders assuming that all bidders bid truthfully, that is, assuming that $b_{ij} = s_{ij}$ for each i and j in $\mathcal{I}$. Formally, given the submitted bids $\mathbf{b} \equiv (\mathbf{b}_i, \mathbf{b}_{-i})$ with $\mathbf{b}_{-i} \equiv (\mathbf{b}_j)_{j \in \mathcal{I} \backslash \{i\}}$ and the induced aggregate bids $\mathbf{B}$, bidder i's payment is

$$p^*_i (\mathbf{b}) = \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{I}} x^*_{jt} (0, \mathbf{b}_{-i}) \alpha_t \left(B_j - b_{ij}\right) - \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{I}} x^*_{jt} (\mathbf{b}) \alpha_t \left(B_j - b_{ij}\right). \tag{2}$$

If all bidders bid truthfully, then the first term in (2) is the total surplus all but bidder i enjoy once the vaccines have been allocated efficiently while ignoring bidder i's bids, whereas the second

Covid Economics 68, 10 February 2021: 100-114

term is the total surplus all but bidder i enjoy once the vaccines have been allocated without ignoring bidder i's bids. By construction, each payment (2) is nonnegative, and, therefore, the auction cannot possibly lose money.

An Equilibrium of the Vaccine Auction

Denote the payment function profile in the auction by $\mathbf{p}^* \equiv (p_i^*)_{i\in\mathcal{I}}$, where each function p_i^* is given in (2). Recall that the auction's allocation rule $\mathbf{x}^*$ is defined in (1). Proposition 1 shows that the game induced by the vaccine auction has a dominant-strategy equilibrium in which each bidder's every bid (on his own behalf and on behalf of others) equals his corresponding valuation.

Proposition 1. *The vaccine auction* $(\mathbf{x}^*, \mathbf{p}^*)$ *is efficient and strategy-proof: irrespective of what other bidders do, each bidder i cannot do better than bid truthfully by setting* $b_{ij} = s_{ij}$ *for each* $j \in \mathcal{I}$. *Moreover, if a bidder i bids truthfully and* $s_{ij} \geq 0$ *for all* $j \in \mathcal{I}$, *then his payoff is nonnegative.*

Proof. The proposition's conclusion lists standard properties of the VCG pivot mechanism, of which the described vaccine auction is an instantiation. A direct proof is in Appendix A. ■

One may intuitively wonder about "free-riding" on others' bids: Why would Alice ever bid her true valuation for getting vaccinated if she knows that her health insurer is prepared to bid tenfold on her behalf? The answer is that free-riding is built into the structure of the auction. Alice does not pay what she bids. She pays the externality that her bids impose. If her bids do not affect the allocation of the vaccine, and the bids of her heath insurer and others are doing all the work, then she pays nothing.

3 Some Practical Considerations

Discount Factors

The efficiency of vaccine allocation and the optimality of truthful bidding rely on the accuracy of the discount factors $(\alpha_t)_{t\in\mathcal{I}}$, which enter both (1) and (2). For efficiency, it suffices to know that the value of getting vaccinated decreases in calendar time, which is a reasonable a priori assumption. Indeed, a vaccine is likely to make the greatest difference as soon as it has been authorized for use,

Covid Economics 68, 10 February 2021: 100-114

and less so as more and more individuals get vaccinated or acquire immunity through infection. In this case, it is efficient to vaccinate individuals in the descending order of the aggregate bids submitted on their behalf; the knowledge of the exact values of $(\alpha_t)_{t\in\mathcal{I}}$ is not required.

The exact values of $(\alpha_t)_{t\in\mathcal{I}}$ may not be available until after the pandemic. That's OK. The computation and the assessment of the payments can be delayed until after the pandemic without compromising truthful bidding at the beginning of the pandemic. In the context of positions auctions, Varian (2009) makes this point. Recall that each α_t is an estimate of the reduction in the average mortality risk for those who get vaccinated at time t. To determine $(\alpha_t)_{t\in\mathcal{I}}$ ex-ante, before the bids have been solicited, would require sophisticated epidemiological modeling and the forecasting of vaccine efficacy, infection mortality, and any medical treatments that could emerge on the way, as well as ample politico-economic guesswork. By contrast, $(\alpha_t)_{t\in\mathcal{I}}$ can be computed mechanically ex-post, once the evolution of the pandemic has been observed. In addition, the payments computed and assessed ex-post are "fair" in the sense that they charge bidders little for a vaccine that has proved inefficacious or has arrived too late to make a difference.

Vaccine Skeptics and Spiteful Individuals

The vaccine auction assumes that everyone participates. The auction guarantees a nonnegative payoff for unspiteful vaccine proponents (the "moreover" part of Proposition 1); such bidders are eager to participate. By contrast, vaccine skeptics and spiteful individuals may end up with negative payoffs.

To illustrate, suppose that there are two bidders (Alice and Bob) and two units of the vaccine with $\alpha_1 = \alpha_2 = 1$. Alice is a vaccine skeptic ($s_{11} = -1$) and does not care about Bob ($s_{12} = 0$). Bob seeks protection ($s_{22} = s_{21} = 2$). The auctioneer vaccinates both Alice and Bob, collects the payment of 1 from Bob and nothing from Alice. Alice's payoff is negative (-1). Even if Alice's objection to getting vaccinated exceeded Bob's benefit (i.e., if $s_{11} + s_{21} < 0$), her payoff would still be negative because she would pay the externality that her skipping the vaccine would impose on Bob. Alice's participation in the auction can be ensured either by coercion (vaccine passports would do the job) or by a subsidy (which may depend on Alice's observable characteristics but not on her bid).

Covid Economics 68, 10 February 2021: 100-114

CEPR PRESS

A bidder's payoff can be negative even if no bidder is a skeptic as long as some bidder is spiteful. To illustrate, modify the example above by assuming that $s_{11} = 2$, $s_{12} = 0$, $s_{22} = 0$, and $s_{21} = -1$ (i.e., Bob is spiteful). The auctioneer vaccinates both Alice and Bob, collects the payment of 1 from Alice and nothing from Bob. Bob's payoff is negative (-1). Even if Bob's objection to Alice's vaccination exceeded her benefit (i.e., if $s_{21} + s_{11} < 0$), his payoff would still be negative because he would pay the externality that his blocking of Alice's vaccination would impose on her. Bob's negative payoff would not discourage him from participating in the auction, for nonparticipation does not mute his hurt when others receive the vaccine.

Both examples illustrate that the auction does not discriminate between "moral" and "immoral" preferences, and takes full advantage of the model's assumption that interpersonal grievances are comparable and transferable.

Just as vaccine skepticism can have an "innocent" explanation (an allergy), so can spitefulness: a hospital may have a financial interests in its patients' contracting the infection thanks to a reimbursement from the insurer. To illustrate, suppose that there are three bidders (Alice, Hospital, and Insurer) and one unit of the vaccine, with $\alpha_1 = 1$ and $\alpha_2 = \alpha_3 = 0$. Alice is a vaccine proponent ($s_{11} = 1$). Hospital profits from treating unvaccinated Alice ($s_{21} = -2$). Insurer must reimburse Hospital and experiences some overhead ($s_{31} = 3$). The remaining valuations are zero. In equilibrium, Alice is vaccinated because $s_{11} + s_{21} + s_{31} > 0$. Alice and Hospital pay nothing, and Insurer pays 1. The example illustrates that the vaccine auction is not deceived by the contractual transfers between the bidders (here, the transfer from Insurer to Hospital, implicit in their valuations) and identifies the efficient allocation correctly.

"Lazy" (or "Inattentive") Bidding

One may be concerned about bidders being insufficiently motivated to discover how much they value the vaccine, especially when it is administered to others. Could a bidders could submit a "lazy"—inaccurate—bid (e.g., a zero) at no cost to himself but to the great detriment to efficiency? No. The concern is unjustified. In the vaccine auction, each bidder's payoff is the total surplus plus a constant. A "lazy" bid that affects the total surplus has payoff consequences.

Covid Economics 68, 10 February 2021: 100-114

CEPR PRESS

Moreover, the result of Bergemann and Välimäki (2002, Corollary 1) applies: the vaccine auction provides bidders with socially optimal incentives to discover their valuations. That is, in the model's extension in which each bidder chooses the precision of a costly signal about his valuations, under appropriate separability assumptions on the costs, the vaccine auction induces a game that has an equilibrium in which each bidder acquires the amount of information that maximizes the expected total surplus generated by the auction. In other words, any bidder inattention to valuations is not only privately rational but is also socially optimal (i.e., efficient).[4]

Self-Image and Privacy

The described vaccine auction forces each bidder to confront the potentially uncomfortable problem of assigning a cash value to his own life, as well as to the lives of relatives, co-workers, and employees. "Am I the kind of person capable of, first, cold-heartedly putting a price tag on a life and then living with the memory of this number for the rest of my life? Would I like others to know whose lives I value and how much?" If the answer to either question is no, then a bidder may refrain from participating in the auction, in order to protect his self-image or privacy, or both. Following Benabou and Tirole (2006), self-image can be modeled as the inference of one's own valuations, suppressed or forgotten once the auction is over, from the allocation and the payments selected by the auction. The loss of privacy refers to the same inference performed by others. A so-called "differentially private" modification of the vaccine auction mitigates both concerns.

Roughly speaking, a differentially private auction ensures that no single bidder can significantly affect the probability distribution over the auction's outcomes, regardless of how others bid. Because the outcome of a differentially private auction depends little on the bids of any single bidder, it is impossible to infer any bidder's bids with much precision by inverting the auction's outcome. The vaccine auction described in this note is not differentially private; a bidder can either deny or guarantee himself an early vaccination by submitting appropriate bids. A differentially private modification of the vaccine auction can be obtained by carefully injecting randomness into the auction's allocation and payments. Huang and Kannan (2012) show how

[4] Bergemann and Välimäki's (2002) result also implies that the vaccine auction motivates the bidders to undertake socially optimal investments in enhancing their valuations of the vaccine.

Covid Economics 68, 10 February 2021: 100-114

this can be done for any VCG mechanism. The implied modified vaccine auction admits a little inefficiency in exchange for some differential privacy while remaining strategy-proof.

4 Concluding Remarks

In conclusion, I address some of the challenges that the adoption of the vaccine auction may encounter.

1. If bidders face credit constraints, the normatively appealing features of the vaccine auction are lost. For an extreme example, suppose that Alice has no spare cash to bid but values the vaccine a lot (e.g., because of old age or a health condition). By contrast, Bob has ample cash but values the vaccine only a little (e.g., because he is young or has already had the virus). Suppose that neither Alice nor Bob cares about the other, and the two are the only bidders. Then, in the auction, Bob outbids Alice for an early vaccination appointment, even though it is efficient to vaccinate Alice first. Even a lottery would do better.
   Without restoring full efficiency, others' bids on behalf of those who are likely to be credit constrained may mitigate the inefficiency. For instance, a healthcare provider (private or public) may bid generously on behalf of the credit constrained as long as the credit constrained are elderly or chronically ill, so that their hospitalization is extremely costly to the provider.
   How pervasive credit constraints are in practice is an empirical matter.[5]

2. The vaccine auction's implicit ethics is that transfers are ethically neutral: a transfer from one bidder to another does not affect social welfare. This neutrality follows by the envelope theorem if the current redistributive taxation in the economy is socially optimal. An ethicist who rejects the neutrality of transfers—for instance, by believing that social welfare would rise if a rich man were to pay a dollar to a poor man—would instinctively oppose the proposed vaccine auction. This ethicist would be advised to investigate the possibility of separating the problem of distributive justice from the problem of efficient allocation of

[5]In a survey of microcredit, Banerjee, Karlan and Zinman (2015) suggest that the effect of relaxing credit constraints for "underserved entrepreneurs" may be limited: "The studies do not find clear evidence, or even much in the way of suggestive evidence, of reductions in poverty or substantial improvements in living standards. Nor is there robust evidence of improvements in social indicators."

Covid Economics 68, 10 February 2021: 100-114

CEPR PRESS

vaccines, instead of subverting efficiency in the name of a little redistribution.[6] That is, if an identifiable group (e.g., the poor, the uninsured, or the unemployed) is deemed to deserve higher welfare, it is best to give them cash, which they can then spend on vaccines or on something else, rather than force them to consume vaccines, which someone else could value more. If one's concern is not so much the welfare of the identified group as their vaccination, then one (e.g., the ethicist, a government agency, or a charity) is free to bid on behalf of this group's members in the vaccine auction. The ethicist would also be advised to ask himself whether any perceived pre-existing distributive injustices would be addressed with greater legitimacy by slow democratic deliberation once the pandemic has been conquered rather than by emergency fiat in the midst of a pandemic.

3. Adoption of a vaccine auction would face political constraints. The government, at its every level, would be reluctant to cede the power to dole out favorable treatment to various interests groups and lobbyists. One could try to convince the government of the merits of the auction by pointing out that any allocation rule that is perceived to be inefficient (e.g., because it panders to select interest groups at the expense of unorganized individuals, and because it neglects privately observed valuations) and that is perceived to be vulnerable to gaming (i.e., is not strategy-proof) will be unpopular. Furthermore, citizens are unlikely to be happy with any priority list that emanates from the government and that appears to rank them according to their worthiness; it is impossible to rank everyone near the top. Therefore, it is politically safest if this ranking is outsourced to an impersonal market, an auction.
   The government should also resist the argument that because vaccines have been purchased from pharmaceutical companies using taxpayers' money, these vaccines should be distributed by the government and free of charge. To accept this argument is to conflate the problems of distributive justice and efficiency. The proposed vaccine auction generates revenue, which can be used to retire some of the national debt; the taxpayers will get their money back.

4. The proposed auction ignores complementarities. Any practical implementation would likely have to address this shortcoming. For instance, a company that values herd immunity in the town where it is located may not value much vaccination of a handful of its employees

[6] Dworczak, Kominers and Akbarpour (2020) discuss when the two problems are inseparable and propose a solution, which calls for rationing and price controls.

Covid Economics 68, 10 February 2021: 100-114

but would be prepared to pay a premium if most of its employees are vaccinated. A bidding language that expresses such complementarities can be developed.[7]

5. One could be tempted to reject the vaccine auction in favor of a compromise solution: first, let the government allocate the vaccines however it sees fit and then allow individuals to trade their vaccination priorities among themselves. The problem with this approach is the ineluctable inefficiency that arises from the Myerson and Satterthwaite (1983) impossibility theorem. Once private ownership is introduced, markets need not put vaccines into the hands of the individuals with the highest valuations when the potential buyers' and the potential sellers' valuations are private information. If one insists on efficiency and strategy-proofness, the proposed vaccine auction is essentially necessary.

A Proof of Proposition 1

Efficiency is by construction of the allocation rule $\mathbf{x}^*$ in (1).

For strategy-proofness, note that, for any collection $\mathbf{b}_{-i}$ of others' bids, for any collection $\mathbf{b}_i$ of bidder i's bids, and for any collection $\mathbf{s}_i \equiv (s_{ij})_{j\in\mathcal{I}}$ of bidder i's true valuations, the payoff from bidding the $\mathbf{s}_i$ weakly exceeds the payoff from bidding $\mathbf{b}_i$:

$$\begin{aligned}
&\sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x^*_{jt}(\mathbf{s}_i,\mathbf{b}_{-i})\,\alpha_t s_{ij} - p^*_i(\mathbf{s}_i,\mathbf{b}_{-i})\\
&\quad= \sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x^*_{jt}(\mathbf{s}_i,\mathbf{b}_{-i})\,\alpha_t\left(\sum_{j'\in\mathcal{I}\setminus\{i\}} b_{j'j}+s_{ij}\right) - \sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x^*_{jt}(0,\mathbf{b}_{-i})\,\alpha_t \sum_{j'\in\mathcal{I}\setminus\{i\}} b_{j'j}\\
&\quad= \max_{\mathbf{x}} \sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x_{jt}\alpha_t\left(\sum_{j'\in\mathcal{I}\setminus\{i\}} b_{j'j}+s_{ij}\right) - \sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x^*_{jt}(0,\mathbf{b}_{-i})\,\alpha_t \sum_{j'\in\mathcal{I}\setminus\{i\}} b_{j'j}\\
&\quad\geq \sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x^*_{jt}(\mathbf{b}_i,\mathbf{b}_{-i})\,\alpha_t\left(\sum_{j'\in\mathcal{I}\setminus\{i\}} b_{j'j}+s_{ij}\right) - \sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x^*_{jt}(0,\mathbf{b}_{-i})\,\alpha_t \sum_{j'\in\mathcal{I}\setminus\{i\}} b_{j'j}\\
&\qquad\qquad= \sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x^*_{jt}(\mathbf{b}_i,\mathbf{b}_{-i})\,\alpha_t s_{ij} - p^*_i(\mathbf{b}_i,\mathbf{b}_{-i}),
\end{aligned}$$

[7] Budish, Cachon, Kessler and Othman (2017) propose a simple bidding language that accommodates complements and substitutes alike.

Covid Economics 68, 10 February 2021: 100-114

where the first and the last equalities use (2), the second equality uses (1), and the inequality replaces an allocation rule that is optimal for $(\mathbf{s}_i, \mathbf{b}_{-i})$ by a rule that need not be optimal for $(\mathbf{s}_i, \mathbf{b}_{-i})$ (but happens to be optimal for $(\mathbf{b}_i, \mathbf{b}_{-i})$). Strategy-proofness follows by combining the chain of equalities and the inequality.

For the "moreover" part, the first line in the display above can be rewritten using (1) as

$$\sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x^*_{jt}\left(\mathbf{s}_i, \mathbf{b}_{-i}\right)\alpha_t s_{ij} - p^*_i\left(\mathbf{s}_i, \mathbf{b}_{-i}\right) = \max_{\mathbf{x}\in X}\sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x_{jt}\alpha_t\left(\sum_{j'\in\mathcal{I}\setminus\{i\}} b_{j'j} + s_{ij}\right) - \max_{\mathbf{x}\in X}\sum_{t\in\mathcal{T}}\sum_{j\in\mathcal{I}} x_{jt}\alpha_t\left(\sum_{j'\in\mathcal{I}\setminus\{i\}} b_{j'j}\right),$$

which is nonnegative when $s_{ij} \geq 0$ for each i and j in $\mathcal{I}$.

References

Ausubel, Lawrence M. and Paul Milgrom, "The Lovely but Lonely Vickrey Auction," in Peter Cramton, Yoav Shoham, and Richard Steinberg, eds., *Combinatorial Auctions*, MIT Press, 2005.

Banerjee, Abhijit, Dean Karlan, and Jonathan Zinman, "Six Randomized Evaluations of Microcredit Six Randomized Evaluations of Microcredit," *American Economic Journal: Applied Economics*, 2015, *7* (1), 1–21.

Benabou, Roland and Jean Tirole, "Incentives and Prosocial Behavior," *American Economic Review*, 2006, *96* (5), 1652–1678.

Bergemann, Dirk and Juuso Välimäki, "Information Acquisition and Efficient Mechanism Design," *Econometrica*, 2002, *70* (3), 1007–1033.

Budish, Eric, Gerard P. Cachon, Judd B. Kessler, and Abraham Othman, "Course Match: A Large-Scale Implementation of Approximate Competitive Equilibrium from Equal Incomes for Combinatorial Allocation," *Operations Research*, 2017, *65* (2), 314–336.

Dworczak, Piotr, Scott Kominers, and Mohammad Akbarpour, "Redistribution through Markets," *Econometrica*, 2020.

Covid Economics 68, 10 February 2021: 100-114

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwartz, "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 2007, *97* (1), 242–259.

Huang, Zhiyi and Sampath Kannan, "The Exponential Mechanism for Social Welfare: Private, Truthful, and Nearly Optimal," in "EEE 53rd Annual Symposium on Foundations of Computer Science" 2012, pp. 140–149.

Myerson, Roger B. and Mark A. Satterthwaite, "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 1983, *29* (2), 265–281.

Varian, Hal R., "Position Auctions," *International Journal of Industrial Organization*, 2007, *25*, 1163–1178.

__ , "Online Advertising Markets," *American Economics Review: Papers and Proceedings*, 2009, *99* (2), 430–434.

Covid Economics Issue 68, 10 February 2021

The impact of Covid-19 on the European short-term rental market[1]

Elisa Guglielminetti,[2] Michele Loberto[3] and Alessandro Mistretta[4]

Date submitted: 2 February 2021; Date accepted: 2 February 2021

The diffusion of COVID-19 and related containment measures practically halted tourism flows, which in many countries generate more than 10% of GDP. By exploiting Airbnb data covering the main touristic destinations in Europe, we investigate how the exposure to COVID-19 and the stringency of containment measures affected the market of short-term rentals over the spring and summer months of 2020. We find that the epidemic reduced dramatically both the supply of apartments available for rents and the consumers' demand, up to 9 months ahead. Prices fell as well, even at long time horizons, but with a delay. All in all, our results point to a persistent impact of COVID-19 and related containment measures on consumers' behaviour, with demand shortages potentially overcoming supply shortfalls.

Covid Economics 68, 10 February 2021: 115-138

1 The views expressed do not necessarily reflect those of the Bank of Italy.
2 Bank of Italy, Directorate General for Economics, Statistics and Research.
3 Bank of Italy, Directorate General for Economics, Statistics and Research.
4 Bank of Italy, Directorate General for Economics, Statistics and Research.

CEPR PRESS

1 Introduction

The spread of the COVID-19 epidemic worldwide has led many governments to adopt drastic containment measures. Social distancing, limits to mobility, and the forced closure of many productive sectors, together with difficult public health conditions, determined a dramatic drop in GDP in the first half of 2020. A key question regards the nature of this unprecedented shock: although the pandemic was a large negative supply shock, due to the forced shutdown of economic activities, demand shortages, connected to consumers' falling income and behavioral changes, also played a very important role (Guerrieri et al., 2020). A related question is whether COVID-19 will have a transitory impact on economic activity or whether it will influence – and to what extent – agents' expectations and behavior, with more persistent effects.

We provide a contribution on these issues by leveraging data on the market of short-term rentals. We use microdata collected by InsideAirbnb.com, which analyzes publicly available information about Airbnb listings for several cities worldwide. Our sample includes listings in 19 major European cities located in 15 different countries, from early 2018 until September 2020. Importantly, we can analyze the impact of COVID-19 on supply, demand (captured by booking and cancellation rates) and prices of short-term rentals. Moreover, the data are forward-looking, as we can look at agents' decisions up to 9 months ahead the observation date. Therefore, we can assess whether the spread of the epidemic influenced only short-term decisions or, instead, led agents to review their plans also several months ahead. These unique features of the data make them suitable to investigate both the short-run impact and the potential long-run consequences of the pandemic, distinguishing between supply-like and demand-like effects.

Although carried out on a specific market segment, our analysis provides insights on the impact of COVID-19 at a more general level. The results likely extend to the whole tourism sector, considering that: i) Airbnb is the largest peer-to-peer platform of short-term rentals and in recent years has become an important player even compared to the traditional hotel industry; ii) accommodation is strictly connected to other activities in the same sector, such as, in particular, restaurants and the cultural and recreational industry. Moreover, the tourism sector is particularly interesting, as it was the hardest hit by the pandemic and related containment measures. Tourism flows stopped almost completely in spring 2020 and only partially resumed in summer; as the uncertainty about the possibility of safe travelling looms large, the pandemic is likely to be having an impact also on medium-run prospects. By studying the tourism sector we get insights on consumers' behavior more in general, in particular concerning spending cuts driven by precautionary motives and fear of infection. Finally, tourism is an important driver of the overall economy, especially in some countries like Italy and Spain,[1] where its share on GDP, taking into account both direct and indirect effects, is estimated at 13.2 and 14.6%, respectively.[2]

The sharp decline in tourism flows and economic activity was determined by both govern-

[1]See MacDonald et al. (2020).

[2]World Travel and Tourism Council (2017).

Covid Economics 68, 10 February 2021: 115-138

ment measures and the diffusion of the epidemic per se. Although disentangling quantitatively these two factors is beyond the scope of this work, we provide a contribution in this direction by taking advantage of the different timing of the epidemic spread across Europe and the different timelines and intensity of the measures adopted to contain it. We measure the severity of the epidemic by infection rates and occupancy of Intensive Care Units (ICUs) by COVID-19 patients. The strictness of the containment measures is captured by the Oxford Stringency Index (Petherick et al., 2020). Since our sample period runs until September 2020, we can investigate the reaction of the short-term rental market both in the early phase of the epidemic, characterized by high degrees of contagion and the enforcement of strict lockdowns in many countries, and during the partial revival of tourism flows in summer. Moreover, we observe market activity over different time-horizons, and we can thus estimate the heterogeneous effects of each driver at different future dates.

The descriptive statistics and the econometric analysis show that the outbreak of the epidemic had a dramatic negative impact on market activity, on all dimensions. The supply of houses for rental decreased and customers reacted to the shock by canceling existing bookings and drastically reducing new bookings across all time horizons. Prices behave somewhat differently, as they started to decrease later in the year, probably reflecting owners' assessment that demand would have been rather inelastic. Government restrictions appear to have played an important role in affecting both demand and supply.

Although becoming smaller at longer time-horizons, the effects of COVID-19 are quite persistent and, in some cases, they remain significant up to 9 months ahead. The impact on prices is not significant, probably because they reacted with a lag to the worsening health conditions.

Our paper speaks to the rapidly expanding literature on the effects of COVID-19. In this work, we provide new evidence by focusing on the accommodation industry and households' travel choices. As in Hu and Lee (2020), we analyze the short-term rental segment based on Airbnb listings. Hu and Lee (2020) use the evolution of daily guests' reviews, considered a proxy for stays in the facilities, until the end of March 2020. Unlike them, we focus on the impact of the epidemic on the longer-term hosts and guests' choices, like booking and cancellations several months ahead, and we exploit data up to September 2020.

As in Brinca et al. (2020) and Guerrieri et al. (2020), we assess the relative importance of demand and supply shocks; however, instead of conducting an analysis across sectors, we focus on a specific industry severely hit by the epidemic.

Finally, our paper is related to the literature that evaluates the economic costs of a lockdown. Many papers show that a lockdown has a very strong negative impact on GDP but increases welfare as it reduces the number of deaths (see, for example, Eichenbaum et al., 2020). Kaplan et al. (2020) show that lockdowns have asymmetric effects among different population groups. Coibion et al. (2020) estimate that the lockdown had a negative impact on both consumption and longer-term expectations. Some studies argue that the effect of lockdown may have been minimal because people had already limited their mobility for fear of becoming infected (see,

Covid Economics 68, 10 February 2021: 115-138

for example, Andersen et al., 2020, Goolsbee and Syverson, 2021, Rojas et al., 2020). Also Buono and Conteduca (2020) show that mobility in Europe has started decreasing before the introduction of government restrictions. However, many of these papers refer to the first wave of the epidemic; Franks et al. (2020) argue that different re-opening strategies, in the period covered by our analysis, significantly affected mobility and infection rates. Relative to these papers, we estimate the effects of containment measures by taking into account not only the initial implementation of lockdown policies but also their subsequent easing. Furthermore, we can analyze their impact in both the short-run and several months ahead.

2 The data

Our sample includes Airbnb listings for 19 major European cities located in 15 different countries.[3] These cities are the capital of their country or major destination of tourist flows. We retrieve web-scraped data on listings from InsideAirbnb.com, an independent project collecting data from the Airbnb website to study the implications of short-term rentals on the housing market. For each city, we have snapshots of listings collected monthly between April 2018 and September 2020, although they may not refer to the same day of the month. A snapshot is the collection of all listings in a city published on www.airbnb.com on a given day, and from each of them, we keep in our sample only listings of entire apartments, excluding shared or private rooms that may be differently affected by the ongoing epidemic due to social distancing measures.

All information in our dataset is public and displayed on the Airbnb site. We observe many details for each listing (for example, information about the host, physical characteristics of the house, location, rental policy), including the calendar availability for the following 365 days, the guests' reviews, and the price per night of the accommodation. We mostly use the information on calendar availability, which allows us to estimate the number of bookings and cancellations and occupancy rates.[4] We can infer the evolution of the occupancy rates also through guests' reviews, although we cannot use them to compare the occupancy rate across cities. Finally, calendar data allow observing prices and estimating hedonic indexes over different time horizons.

The analysis of these data requires some caution. Airbnb is an online marketplace that allows homeowners (hosts) to offer rental services to other people (guests). Differently from the standard hotel and accommodation industry, a prominent share of suppliers consists of non-professional operators.[5] Those hosts may not rent the house regularly and may not update

[3]The cities are Amsterdam, Athens, Barcelona, Berlin, Brussels, Copenhagen, Dublin, Florence, Geneva, Lisbon, London, Madrid, Milan, Oslo, Paris, Rome, Stockholm, Venice and Vienna.

[4]For each day in the calendar, we can only observe if a house is available for booking or not. An apartment cannot be available for booking because it was already booked or because the host does not accept reservations for that particular day.

[5]Recently, intermediaries have emerged who specialize in taking empty houses over from owners and managing short-term rentals on the Airbnb marketplace.

Covid Economics 68, 10 February 2021: 115-138

the calendar frequently; thus, calendar data can be the same in two different snapshots, and we cannot infer activity in the time-lapse. Furthermore, some hosts deliberately limit the time horizon for which they accept bookings (for example, giving the possibility to book only up to 90 days ahead). In this case, the house appears unavailable beyond this time horizon.

To limit these issues and identify hosts that are more active in the market, we make the following choices. First, from each snapshot, we keep only listings with an updated calendar compared to the previous observation. We consider the calendar updated if at least one of the following conditions is satisfied: (i) there is a change in the availabilities compared to the previous snapshot; (ii) the host has changed its prices. Second, if the calendar dates for a listing are all unavailable from a specific date onwards, we assume they are not bookable; hence, we include these listings in our analysis only for the period in which they are bookable.

Using these criteria, we select the most active listings and hosts on the market. In this work, we are interested in quantifying the effects of the COVID-19 epidemic on the supply and demand for short rentals promptly. Considering listings with outdated calendars, or more generally host not very active on the market, would have led to underestimating the effects of the shock.[6]

3 The evolution of the epidemic and containment measures

To gauge the impact of the epidemic on short-term rental markets, we rely on data from the European Center for Disease Prevention and Control (ECDPC) and from the Institute of Health Metrics and Evaluation (IHME). ECDPC collects daily data on the number of COVID-19 cases and deaths and weekly data on the number of tests. We complement these data with information on intensive care units (ICUs) used for COVID-19 patients taken from the IHME.[7] To measure the spread of contagion we focus on the number of COVID-19 infections and occupation of ICUs; although statistics may be not fully comparable across countries, for the purpose of our analysis what matters the most is the perceived risk of contagion, which depends on reported data. The number of reported infections does itself depend on the testing capacity, which has strongly increased compared to the early phase of the pandemic. For this reason, we focus on the positivity rate, namely the number of infections out of the number of tests. We measure the restrictiveness of government policies adopted to fight the epidemic through the Oxford Stringency Index. The index is measured on a scale from 0 to 100, where higher values denote more restrictive policies, and is broadly comparable across countries.[8] In our econometric model,

[6]We do not believe that these are universal criteria to adopt in all circumstances, as they may not be necessary for other analyses. After all, Airbnb aims to allow people (not professionals) to share their house with other people in return for a fee, and the calendar may be outdated because bookings are very low in some periods of the year.

[7]For countries and/or time periods for which ECDPC data are not available (e.g. Switzerland) we use IHME data or data taken from Our World in Data (OWD).

[8]Conteduca et al. (2020) identify come limitations of the Oxford Stringency Index. They propose to redefine the score of some variables and to weigh the sub-national restrictive measures by the share of the targeted

Covid Economics 68, 10 February 2021: 115-138

we consider separately the sub-component of the Stringency Index, capturing restrictions to international travels and the overall index purged from the former.[9]

The first COVID-19 case was reported in China, in the Hubei region, on January 9, 2020.[10] Among the selected European countries considered here, Italy was the first one, on January 23, to introduce measures aimed at detecting passengers with symptoms possibly related to COVID-19 flying from China. In Europe, the first COVID-19 cases were identified at the end of January; however, the disease remained mostly undetected until the end of February, when infection clusters in Northern Italy became apparent. In Italy the epidemic gained momentum at a fast pace: in the first week of March the number of new cases per day more than tripled and the positivity rate exceeded 20%. This prompted the Italian authorities to impose a nationwide lockdown since March 10. In March and April the epidemic spread in all European countries, albeit with different timings and intensities. Beside Italy, the hardest and earlier hit countries were Spain, Austria and Belgium, where the positivity rate reached its peak already in late March. On the other side of the spectrum, in Germany, Portugal and Greece the positivity rate never exceeded 10% during the first wave. In the first two panels of Figure 1 the spread of the contagion is represented by the evolution of the infection and the positivity rates, where the latter controls for the testing capacity.

The severity of the epidemic can be further judged with respect to the strain on the health-care system. This can be measured by the share of ICUs dedicated to COVID-19 patients over total population, and depends not only on the severity of contagion, but also on the ability of local authorities to keep the sanitary situation under control and on pre-existing conditions. ICU occupation generally peaks later than the positivity rate and is more persistent, given the length of stay of the patients (Figure 1, third panel). These somewhat different patterns allow us to distinguish the impact of the spread of contagion from that of overwhelmed ICUs even in a context of declining infection rates.

Another dimension of heterogeneity regards the containment measures adopted in response to the spread of the virus. Italy, Spain and Austria enforced a strict lockdown already in March (Figure 1, fourth panel). Other countries adopted different approaches: the UK government, for instance, imposed shutdown orders only when the strategy of achieving herd immunity appeared to be too costly in terms of human lives. Sweden was the only European country that relied on mild mitigation policies, notwithstanding relatively high levels of contagion and ICU occupancy.

In May, the epidemiological situation improved and containment measures were eased; since June, people were allowed to travel also outside their country of residence, although with some restrictions. However, some European countries, like Spain, experienced a new surge in COVID-

population. Though acknowledging its limitations, we use the Oxford Stringency Index to ensure comparability with other studies.

[9]The Oxford Stringency Index is a weighted average of eight subcomponents: for a detailed description, see Petherick et al. (2020). In the Appendix we describe how we separate the domestic and the international sub-components.

[10]Several studies have then proved that the coronavirus was already diffused; however, here we are interested in the official communication and the spread of news to the public.

Covid Economics 68, 10 February 2021: 115-138

CEPR PRESS

19 cases already in early July. In France, the UK, the Netherlands and Belgium the contagion quickly spread in August and September, while in Italy and Germany the diffusion of the infection remained modest until the beginning of October.[11] Despite rising infection rates, until September - the end of our sample period - governments avoided a new tightening, hoping to keep the evolution of the epidemic under control while preserving economic activities.

Figure 1: The evolution of the epidemic and policy responses

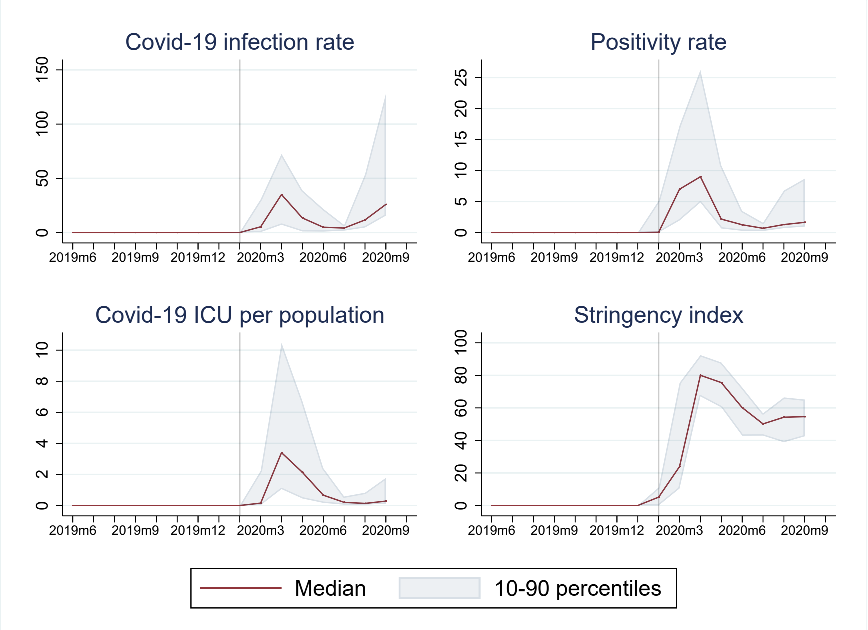

Note: The infection rate is measured as COVID-19 new cases per population (in thousands). The positivity rate is the share of detected infections over total tests.

4 The effects of the epidemic on short-term rental markets: stylized facts

Difficult health conditions and containment measures adopted by governments led to a collapse in the mobility of people between cities and between countries. Consequently, levels of activity in the accommodation sector, and thus also in the short-term rental market, collapsed.

This section describes the evolution of the short-term rental market after the outbreak of the coronavirus epidemic under different dimensions. The epidemic led to a gradual decrease in the number of houses on the market. Customers reacted to the shock by canceling existing bookings and drastically reducing new bookings across all time horizons. Therefore, the accommodation facilities' occupancy rates dropped, and owners progressively reduced the required prices, even on longer time horizons. For each variable, we take into account the possibility

[11]See Borin et al., 2020 for a detailed analysis of worldwide evolution of the pandemic since the summer.

Covid Economics 68, 10 February 2021: 115-138

Figure 2: Number of short-term rentals

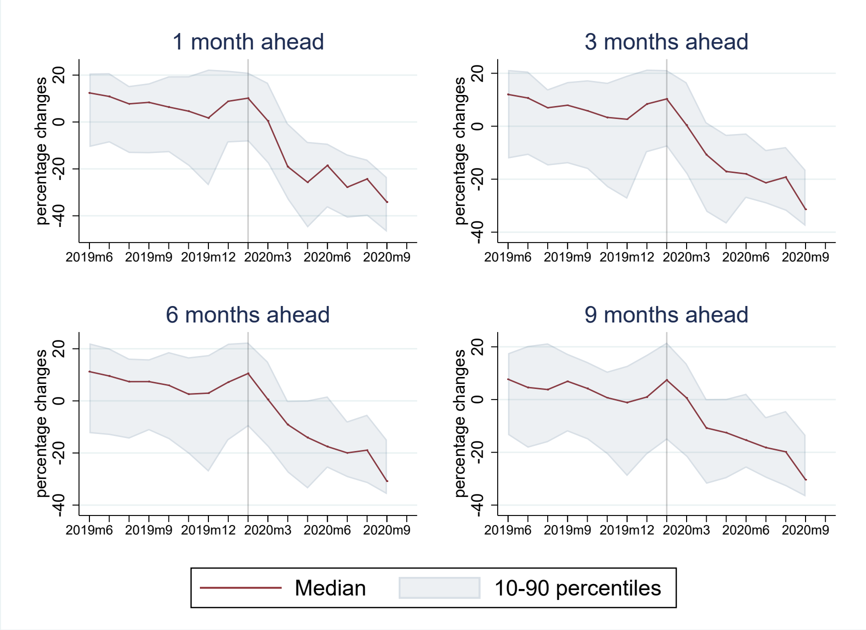

Note: Y-o-y percentage changes. Supply can change over different time horizons because some hosts allows reservations only for a limited period forward in time.

that the effects of the epidemic may have been heterogeneous at different time horizons (denoted by h), going from 1 month ahead up to 9 months ahead. Then, we compute for each city the y-o-y monthly variations and show how the 10th, 50th and 90th percentiles of the distribution of these variations have changed over time.

The supply of short-term rentals. Before the COVID-19 outbreak, the market for short-term rentals was expanding in most of the cities in our sample. In January and February, the number of listings was decreasing significantly only in Dublin and Amsterdam, due to pre-existing downward trends. The y-o-y median growth rate of available accommodations across cities was about 10%. In Vienna, Athens and Brussels, growth rates were close to 20%.

This positive trend was interrupted in March (Figure 2). Initially, the supply contracted in particular in Italian cities, which were the first to be affected by the epidemic. The growth rate turned negative, from 9% to -14% in Rome and from 17% to -9% in Milan. Subsequently, the worsening affected all the other cities; in May the growth rate of the number of listings was negative everywhere.

The drop in supply was quite strong between March and May, but in most cities the decline continued even afterward. Between May and September, the offer was about a quarter lower than the corresponding period in 2019. In Dublin and Amsterdam, the offer roughly halved.

The sharp drop in supply observed in the spring can be mostly explained by governments' restrictive measures, which reduced the number of hosts actively operating in the market. De-

Covid Economics 68, 10 February 2021: 115-138

Figure 3: Cancellation rates

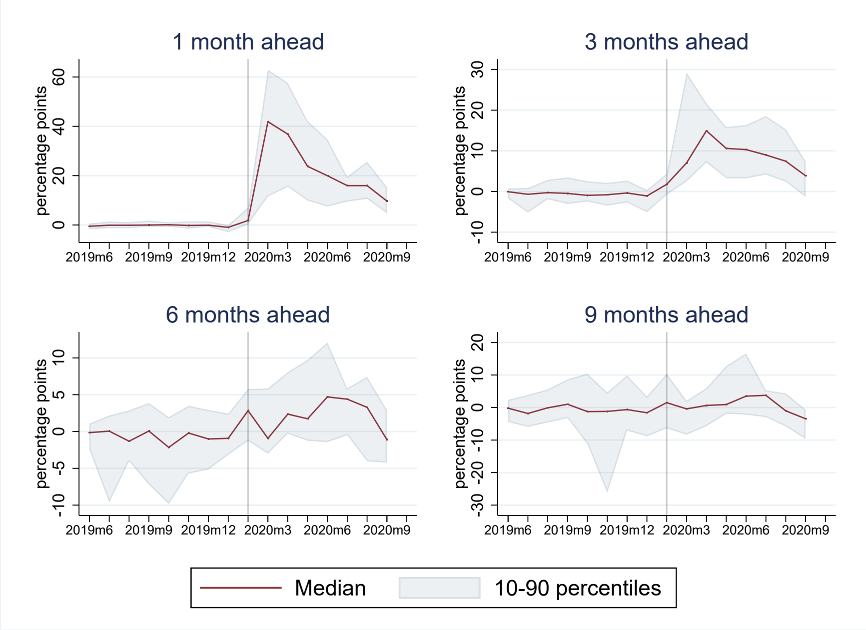

Note: Absolute differences compared to the previous year (percentage points).

Covid Economics 68, 10 February 2021: 115-138

spite the improvement in health conditions and the loosening of containment measures, the downward trend persisted also in the summer months, probably due to the uncertainty regarding the evolution of the epidemic, the travel restrictions in many countries and the unfavorable perspectives for the tourism sector. Many owners may have in fact preferred to rent their homes in the long-term rental segment, waiting for a return to normality. Moreover, the epidemic led to a worsening of the average quality of the apartments on the market, as can be deduced from the analysis of price trends (see below).

Cancellations and new bookings. The effects of the epidemic on current and future demand are visible. First, we estimated the share of booking cancellations over different time horizons.[12]

In March, the share of one-month-ahead bookings that were canceled increased markedly in all cities (Figure 3). The median increase compared to March 2019 was about 40 percentage points. In some cities, the increase in the cancellation rate exceeded 60 points.[13] The one-month-ahead cancellation rate was much higher compared to normal times, even in the following months, although the gap gradually narrowed.

Initially, the impact on cancellations was higher in the Italian touristic cities. In Rome, Venice and Florence the one-month-ahead cancellation rate was about 70% in March. The

[12]We estimate the share of cancellations instead of the absolute number to smooth out the strong seasonal pattern on these statistics. Technical details about the estimation of all indicators discussed in this note are in the Appendix.

[13]In normal times, the one-month-ahead cancellation rate varies between 2% and 10%.

Figure 4: Booking rates

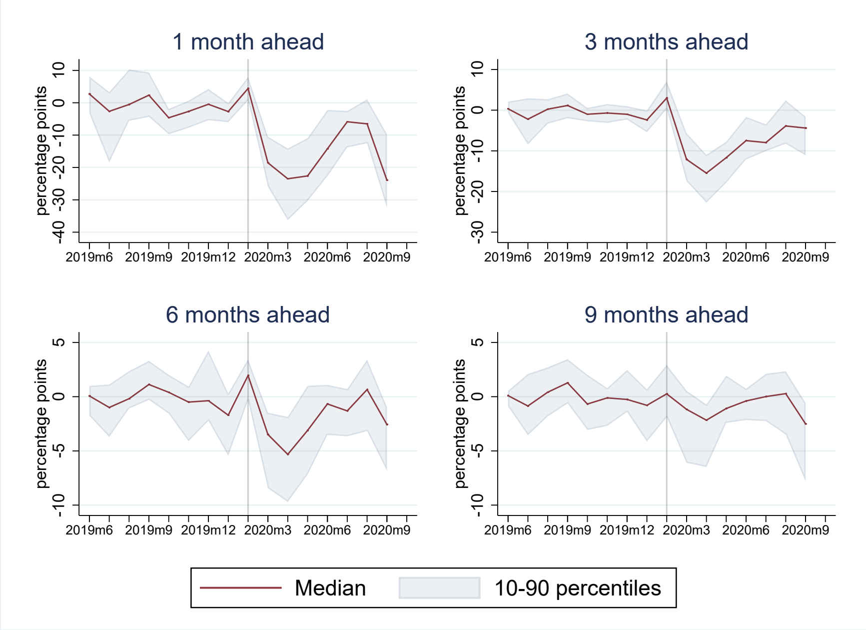

Note: Absolute differences compared to the previous year (percentage points).

worsening was more limited in the capitals of Northern European countries. Except for March, the impact was also lower in Milan, although Lombardy was the first Italian region hit by the epidemic and the most affected one. That could be due to the stronger business orientation of Milan, which plausibly makes the hosts' activity relatively less dependent on the dynamics of tourism flows compared to the other cities under scrutiny.[14] We observe a similar pattern for London, although the pandemic severely hit the UK.

The sharp increase in the share of pre-existing bookings that were cancelled in March and April is plausibly due to the mobility restrictions imposed by the governments.[15] However, travelers also canceled many reservations up to 3 months ahead, at a rate significantly higher than in normal times. Cancellation rates converge to those prevailing in normal times from the six-months ahead horizon.

The number of estimated new bookings also fell sharply since March (Figure 4). In April, the booking rate's median gap across cities compared to a year earlier reached -23 percentage points. Venice recorded the worst performance: the ratio of new bookings over previously

[14] According to Istat data, in 2018 the percentage of nights spent by foreigners in the province of Milan was almost half of the average share in the provinces of Rome, Venice and Florence.

[15] Formally, Airbnb remained active throughout the lockdown period because the short-term rentals offered through the platform could provide accommodation also for health workers. In Italy, for instance, Airbnb rentals are not formally considered as a part of the accommodation sector but can be run as a private rental activity. In practice, however, the majority of Airbnb services target the tourism sector (Guttentag et al., 2018), which was forcefully shutdown.

Covid Economics 68, 10 February 2021: 115-138

Figure 5: Vacancy rates

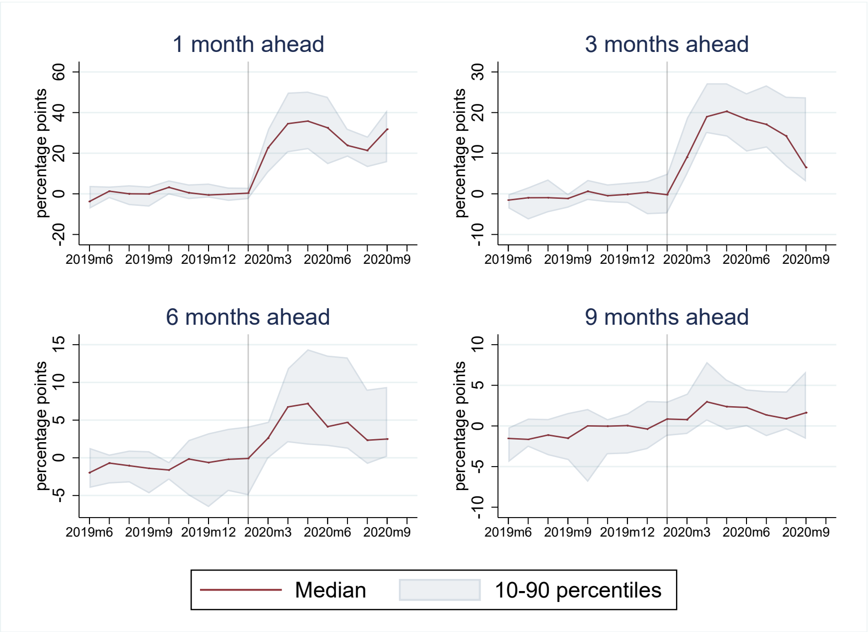

Note: Absolute differences compared to the previous year (percentage points).

Covid Economics 68, 10 February 2021: 115-138

available accommodations one-month-ahead fell to 6%, about 40 points lower than in 2019. All the main tourist destination cities were hardly hit also in terms of cancellations, mostly because many bookings for the summer months are usually made in spring.

Differently from cancellations, booking rates decreased significantly even at longer horizons (beyond 6 months). This indicates that customers significant revised their travel plans since the pandemic outbreak without waiting for the crisis's possible resolution. The fact that cancellations were instead limited to a shorter horizon may be explained by cancellation policies' flexibility: in March and April, some guests may have decided to wait before cancelling their reservations for the summer months as they could do it without incurring in immediate penalties.

Occupancy. Because of the joint occurrence of more cancellations and fewer new bookings, the vacancy rate – measured as the share of listings available for booking – skyrocketed in all cities (Figure 5). The impact is visible up to six months ahead. In the aftermath of the COVID-19 outbreak, the vacancy rate's median increase was about 40 percentage points for one month ahead reservations and 20 percentage points over a three-month horizon. As expected, the largest increase occurred in the main touristic destinations. The vacancy rate somewhat decreased during summer, although remaining far larger than in 2019 (the median variation was about 20 percentage points for one-month ahead reservations). In Barcelona, during the summer months, the vacancy rate was 40 percentage points larger than in 2019.

Figure 6: Reviews

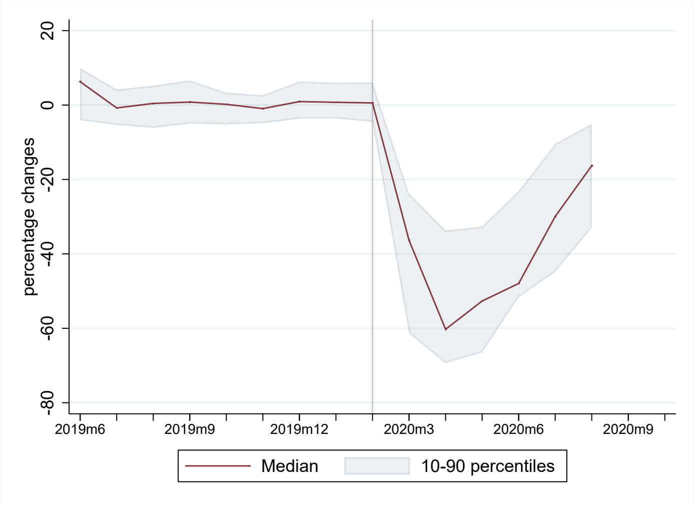

Note: Y-o-y percentage changes.

Figure 7: Total bookings

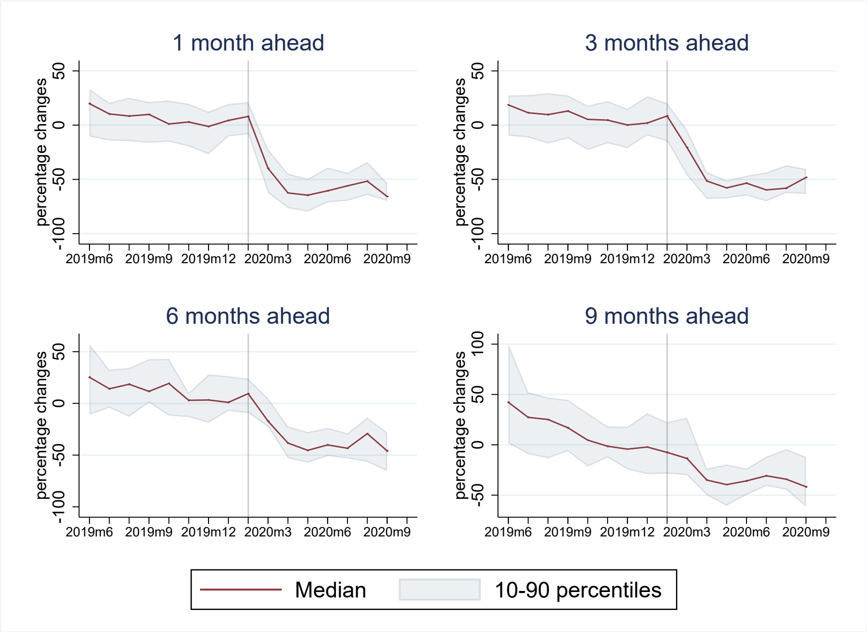

Note: Y-o-y percentage changes.

Covid Economics 68, 10 February 2021: 115-138

The insights derived from the vacancy rate are consistent with those coming from the average monthly number of reviews per apartment, which is a proxy of the occupancy rate.[16] This indicator decreased markedly during the lockdown, reaching a median y-o-y reduction of about 60% in April (Figure 6). Subsequently, the indicator improved, especially in July and August, although remaining significantly lower than a year earlier (the median y-o-y percentage variation was -17%). However, the vacancy rate and the average number of reviews measure only the degree of utilization of the apartments offered on the market. By taking into account the simultaneous reduction in the supply of houses (extensive margin), the contraction in the absolute number of bookings was even stronger (Figure 7). The median reduction in one-month ahead total bookings was about 60% compared to 2019. The effect was significant over all time-horizons, up to 9-months ahead, and more persistent compared to the previous indicators: bookings increased in July and August, but very modestly.

Prices. Before the epidemic, prices were growing markedly in most cities: the median variation was almost 10% y-o-y. In March and April, prices were not significantly affected, suggesting that the hosts considered demand to be largely inelastic, which is reasonable given that the fall in demand was due to both fear of contagion and limitations to mobility imposed by law (Figure 8). However, prices gradually decreased in the following months: the median variation of one-month ahead prices reached -17% y-o-y in September. The Italian and Spanish cities in our sample are those where prices fell the most. In Barcelona, in particular, during the summer months, prices were more than a third lower than in the previous year.

The drop of prices is quite significant, considering the upward trend before the epidemic and given the simultaneous sharp contraction in apartments' supply. Hosts reduced prices markedly up to one year ahead. That is important for two reasons. In our observation sample, future prices are excellent proxies for the spot prices in the following months. That is to say that the median price for renting a house in August is similar when posted in May, in July or in August. Second, accommodation prices published on Airbnb are significantly correlated with hotel prices.[17] Therefore, our results have implications beyond the Airbnb marketplace, pointing to a decrease in prices for accommodation services even when health conditions have normalized.

Price dynamics also makes it possible to indirectly measure changes in the accommodation's average quality. The median price of the apartments available on the market commented so far is affected, by construction, by composition effects. We thus build a second indicator, based on price changes for the same apartment in two contiguous months, to take into account the

[16]The vacancy rate measures perspective bookings observed at given date. In this case, the reference month refers to the date of web scraping. This indicator can be noisy, because data are not downloaded always in the same day of the month. At the opposite, the indicator based on the reviews refers to all the reviews observed in a given month. For this reason, reviews-based indicator can be more effective in measuring the occupancy rate.

[17]The correlation between the prices of Airbnb apartments and the price index of accommodation services of the respective countries is 0.4 in the period under consideration. For Italian cities, for which we have the corresponding price index of accommodation services at provincial level, the correlation is 0.6.

Figure 8: Prices

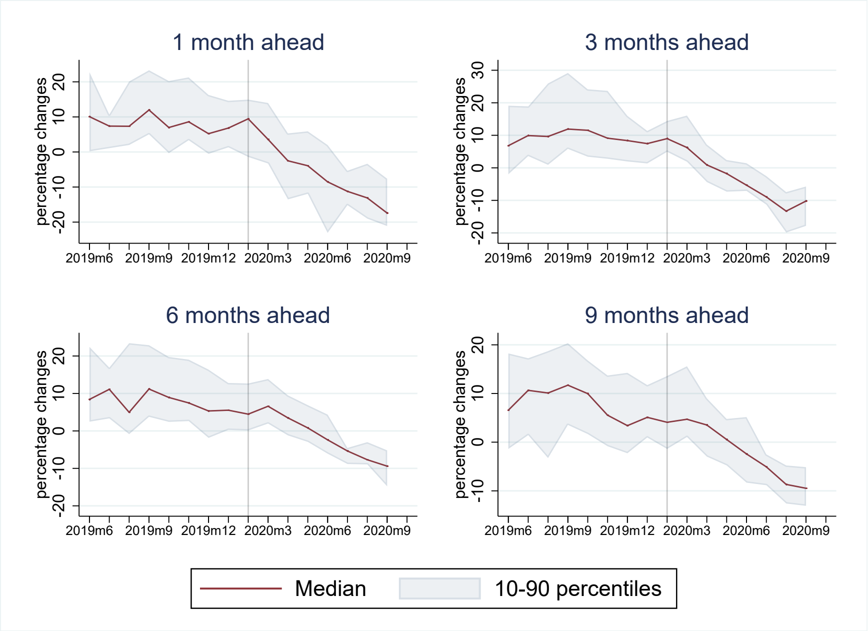

Note: Y-o-y percentage changes.

Figure 9: Hedonic prices

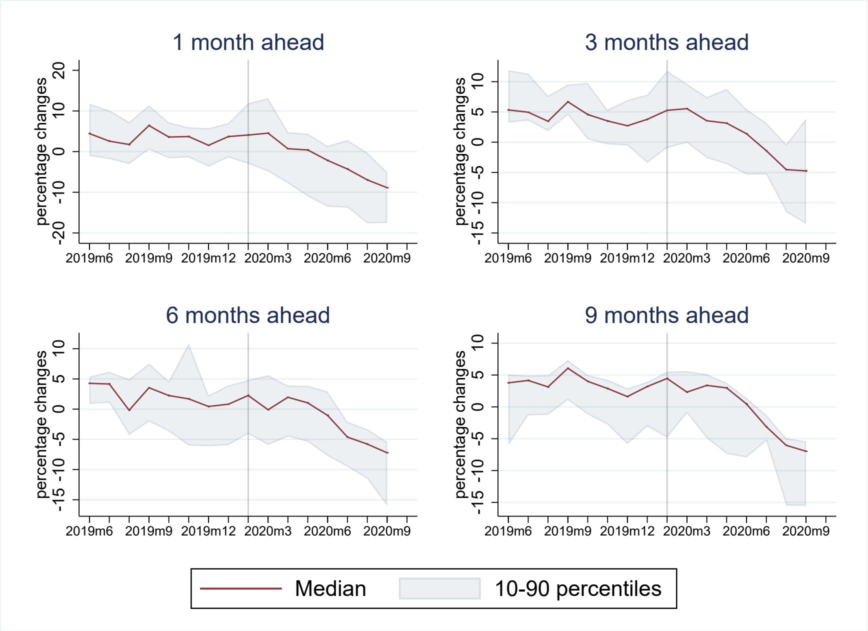

Note: Y-o-y percentage changes.

CEPR PRESS

characteristics of the houses offered.[18] In most cities, this indicator - which we improperly call "hedonic" - decreased less markedly than the other (Figure 9). This signals that composition effects contributed negatively to the dynamics of observed prices, implying that the apartments that left the market were of better quality than the remaining ones.

5 Econometric strategy

5.1 Model setup

Our econometric strategy takes advantage of the different timing and intensity of both the epidemic spread and counteracting containment measures across Europe, as described in the previous section. For each city, we estimate how prospective indicators of market activity at different horizons varied over time and to what extent the changes observed since March 2020 can be due to either the coronavirus diffusion or the containment policies. We measure the evolution of epidemiological conditions through the average infection rate and COVID-19 ICU per population between two close observation dates.[19] We proxy government containment measures with the average stringency index both for domestic and international travel restrictions.

Isolating the separate contributions of government measures, which limited both domestic and international travels, and of worsening health conditions, that could have reduced people's mobility even in absence of shutdown orders, is challenging. Given that our data are collected at monthly frequency and the sample period affected by COVID-19 is relatively short, we do not claim to achieve a sharp identification of the distinct impact of these two factors. However, the identification problem is mitigated for three reasons. First, some countries delayed the adoption of containment measures, or their intensity was initially low. Second, although data are collected at monthly frequency, the dates of web-scraping usually differ among cities, providing us with additional variation within the same month.[20] Third, as our dataset runs until September 2020, we can observe the effects of the epidemic at the beginning of the second wave of contagion in Europe, when containment measures had not been tightened yet. Overall, these sources of heterogeneity allow us to consider jointly the severity of epidemic and the degree of policy tightening within a unique regression framework.

As dependent variables, we consider the number of available houses, the price level (both as median of observed prices and median of the hedonic index), the booking rate, the cancellation

[18]The methodology for estimating this indicator is in the Appendix. One problem with this indicator is that it considers in each month only those apartments on offer in the previous month. As a result, selection effects may bias the indicator.

[19]Due to the poor knowledge of the epidemiological phenomenon at the beginning of the pandemic spread, it is possible that these indicators were subject to a greater measurement error during the first period of the pandemic.

[20]For example, data for Paris and London were scraped on March 15, before the announcement of the national lockdown. The March snapshot is after the outbreak of the epidemic and prior to the lockdown also for other cities. Finally, our sample includes Stockholm, for which government measures were very modest compared to other countries.

Covid Economics 68, 10 February 2021: 115-138

rate, and the vacancy rate. For each of these variables, we estimate the following linear regression models:

$$Y_{ith} = \alpha_{ith} + \beta_{1h} INF_{it} * h + \beta_{2h} ICU_{it} * h + \beta_{3h} STR_{it}^{INT} * h + \\ + \beta_{4h} STR_{it}^{DOM} * h + \gamma \mathbf{X}_{it} + \varepsilon_{ith} \quad (1)$$

where Y_{ith} is dependent variable (e.g. the vacancy rate) h-months ahead for city i and data scraped during month t. INF_{it}, ICU_{it} and STR_{it} are the average infection rate, the COVID-19 ICU per population and the stringency index in the country of city i, respectively;[21] as previously argued, we distinguish between domestic (DOM) and international travel (INT) restrictions using the corresponding components of the stringency index. Each variable is interacted with a set of dummies h, one for each time horizon, to obtain a marginal effect varying with the horizon.

For any snapshot of data, we take the average of the epidemiological and stringency over the period between the date of the previous snapshot and the previous day of the current snapshot. We also consider city-month-horizon fixed effects (α_{ith}), to account for heterogeneity across cities and for city-specific monthly seasonal effects specific for any time-horizon. Finally, $\mathbf{X}_{it}$ is a vector of additional control variables, such as the number of total tests per population, the total houses used to compute the dependent variable, the span period between the current and previous snapshots and the decade of the web-scraping date. We do not use controls related to population characteristics since they are indirectly captured by the fixed effects already included in the model.

5.2 Results

In this section we present the estimates of the coefficients of interest in eq. (1), namely those capturing the effects of the sanitary variables and the stringency index. As previously explained, the empirical strategy gives the possibility to estimate a different impact for any time horizon, up to 9 months ahead the observation date. This feature allows us to distinguish short-term effects (1-2 months ahead) from more persistent ones, more likely influenced by agents' expectations on future developments. We graphically show estimates for our parameters of interest in Figures 10-15, where we plot both point estimates and confidence intervals at 10% level obtained using errors clustered at the city level. The total effect associated to the stringency index combines the domestic and the international component. Our results show that the outbreak of COVID-19 determined a sharp reduction in the supply of short-term rentals; this effect seems more related to the restrictions adopted by the governments to fight the epidemic rather than the spread of virus itself (Figure 10). The effect of the stringency index is negative and significant over all time horizons, though becoming smaller over time: the impact 9 months ahead is roughly half that for the short-term. Our estimates imply, for instance, that an increase of the

[21] Since the data for the two cities may have been scraped at different dates, these variables are different even for cities belonging to the same country.

Covid Economics 68, 10 February 2021: 115-138

Covid Economics 68, 10 February 2021: 115-138

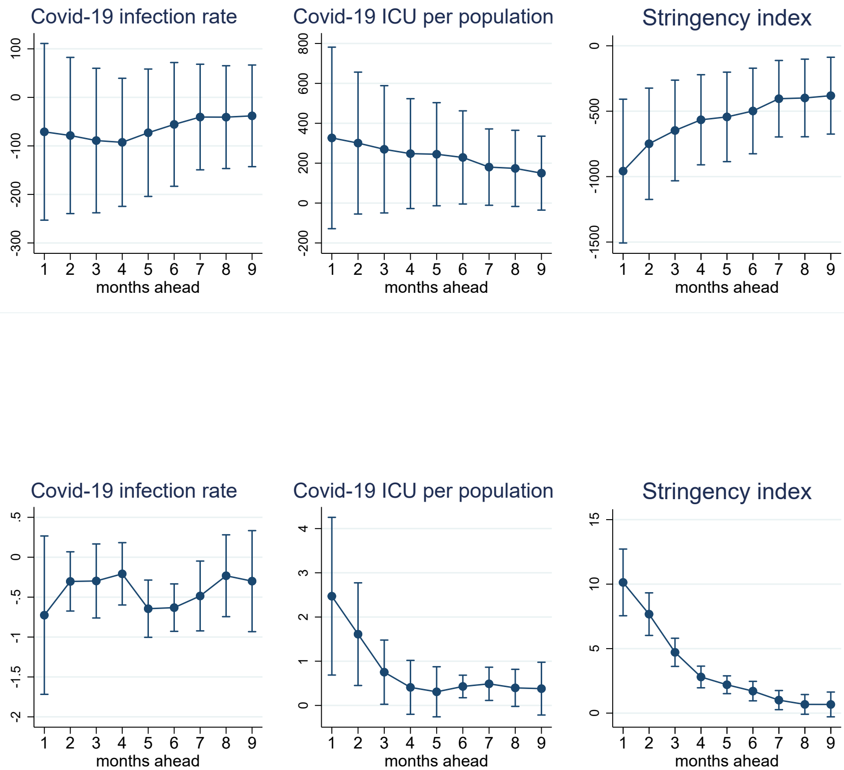

Note: The graph shows the coefficients of the regression of cancellation rates on sanitary variables and the stringency index for different time horizons. Regressors are standardized. Confidence intervals are at 10% level.

stringency index by 28 points, like that occurred on average in March 2020, is associated with a reduction on impact of about 17% and 7% for one and 9-months ahead supply, respectively. Our analysis further suggests that these effects are more related to domestic restrictions rather than international ones. This result is possibly due to the fact that domestic restrictions apply to both domestic tourists, which on average represent about half of tourist presences in the countries we consider,[22] , and international tourists, who take into account not only the difficulty of accessing a given country but also the possibility of visiting touristic attractions once they have crossed the border.[23]

By considering the combined impact of infection rates and ICU occupancy, we find a strong effect of the pandemic also on the demand for vacation rentals, with an increase in the cancel-

[22]Source: Eurostat, nights spent at tourist accommodation establishments by country/world region of residence of the tourist in 2019.

[23]Given the high covariance between STR_{it}^{INT} and STR_{it}^{DOM}, these two components are less significant when taken in isolation. Results about the separate impact of domestic and international restrictions are available upon request.

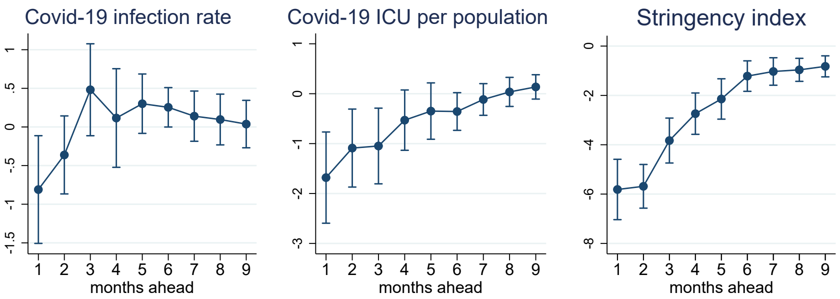

Note: The graph shows the coefficients of the regression of booking rates on sanitary variables and
the stringency index for different time horizons. Regressors are standardized. Confidence intervals

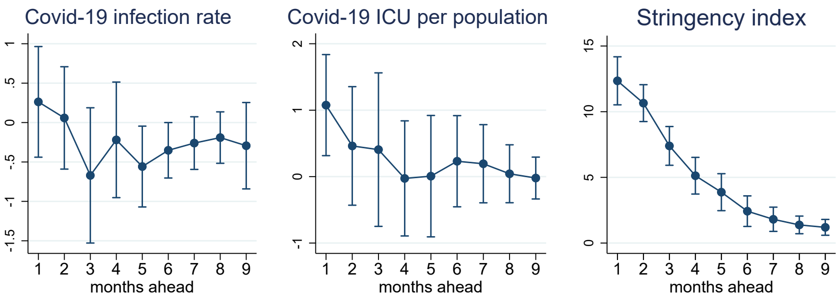

Note: The graph shows the coefficients of the regression of vacancy rates on sanitary variables and
the stringency index for different time horizons. Regressors are standardized. Confidence intervals

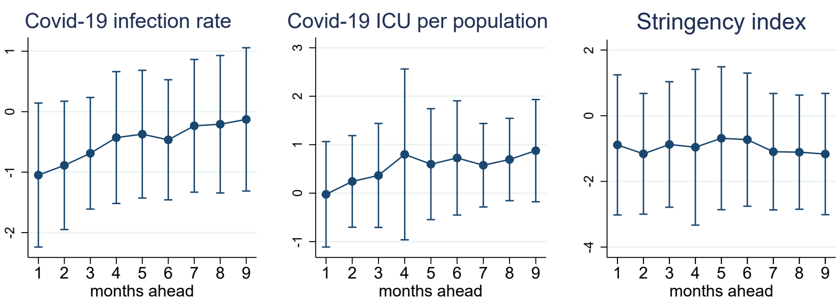

Note: The graph shows the coefficients of the regression of median prices on sanitary variables and the stringency index for different time horizons. Regressors are standardized. Confidence intervals are at 10% level.

Covid Economics 68, 10 February 2021: 115-138

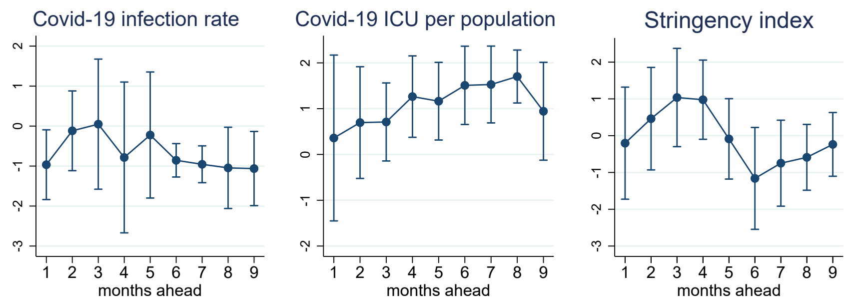

Note: The graph shows the coefficients of the regression of the median hedonic price index on sanitary variables and the stringency index for different time horizons. Regressors are standardized. Confidence intervals are at 10% level.

Covid Economics 68, 10 February 2021: 115-138

lation rates and a decrease in new bookings (Figures 11-12). The stringency index and ICU occupancy both play an important role; COVID-19 infections, instead, have a negative impact only on new bookings in the very short term. The impact of the stringency index is quite persistent: for the booking rate it remains significant over all time horizons, while cancellations are affected up to 7 months ahead. The impact of health conditions is more short-lived, vanishing for horizons longer than one quarter. To provide an example, consider what happened in March 2020. The one-month ahead cancellation rate increased by about 35 percentage points (from 7 to 43%); according to our estimates, considering only the variables of interest, about 12 points out of 35 were due to government restrictions and 2 points to worsening health conditions. In the same period, the one-month-ahead booking rate halved (to 16%): the rise in the stringency index explains about 6 points of this drop, the variation in ICU occupancy about 0.8 points and the spread in infection another half percentage point. As a result of the reduction in the demand for short-term rentals, the vacancy rate is also positively affected by the stringency index and, in the short-term, by the sanitary situation (Figure 13). As it is the case for the supply, also for the demand of short-term rentals the impact of domestic restrictions seems stronger than that of international controls.

Lastly, the effects on prices are mostly not significant because of the large uncertainty surrounding the coefficients estimates (Figures 14-15). This result could be due to a delayed reaction of prices, as apparent from the descriptive statistics (Figures 8-9). Indeed, in the early phase of the pandemic demand reaction was likely to be almost nil, because strict lockdowns prevented any mobility and there was high uncertainty regarding the possibility of travelling in the future. As containment measures were eased and households started to plan their summer holidays, owners probably reacted by reducing prices. However, our regression is not able to capture such delayed effects because it includes only contemporaneous variables.

CEPR PRESS

6 Conclusions

This work exploits data on the market for vacation rentals to investigate the impact of COVID-19 on consumers' behavior and expectations. This market provides an ideal setting for addressing this issue, since we have data on supply, demand and prices at different time horizons, allowing us to distinguish the short and the medium run effects. We find that the pandemic and the related containment measures strongly reduced both the supply of apartments available for rent and households' demand, even up to 9 months ahead. Descriptive evidence further shows a drop in asked prices, albeit with a delay. Overall, our results suggest that the pandemic resembles more a demand shock rather than a supply shock and its effects could be quite persistent.

We think that in the current juncture this type of analyses based on micro-level data are particularly informative about the macroeconomic impact of COVID-19. Looking ahead, the availability of new data on the second wave of contagion, when policy measures have been generally milder and more heterogeneous among countries and sub-national entities, will provide researchers with additional power for identifying the separate impact of health conditions and policy responses.

Covid Economics 68, 10 February 2021: 115-138

References

Andersen, Asger Lau, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan, "Pandemic, Shutdown and Consumer Spending: Lessons from Scandinavian Policy Responses to COVID-19," Papers 2005.04630, arXiv.org May 2020.

Borin, Alessandro, Francesco P. Conteduca, Michele Mancini, and Luca Rossi, "Recenti Sviluppi della Pandemia COVID-19 nel Mondo," COVID-19 Note, Bank of Italy October 2020.

Brinca, Pedro, Joao B. Duarte, and Miguel Faria e Castro, "Measuring Labor Supply and Demand Shocks during COVID-19," Working Papers 2020-011, Federal Reserve Bank of St. Louis May 2020.

Buono, Ines and Francesco P. Conteduca, "Mobility before Government Restrictions in the Wake of COVID-19," COVID-19 Note, Bank of Italy November 2020.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber, "The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending," Covid Economics 20, CEPR May 2020.

Conteduca, Francesco P., Luca Rossi, and Flavia Tonelli, "Fighting COVID-19: Measuring the Restrictiveness of Government Policies," COVID-19 Note, Bank of Italy May 2020.

Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt, "The Macroeconomics of Epidemics," NBER Working Papers 26882, National Bureau of Economic Research, Inc March 2020.

Franks, Jeffrey R., Bertrand Gruss, Carlos Mulas-Granados, Manasa Patnam, and Sebastian Weber, "Exiting from Lockdowns: Early Evidence from Reopenings in Europe," Working Paper, International Monetary Fund October 2020.

Goolsbee, Austan and Chad Syverson, "Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020," *Journal of Public Economics*, 2021, *193* (C).

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning, "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?," NBER Working Papers 26918, National Bureau of Economic Research, Inc April 2020.

Guttentag, Daniel, Stephen Smith, Luke Potwarka, and Mark Havitz, "Why Tourists Choose Airbnb: A Motivation-Based Segmentation Study," *Journal of Travel Research*, 2018, *57* (3), 342–359.

Hu, Maggie and Adrian D. Lee, "Airbnb, COVID-19 Risk and Lockdowns: Local and Global Evidence," 2020.

Covid Economics 68, 10 February 2021: 115-138

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante, "The Great Lockdown and the Big Stimulus: Tracing the Pandemic Possibility Frontier for the U.S," NBER Working Papers 27794, National Bureau of Economic Research, Inc September 2020.

MacDonald, Margaux, Roberto Piazza, and Galen Sher, "A Simple Guide to Estimating the Impact of COVID-19 on Travel and Hospitality Activity," Special Series on Covid-19, International Monetary Fund 2020.

Petherick, Anna, Beatriz Kira, Thomas Hale, Toby Phillips, Samuel Webster, Emily Cameron-Blake, Laura Hallas, Saptarshi Majumdarand Helen Tatlow, Thomas Boby, and Noam Angrist, "Variation in Governement Responses to COVID-19," Working Paper, Blavatnik School of Government December 2020.

Rojas, Felipe Lozano, Xuan Jiang, Laura Montenovo, Kosali I. Simon, Bruce A. Weinberg, and Coady Wing, "Is the Cure Worse than the Problem Itself? Immediate Labor Market Effects of COVID-19 Case Rates and School Closures in the U.S," NBER Working Papers 27127, National Bureau of Economic Research, Inc May 2020.

Covid Economics 68, 10 February 2021: 115-138

CEPR PRESS

A Indicators

Below we describe the methodology underlying the indicators presented in the main text.

Cancellation rate. We can only observe if a given apartment is available or unavailable for booking on a given day. Then, we estimate cancelled reservations based on the number of newly available slots for reservations compared to the previous month. To compute the indicator for the month T, we consider listings present both in the snapshot for month T and the snapshot for month $T-1$. For each listing, we observe the calendar both at time T and $T-1$, and we compute the number of days that at time T are available for reservation out of those that at time $T-1$ were unavailable. We compute for each listing the ratio between the newly available slots for reservations at time T and the number of unavailable slots at time $T-1$ and then take the average ratio across all listings. We consider a ratio instead of the absolute number of cancellations to smooth out seasonal patterns. Finally, we partition future dates on the calendar in different months (30 days period) starting from day T, and we compute the indicator for different time horizons. For example, we compute the share of cancelled bookings five months ahead by considering the number of newly available slots between 120 and 150 days since day T out of those that were unavailable for the same days at time $T-1$.

Booking rate. The methodology underlying the estimation of new bookings is the same used to estimate cancellations, but we consider the number of newly unavailable slots at time T out of those that were available at time $T-1$.

Vacancy rate. We compute the daily share of available slots for reservations across all listings. Then, we partition the calendar in different months (30 days period) starting from day T, and we compute the average for each period.

Occupancy rate. Although we cannot observe ex-post how many days guests have occupied a listed house, we can figure out the evolution of the occupancy rate by using guests' reviews. Airbnb urges guests to give their opinion on their stay within 14 days after checking out; beyond this time, it is no longer possible to provide a review. Since hosts are not professionals, guest reviews are essential for this market to work; otherwise, it would become a "lemon" market. As a result, Airbnb sends several reminders to users to leave their review: according to Airbnb estimates, users review 75% of their stays. Reviews allow tracking the number of stays across different months (with a potential 14 days' lag) in the same city. In particular, we compute for each listing the number of monthly reviews and take the average across all listings.

Unfortunately, we cannot use reviews to compare the occupancy rate across cities because the average period of stay is different, as the reasons why people travel to that city. For example, we expect the average monthly number of reviews per listing to be higher in Venice than in Milan. In the first case, guests are mostly tourists staying in the city for a limited time. In the second case, guests may be workers temporarily moving to the city for job-related reasons, and their period of stay could be longer than for tourists.

Covid Economics 68, 10 February 2021: 115-138

Oxford Stringency Index (sub-components). As described by Petherick et al. (2020), the Stringency Index at a given date is the average of 9 sub-indices, each taking a value between 0 and 100: $I = \frac{1}{9}\sum_{j=1}^{9} I_j$. Among the 9 sub-indices, we separately consider the one referring to international travel controls (STR_{it}^{INT}) and compute the residual one, which captures domestic restrictions (STR_{it}^{DOM}), using the following relationship: $I = \frac{1}{9}\left(\sum_{j=1}^{8} I_j^{DOM} + STR_{it}^{INT}\right)$. By defining $STR^{DOM} = \frac{1}{8}\sum_{j=1}^{8} I_j^{DOM}$, we get: $STR^{DOM} = \frac{9}{8}\left(I - \frac{1}{9}STR_{it}^{INT}\right)$.

Prices. We compute for each listing the average price in euro at different time horizons, as explained above. We consider only prices for dates that are available for booking. Then, we compute the average across listings.

To control for composition effects, we build an additional indicator using the following procedure. For each snapshot, we use only the listings that were present in the previous snapshot. For each listing, we calculate the price variation over all time horizons compared to the previous snapshot. Then, we calculate the median percentage change for each time horizon. Finally, we use these variations to build a price index recursively.

Covid Economics 68, 10 February 2021: 115-138

Covid Economics Issue 68, 10 February 2021

Cash management in times of Covid-19[1]

Fernando Alvarez[2] and David Argente[3]

Date submitted: 1 February 2021; Date accepted: 2 February 2021

The COVID-19 pandemic has decreased households' needs for cash and has forced households to use other payment methods. We argue that these patterns are consistent with a variety of cash management models and, in particular, that these patterns indicate an increase in the transaction cost of adjusting a household's stock of cash. The model allows us to separate the contributions to the observed decline in cash transactions from reduced total spending and from cash management decisions. We use detailed data on ATM cash disbursements in Argentina and the US to estimate how much the pandemic has changed the transaction cost of using cash. This estimation shows that if the intensity of the virus doubles in a county, cash transaction cost increases by approximately 2%. The results from Argentina and the US are quantitatively consistent and imply that, given that cash and other payment methods are imperfectly substitutable, the recent increase in contactless payments due to health risks is not without cost to households.

1 We want to thank Andy Neumeyer for helpful comments and suggestions. We also thank Banco Bilbao Vizcaya Argentaria, S.A © – "BBVA" – for providing us with the data for Argentina and SafeGraph for making their US data freely available to the research community. We also thank Ignacia Cuevas and Agustin Gutierrez for excellent research assistance.

2 University of Chicago.
3 Pennsylvania State University.

Covid Economics 68, 10 February 2021: 139-171

1 Summary and Introduction

The recent COVID-19 pandemic has altered households' use of cash in at least three important ways. First, as expenditures drop, households need less cash. In the US, consumption expenditures dropped almost 9 percentage points from the first to the second quarter of 2020. Furthermore, lockdown policies implemented in many cities to combat the virus most affected the sectors of the economy that traditionally account for much of the cash exchanged in the US, like retail and restaurants.

Second, the pandemic has forced many consumers and businesses to embrace alternative payment methods. In the early stages of the pandemic, several central banks began disinfecting and quarantining paper money with the idea of stopping the spread of the virus. And, though the risk of infection from cash is low, these actions have prompted many establishments to avoid using cash. Even the Centers for Disease Control and Prevention (CDC) recommended the use of touch-free payments whenever possible. In fact, in the Federal Reserve's May 2020 Diary of Consumer Payment Choice, 63% of respondents reported making no in-person payments in the first two months of the pandemic and 28% reported that they had been avoiding the use of cash.[1]

Third, households are holding more cash for precautionary reasons. The amount of currency in circulation in the US is at an all time high and increased more than 5 percentage points from the first to the second quarter of 2020. Respondents in the Diary of Consumer Payment Choice were carrying 17 percent more cash and had stored nearly twice as much cash elsewhere.

We argue that the pandemic has affected the use of cash via channels that are consistent with the structural properties of cash management models (e.g. Baumol (1952), Tobin (1956), Miller and Orr (1966), and Alvarez and Lippi (2009)). In particular, they are consistent with a significant increase in the households' transaction cost of adjusting their stock of cash (e.g. "making trips to the bank" to obtain cash). Cost management theories provide a way to "difference out" the overall decline in the use of cash and the cash-management decisions of households regarding the size and frequency of cash withdrawals. This aspect of the theory allows us to overcome the challenge of measuring cash expenditures.

Using this insight, we quantify the overall change in the cost of cash management during the COVID-19 pandemic using detailed data on ATM disbursements from Argentina and the US. First, we find evidence that the cash-management behavior of households matches the predictions of the theory. We find evidence of an relevant increase in the average size of withdrawals and a relevant decrease in the number of transactions. Second, we find a

[1] Refer to Kim, Kumar and O'Brien (2020) for details on the survey.

Covid Economics 68, 10 February 2021: 139-171

substantial increase in the transaction cost of adjusting a household's stock of cash. If the intensity of the virus doubles in a given county in a two-week period, the transaction costs of using cash will increase by almost 2-3%. Overall, we find that the cost of cash management has increased by approximately 1.5%.[2]

Our results indicate that the pandemic and ensuing lockdown policies have exacted unexpected costs on households, given that both disrupt optimal cash-management decisions. In particular, given the degree of imperfect substitution across payment methods, the prevalence of alternative payment methods during the pandemic is not without cost to households. Our findings also imply an increasing welfare cost of inflation. Inflation erodes a household's stock of cash at a higher speed, forcing more trips to the bank and increasing transaction costs.[3]

Covid Economics 68, 10 February 2021: 139-171

2 Cash Management Models

2.1 A Simple Generalization of Baumol-Tobin

Assets include money M, which does not pay any interest, and those that do pay interest, all of which pay the same nominal interest rate. The opportunity cost of holding money is therefore proportional to the nominal interest rate. Holding cash carries other opportunity costs as well, such as the probability of losing cash per unit of time. We denote R as the total opportunity cost of holding cash per unit of time. Households spend C units of cash per unit of time; C is expressed in real terms. We emphasize that C is a flow per unit of time and it represents only cash expenditures.[4] Households can withdraw more cash by paying a fixed adjustment cost B, measured in real terms. B represents the transaction cost of "making trips to the bank" to adjust the stock of money.

Households make N withdrawals of size W per unit of time. We impose the constraint that $W \times N = C$, meaning that households spend the cash they withdraw. We assume that the average cash balance is $M(W)$, and this function satisfies the following properties: i) $0 \leq M(0) < \infty$, ii) $M'(W) \geq 0$, iii) $WM''(W)/M'(W) = 1$, and iv) M is independent of C and N.

Households choose W and N to minimize the transaction cost per unit of time BN plus

[2]This result is consistent with increased transaction costs in the corporate bond market (i.e. O'Hara and Zhou (2020) and Kargar, Lester, Lindsay, Liu, Weill and Zúñiga (2020)) during the COVID-19 pandemic. Traders have been forced to shift to slower agency trades as a result.

[3]The welfare cost of inflation can be approximated as the area under the money-demand curve derived from the cash management models studied here. See Alvarez, Lippi and Robatto (2019) and Lucas and Nicolini (2015) for more details.

[4]We take the determination of C as given. Clearly, some of the factors that determine C are common to the ones that determine other aspects of cash management. This is the main reason for our review of cash management models.

the opportunity cost of holding cash per unit of time $M(W)R$:

$$\min_{N,W} BN + RM(W) \;\text{ subject to } WN = C \tag{1}$$

where, since the objective function is homogeneous of degree one in (B, R), we can treat B/R as a parameter.

This version of the model is a small generalization of Baumol-Tobin since, in that model, it is assumed that households withdraw cash when their balance hits zero and the cash balance decreases by a constant amount per period so that $M(W) = W/2$. We can allow $M(W) = W/2 + \underline{M}$ where $\underline{M}$ is some minimum amount of cash held by households (e.g. cash hoarding), so long as $\underline{M}$ is independent of C and N.

Replacing $N = C/W$ we can also write:

$$\min_{W} \frac{C}{W}\frac{B}{R} + M(W)$$

whose first-order condition gives:

$$-\frac{C}{(W^*)^2}\frac{B}{R} + M'(W^*) = 0 \;\text{ or }\; \frac{(W^*)^2}{C} = \frac{B/R}{M'(W^*)} \tag{2}$$

where we denote the optimal policy as W^* and we note that $(W^*)^2/C = W^*/N^*$.[5]

Implications of the model. To simplify our discussion of the implications of the model we introduce the following notation: $n = \log N^*$, $w = \log W^*$, $c = \log C$, and $b = \log B/R$ and

$$w_b \equiv \frac{\partial \log W^*}{\partial \log B/R},\; w_c \equiv \frac{\partial \log W^*}{\partial \log C},\; n_b \equiv \frac{\partial \log N^*}{\partial \log B/R},\; n_c \equiv \frac{\partial \log N^*}{\partial \log C}$$

Problem 1 gives the following four implications:

1. The optimal $\frac{W^*}{N^*}$ is independent of C, and is strictly increasing in B/R.
2. The elasticity of W^* with respect to C is 1/2, i.e.: $w_c = 1/2$.
3. The elasticity of W^* with respect to B/R is positive, i.e.: $0 < w_b$.
4. The sum of the elasticities of W^* and N^* with respect to B/R is zero, and the sum of the elasticities of W^* and N^* with respect to C are one, i.e. $w_b + n_b = 0$ and $w_c + n_c = 1$.

[5]There exists a unique solution for W^* since $(W^*)^2/C$ is increasing in W^*, with elasticity 2, and $\frac{B/R}{M'(W^*)}$ has elasticity smaller than 2 with respect to W^*.

Covid Economics 68, 10 February 2021: 139-171

Thus, the elasticity of the ratio W^*/N^* with respect to B/R equals twice the elasticity of W^* with respect to B/R, i.e.

$$\frac{\partial \log(W^*/N^*)}{\partial \log(B/R)} = 2\, w_b > 0 \text{ which is independent of } C\,. \tag{3}$$

We will use this property in our estimation.[6]

Cash Management and Covid-19. We will next use the above implications to estimate how the COVID-19 pandemic has affected cash management. We hypothesize that the intensity of COVID-19 affects both total cash expenditures C and the cost of cash withdrawals B. The pandemic affects C changing total expenditures and the the fraction of total expenditures paid in cash. We might naturally assume that both channels imply that increasing infections in a community will decrease C. Additionally, we hypothesize that increased infection rates will increase the cost of adjusting a cash stock B since adjustments such as ATM withdrawals increase a household's exposure to the virus. We concentrate on measuring this second effect, using the property implied by the theory to separate out the effect of cash expenditures on C. In particular, in the case of the model of this section, the ratio W^*/N^* does *not* depend on C. Under these assumptions, the direct effect of the pandemic on W^*/N^*, divided by twice the elasticity of W^* with respect to B/R, gives the increase in B/R due to the risk of COVID-19.

2.2 Beyond Baumol-Tobin

In this section, we derive similar implications for any model in which the constraint $WN = C$ holds and W and N are functions of B/R and C (e.g. Miller and Orr (1966) and Alvarez and Lippi (2009)). Differentiating the log of $W^*(C, B/R) \times N^*(C, B/R) = C$:

$$w_b + n_b = 0 \text{ and } w_c + n_c = 1 \tag{4}$$

Let $\hat{w}$ and $\hat{n}$ be the total difference in w and n with respect to changes on b and c. Then,

[6] Implication 1 follows directly from the first-order condition (2). Implication 2 follows from noticing that W^2/C does not depend on C and hence the elasticity of W^* with respect to C must be 1/2. Implication 2 follows because W^2/C is independent of C and hence the elasticity of W^* must be $\frac{1}{2}$. Implication 4 follows immediately from differentiating the constraint $W^* \times N^* = C$ with respect to B/R and with respect to C. Baumol-Tobin corresponds to $M(W) = W/2$ and thus $M'\frac{1}{2}$. In this case, the elasticity of W^* with respect to B/R is 1/2 and W^*/N^* is not just increasing in B/R, but exactly twice B/R, i.e the elasticity of W^*/N^* to B/R is one, i.e $w_b - n_b = 1$. Importantly, the change in households' cost of cash management equals one-half the change in B/R.

Covid Economics 68, 10 February 2021: 139-171

Covid Economics 68, 10 February 2021: 139-171

we can write:

$$\hat{w} = w_b \hat{b} + w_c \hat{c} \text{ and } \hat{n} = n_b \hat{b} + n_c \hat{c}$$

and using equation (4) we get:

$$\hat{w} - \hat{n} = 2w_b \hat{b} + [2w_c - 1]\, \hat{c} \tag{5}$$

Note that in Section 2 we use $w_b > 0$ (Implication 3) and $w_c = 1/2$ (Implication 2) to recover that the logarithm of the changes of the ratio W^*/N^*, i.e. $\hat{w} - \hat{n}$, depends only on $\hat{b}$. This is $\hat{w} - \hat{n}$ does *not* depend on $\hat{c}$. However, if the elasticity $w_c \neq 1/2$, then changes on $\hat{c}$ will have an impact on $\hat{w} - \hat{n}$, and $\hat{w} - \hat{n} \neq \hat{b}$. This theoretical property suggests that we should include $\hat{c}$ in our empirical specifications in order to avoid a potential omitted-variable bias.

We finish this section by noting that the models described so far apply to a steady state: they describe households' decisions taken to minimize cost with constant parameters. We use the model to describe data from short periods of time (every two weeks) for comparative statics while ignoring the dynamics, i.e ignoring the effects of past decisions and of expectations of the future values of relevant variables such as B and R. Two reasons justify our choice to focus on statics, which are not different from the reason why the Diamond-Mortensen-Pissarides model is often used in similar fashion. First, we focus on measuring W and N in the data. The state of the households' problem is the stock of cash or cash balances. After any adjustment, this state value is reset, so that there is no memory. Given that adjustments (e.g. visits to a bank branch or ATM) occur approximately twice per month in the data, two-week periods are a good approximation. Second, for essentially the same reasons, expectations about values of B and R far away in the future are not relevant to current decisions, i.e. the optimal decisions in this case are almost identical to those taken in the steady state under the current parameter values.

3 Data

We use data about ATM transactions and withdrawals from Argentina and the US. We see these data sets as complementary. The countries differ greatly in the usage of cash. Cash is the main payment method used in Argentina, both in terms of the number of transactions and the value of payments, whereas cash accounts for only a small share of the total value of payments in the US. The data for Argentina comes from a large bank, covers all provinces in the country, and is collected directly from *ATM transactions*. However, since the data end the second week of August and the virus was most relevant in Argentina during the months of June and July, we have only limited variation with which to study the pandemic's impact on cash management decisions. The data for the US, on the other hand, is collected from *card transactions* and covers most counties in the US. Since the number of US counties far exceeds the number of Argentinean localities and the virus spread much earlier in the US, the data offer more variation for our analysis. Importantly, the US data include total expenditures and card expenditures, which allow us to estimate the changes in the transaction cost of using cash through the lens of models that generalize Baumol-Tobin. Remarkably, results for Argentina and the US are similar in qualitative and quantitative terms.[7]

3.1 ATM Data: Argentina

We use proprietary data of ATM transactions from Banco Bilbao Vizcaya Argentaria (BBVA) Argentina. BBVA is the third-largest private financial institution in Argentina. The data include all transactions at BBVA's ATMs, including those from clients and non-clients. The data set includes information on the number of transactions and size of withdrawals at the branch level. The data set includes information from 267 branches, 139 localities, and 24 provinces. We also obtain information on the daily COVID-19 cases and deaths from COVIDSTATS. The source of the information is the Ministry of Health of Argentina. The data begin the day the first case was confirmed in the country, March 3rd, 2020.

Table B1 shows daily averages of our main variables at the locality level. The average locality in our data has 627.95 (std. 467.02) ATM transactions per day. The size of the disbursement per transaction is 92.61 USD (std. 14.26). The table also reports the average changes in confirmed cases and deaths in each 14-day period. Over our sample period, the average locality saw an increase of approximately 126.87 new confirmed cases every two weeks.

3.2 Card Data: United States

Our data on ATM withdrawals come from Facteus, a provider of financial data for business analytics. The data set contains information on the total expenditures, total number of transactions, and total number of cards, at the zip-code level and with daily frequency. Approximately 10 million debit cards are included. The data set begins in 2017 and ends in the first week of July 2020. It contains information of about 32,285 zip codes out of which 28,104 saw at least one ATM transaction in 2020.[8] The debit cards in the Facteus

[7]Although we do not have information on the size of withdrawals, in the Appendix we include evidence from ATM transactions in Mexico, which is also consistent with the evidence from the US and Argentina.

[8]Figure A1 shows that the data cover approximately 3,199 counties in the US.

Covid Economics 68, 10 February 2021: 139-171

panel are issued by "challenger banks," which are newer banks that tend to serve underbanked consumers, payroll cards issued by employers for direct debit of wages to employees, government cards issued to access funds from garnished wages, and general-purpose debit cards that can be loaded with cash deposits or via direct deposit and can be used at ATMs to withdraw cash. The cardholders whose transactions are in the data tend to come from the middle- and lower-income brackets, a segment of the population that is both more likely affected by COVID-19 financially and more likely to make cash payments.[9]

The data set includes information of more than 200 Merchant Category Codes (MCCs), which correspond to the MCC standard as maintained by Visa and Mastercard. Every transaction processed by the card networks is assigned a MCC, which is a four-digit number that denotes the type of business providing a service or selling merchandise. MCCs determine whether a business transaction needs to be reported to the IRS and the percentage of each transaction a business needs to pay to the credit-card processor. To select records of households' cash disbursements, we use MCC 6011 ("ATM Cash – Disbursements"), which include cash disbursements at automated teller machines (ATMs) owned, leased, controlled, or sponsored by banks, savings and loans, thrifts, and credit unions, including face-to-face transactions.[10]

We also use a database of daily cumulative counts of coronavirus cases and deaths collected by the New York Times. The data set begins with the first reported coronavirus case in Washington State on January 21, 2020, and has been compiled from state and local governments and health departments. Since the data is aggregated at the county level, we aggregate the zip codes of the Facteus panel to the county level using the U.S. Department of Housing and Urban Development (HUD) United States Postal Service ZIP Code Crosswalk Files.[11]

Table A1 shows summary statistics of our main variables at the county level. The average

[9]The data set offers better coverage of small and mid-sized counties. Figure A2 in the Appendix shows that, as the size of the county increases both in terms of total income and population, the coverage of the data, measured as the share of total expenditures covered in the data relative to the income of the county reported to the IRS, decreases.

[10]The data also include MCC 6010 ("Manual Cash Disbursements") , which points to face-to-face cash disbursements at financial institutions. We focus on ATM cash disbursement for two reasons. First, given the sources of the cards included in the data, the coverage of face-to-face cash disbursements is very limited. For example, in the data, debit cards come primarily from challenger banks, which are typically mobile-only banks with no physical branches. As a result, despite the fact that these cards are typically the primary card for the cardholders who own them, over-the-counter cash transactions with these cards are not common. Also, face-to-face transactions are likely to be more affected by temporary closures of bank branches instead of changes in cash-management decisions. Nonetheless, Table A5 shows our main results are robust to including these transactions.

[11]Since zip codes typically overlap with many counties, we use the ratio of residential addresses in the zip-county to the total number of residential addresses in the entire zip code to proportionally distribute the total transactions and disbursements of a zip code to the proper county.

Covid Economics 68, 10 February 2021: 139-171

Covid Economics 68, 10 February 2021: 139-171

county in our data has 8.64 (std. 10.75) total ATM transactions per day. The distribution of transactions is right-skewed; the median county has substantially fewer transactions (4.73) than the average county. The average disbursement per transaction in the average county is 149.97 USD (std. 26.3), which is close to the average value of ATM cash withdrawals of 156 USD reported in the 2019 Federal Reserve Payments Study. The share of cash expenditures, measured as the ratio of total ATM disbursements to total expenditures, is 0.13. The share of transactions at ATMs relative to total transactions is 0.03; transactions that do not include ATM disbursements are small in size. Indeed, the average transaction in the average county is approximately 37.78 USD. The table also reports the average changes in the confirmed cases and deaths in a 14-day period. Over our sample period, the average county recorded approximately 64 new confirmed cases every two weeks.

Panel (a) of Figure 2 shows the relationship between the share of cash expenditures and the income per capita of each county. The panel shows a negative relationship; counties with higher income per capita have a lower share of cash expenditures in our data. This is consistent with Kumar, Maktabi and O'Brien (2018) who show using the 2018 Diary of Consumer Payment Choice (DCPC) that for individuals in households below the median of the income distribution, cash is the most common form of payment and that, as income rises, other payment methods replace cash as the most commonly used payment instrument.[12]

4 Empirical Strategy

4.1 COVID index

We begin by defining a measure that summarizes the intensity of the COVID-19 pandemic. The total number of confirmed cases or the total number of deaths would be natural choices, but both present measurement challenges. At the local level, the total counts of cases and deaths are often updated when local governments correct errors or when they relocate cases to other regions.[13] Furthermore, accurate measurements of the total number of cases depend on the amount and accuracy of the testing taking place in the country or region. Moreover, the total number of deaths registered may be subject to large percentage fluctuations because the underlying number of deaths in a region may be small. To alleviate these concerns we define: COVID $\text{index}_{it} \equiv (Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ stands for the total confirmed

[12]Cash is still the most frequently used payment instrument, representing 30 percent of all transactions. It is used, however, predominantly for small-value purchases; its share of value is approximately 8-9%.

[13]For Argentina we mainly use localities, except for the Autonomous City of Buenos Aires where we can obtain ATM transactions and COVID-19 information at the neighborhood ("barrio") level. For the US we focus on counties.

CEPR PRESS

cases in region i over the last 14 days and $Deaths_{it}$ stands for the total confirmed deaths. We choose the 14-day difference to increase the accuracy of the measurements and because it is the length of the observation period for people who have been exposed to the virus.

We also construct a leave-out COVID index, which we use as an instrument to further alleviate concerns around classical measurement errors. The leave-out COVID index is constructed using $\overline{Cases}_{it} = \sum_{j \neq i} \omega_{ij} Cases_{jt}$ and $\overline{Deaths}_{it} = \sum_{j \neq i} \omega_j Deaths_{jt}$, where ω_j represents the share of workers commuting to county i from county j normalized so that $\sum_{j \neq i} \omega_{ij} = 1$.[14]

4.2 Motivating Facts: Argentina

The left panel of Figure 1 shows the evolution of both total transactions and the average size of withdrawals. The dashed vertical line marks the day of the first positive case of COVID-19 in Argentina. The panel shows that during the pandemic the average size of a withdrawal has increased considerably and the number of ATM transactions has decreased.

Viewed through the lens of the Baumol-Tobin model, the pandemic has increased households' transaction cost of adjusting their stock of cash (e.g. "making trips to the bank" to obtain cash). In the theory, this can be denoted as B/R. Since W/N, the ratio of the size of each withdrawal to the number of withdrawals per unit of time, is strictly increasing in B/R and is observable in our data, we study the relationship between this ratio and COVID index. The right panel of Figure 1 shows that, in the time-series, the households' transaction cost of adjusting their stock of cash has increased and that the timing of this change coincides with the increasing spread of COVID-19.

[14] We obtain the commuting flows for each county from the 2011-2015 American Community Survey (ACS). The ACS asks respondents about their primary workplace location. When information about workers' residence location and workplace location are coupled, a commuting flow is generated.

Covid Economics 68, 10 February 2021: 139-171

Covid Economics 68, 10 February 2021: 139-171

Figure 1: COVID-19 and the Use of Cash: Argentina

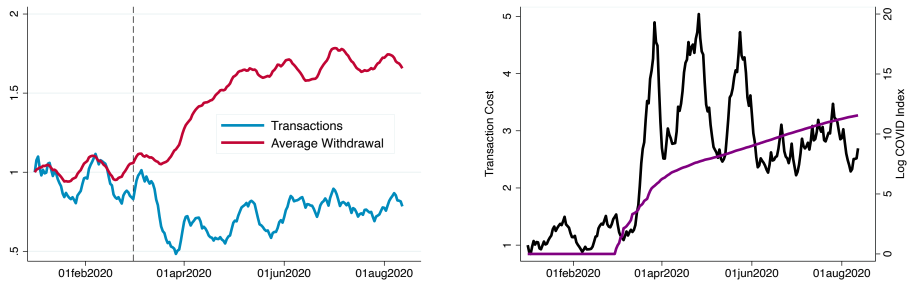

(a) Transactions and Withdrawals (b) Transaction Costs and COVID-19 Cases

Note: Panel (a) shows the evolution of both total transactions and the average withdrawal size in Argentina. Both are two-week moving averages normalized to 1 on the first day of 2020. The dashed vertical line marks the day of the first positive case of COVID-19 in Argentina. Panel (b) shows the relationship between the transaction cost of adjusting the stock of cash (W/N) and the logarithm of COVID index (i.e. COVID index $= (Cases)^{1/2}(Deaths)^{1/2}$). The transaction cost is approximated using the ratio of the average size of withdrawals and the total ATM transactions. Both are two-week moving averages normalized to 1 on the first day of 2020.

4.3 Motivating Facts: United States

In the US, the pandemic has decreased total spending significantly, which in turn has decreased the need for cash. Panel (b) of Figure 2 illustrates this pattern with the relationship between our COVID index and the total expenditures in each county and two-week period. After controlling for county- and period-fixed effects, we can see that the intensity of the pandemic and total expenditures have a strong negative relationship.

The pandemic has also shifted transactions from cash to alternative methods of payment. This is depicted in Panel (c) of Figure 2 which shows the relation between the cash-credit share and COVID index, where the share of cash expenditures is measured using ATM disbursements and the credit share covers every other payment using a card. The figure shows that, as the pandemic worsens, households prefer to use cards instead of cash, which can be the result of households making more online purchases or simply following the CDC recommendation to use touch-free payment methods.[15]

[15] Households have also been hoarding cash for precautionary reasons. Unfortunately, we are unable to verify this pattern directly from our data, since we can only observe ATM transactions and disbursements. Nonetheless, both the Federal Reserve System and Central Bank of Argentina report that the currency in circulation during the pandemic has surged to an all time high.

Covid Economics 68, 10 February 2021: 139-171

Figure 2: COVID-19 and the Use of Cash: US

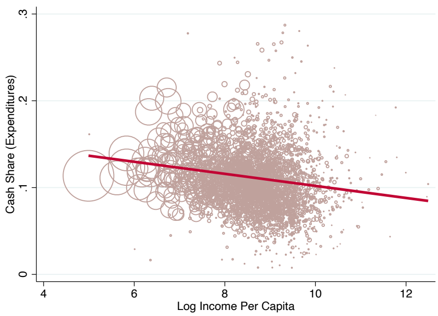

(a) Share of Cash Expenditures and Income

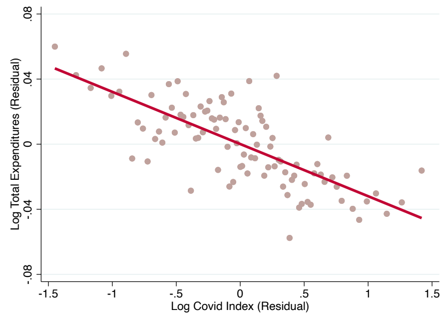

(b) Total Expenditures

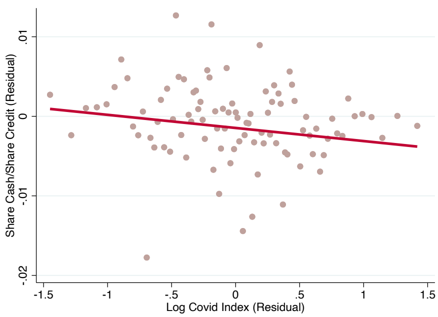

(c) Cash-Credit

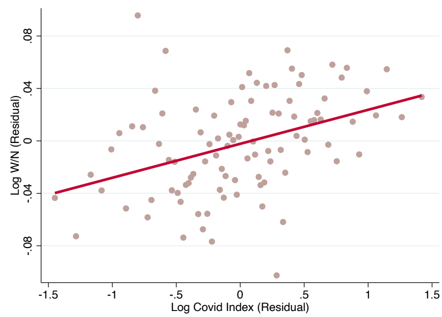

(d) Transaction Cost

Note: Panel (a) shows the average share of cash expenditures and the income per capita at the county level. The share of cash expenditures is computed by averaging across 2017-2019, the years before the virus outbreak. Income is measured using individual income tax returns (Forms 1040) filed with the Internal Revenue Service (IRS) during the 12-month period, January 1, 2017 to December 31, 2017. The population totals come from the US Census. Panel (b) shows the relationship between total spending and the COVID index. Panel (c) shows the relationship between the ratio of the share of spending in cash (ATM disbursements) and the share of spending in card and the COVID index. Panel (d) shows the relationship between the transaction cost of adjusting the stock of cash (W/N) and the COVID index. The transaction cost is approximated using the ratio of the daily average size of withdrawals and the daily average of total ATM transactions for each county at the bi-weekly level. The variables in panels (a)-(d) are plotted after controlling for county and time effects. COVID index$_{it} = (Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ are the total confirmed cases in the county over the last 14 days and $Deaths_{it}$ are the total confirmed deaths over the last 14 days in county i and period t.

Panel (d) of Figure 2 shows a strong positive relationship between the logarithm of W/N and the COVID index after controlling for county- and time-differentiated effects, suggesting that the adjustment cost of withdrawing more cash has gone up during the pandemic. Note, however, that in our generalization of Baumol-Tobin, W/N is not independent of C, the units paid in cash per unit of time, which can be affected in the pandemic by changes in total expenditures as well as the cash-credit substitution depicted in Panel (c). This information is available in the US data. Thus, the next section will explore this relationship using a reduced-form approach in order to isolate the impact of COVID-19 infections on the cost of cash adjustment B/R.

4.4 Transaction Cost of Obtaining Cash

To study the effect of COVID-19 intensity on the fixed cost of obtaining cash, we use the following specification

$$\ln Y_{it} = \alpha + \beta \ln \text{COVID index}_{it} + \theta \ln C_{it} + \lambda_i + \theta_t + \epsilon_{it} \tag{6}$$

where Y_{it} represents dependent variables such as the number of ATM withdrawals, N_{it}, the average size of cash withdrawals, W_{it}, or the transaction cost of adjusting a stock of cash, $\frac{W_{it}}{N_{it}}$. C_{it} are total cash expenditures approximated using all ATM disbursements in county i at time t. All our specifications include region effects, λ_i, and time effects, θ_t. Since the theoretical results that motivate this specification apply to the steady state, we focus on two-week periods.[16] Moreover, given that the error term could be both serially and cross-sectionally correlated, we use Driscoll and Kraay standard errors.[17]

We begin with the changes in cash management decisions in Argentina. Column (1) in Table 1 indicates that if the intensity of the pandemic doubles, the ratio of W/N increases by approximately 2%. This column corresponds to the Baumol-Tobin case; recall that an implication of this model is that the ratio of W/N does not depend on the total cash expenditures C and is strictly increasing with the transaction cost B/R. Columns (2) and (3) show that, as predicted by the model, an increase in the prevalence of the virus increases the average size of withdrawals and decreases the number of ATM transactions. Columns (4)-(6) show that these results are robust to controlling for the total cash expenditures, C. This case corresponds to the generalization of Baumol-Tobin in which W/N can be recovered only after controlling for C in order to avoid potential omitted-variable issues. These findings are consistent with the predictions of a wider class of cash management models.[18]

[16] Households in the US, for example, withdraw cash from ATMs more than once per month on average (Bagnall et al., 2014).

[17] Driscoll-Kraay standard errors tend to be conservative. Our results remain if we use robust standard errors or if we cluster them at the county level; these results are presented in Table A6.

[18] Table B2 and Table B3 show these findings are robust to using cases or deaths as the dependent variable.

Covid Economics 68, 10 February 2021: 139-171

CEPR PRESS

Covid Economics 68, 10 February 2021: 139-171

Table 1: COVID-19 and the Use of Cash: Argentina

Note: The table reports the estimates of equation (6) for Argentina. The dependent variable in columns (1) and (4) is the transaction cost of adjusting the stock of cash, which is approximated using the ratio of the daily average size of withdrawals and the daily average of the total ATM transactions for each county at the bi-weekly level. The dependent variable in columns (2) and (5) is the average size of withdrawals and in columns (3) and (6) the dependent variable is the total ATM transactions. The independent variable is the logarithm of the COVID index$_{it}$ = $(Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ are the total confirmed cases in the county over the last 14 days and $Deaths_{it}$ are the total confirmed deaths over the last 14 days in county i and period t. In columns (4)-(6) we control for the logarithm of total expenditures paid in cash. We use Driscoll and Kraay standard errors with four lags. All the specifications include locality and time effects.

	(1) Log $\frac{W}{N}$	(2) Log W	(3) Log N	(4) Log $\frac{W}{N}$	(5) Log W	(6) Log N
Log COVID index(t)	0.020***	0.008***	-0.011*	0.018***	0.009***	-0.009***
	(0.005)	(0.002)	(0.005)	(0.004)	(0.002)	(0.002)
Log C(t)				-0.749***	0.126***	0.874***
				(0.068)	(0.034)	(0.034)
Observations	1,867	1,867	1,867	1,867	1,867	1,867
Number of groups	169	169	169	169	169	169
County	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y

Table 2 shows how the transaction cost for cash changed in the US during the pandemic. Column (1) indicates that if the intensity of the pandemic doubles, the ratio W/N increases by 3%. This estimate is remarkably similar to the estimate from Argentina. In column (2), we control for the total cash expenditures, C. In this case, the coefficient of the COVID index decreases but it is still positive and significant.

As discussed above, a potential concern is that the number of cases and number of deaths in a given county are measured with error, thus leading to bias in the coefficient of the COVID index. Columns (3) and (4) address this issue by instrumenting the COVID index with a second measurement of COVID index correlated with the original but with an independent measurement error.[19] In column (3) we instrument the COVID index with a one-period lag. In column (4) we instrument the COVID index with the county-level leave-out COVID index. Consistent with the presence of classical measurement error, the coefficient increases in both cases, even after conditioning for total expenditures paid in cash.[20]

[19] The first stage for all the specifications presented in Table 2 are presented in Table A8.

[20] Table A4 shows that we obtain similar results when, instead of using COVID index, we use the total confirmed cases or the total deaths as dependent variables.

Covid Economics 68, 10 February 2021: 139-171

Table 2: COVID-19 and the Use of Cash: Transaction Cost (W/N)

Note: The table reports the estimates of equation (6). The dependent variable is the transaction cost of adjusting the stock of cash, which is approximated using the ratio of the daily average size of withdrawals and the daily average of the total ATM transactions for each county at the bi-weekly level. The independent variable is the logarithm of the COVID index$_{it}$ = $(Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ are the total confirmed cases in the county over the last 14 days and $Deaths_{it}$ are the total confirmed deaths over the last 14 days in county i and period t. In column (2) we control for the logarithm of total expenditures paid in cash. In column (3) we instrument the logarithm of the COVID index$_{it}$ with its one-period lagged value. In column (4) we instrument the logarithm of the COVID index$_{it}$ with a leave-out instrument as described in the main text. In column (5) we instrument the logarithm of the total expenditures paid in cash with the logarithm of total expenditures and its lagged value. In column (6) we instrument the logarithm of total expenditures paid in cash with a leave-out instrument of the logarithm of total expenditures and its lagged value. In column (7) we instrument both the COVID index$_{it}$ and the logarithm of total expenditures paid in cash using the leave out instrument and the logarithm of total expenditures and its lagged value. We consider county-two-week pairs with at least 5 ATM transactions and use Driscoll and Kraay standard errors with four lags. All the specifications include county and time effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log COVID index(t)	0.030***	0.015***	0.026**	0.033***	0.008**	0.006*	0.021**
	(0.004)	(0.002)	(0.009)	(0.008)	(0.003)	(0.003)	(0.008)
Log C(t)		-0.344***	-0.349***	-0.339***	-0.511***	-0.564***	-0.504***
		(0.020)	(0.025)	(0.026)	(0.039)	(0.052)	(0.041)
Observations	21,009	21,009	17,698	20,914	20,863	20,863	20,856
County	Y	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y	Y

Column (5) addresses the endogeneity of the total cash expenditures, which we must address for two reasons. First, some changes in the transaction cost of cash might not be captured by COVID index, and these could be correlated with the total cash expenditures. This effect would yield biased and inconsistent estimates of θ and β, the coefficient of the COVID index. We address this potential bias by instrumenting the total expenditures in cash with the total expenditures, E, and its lagged value. Our identifying assumption is that the unobserved changes in transaction cost affect cash-credit substitutions, but not total household expenditures. In this case, the total expenditures would affect W/N only through cash expenditures. Column (5) shows that in this case the coefficient of the COVID index reduces but remains positive and significant.

The second concern is that households hoard more cash, not only in response to the overall presence of the virus, but also in response to an increase in the intensity of infection in their respective counties. In this case, total ATM disbursements would include both the cash households spend and the cash that they hoard. Since our measure of total expenditures suffers from the same issue, we address it by instrumenting the total cash expenditures with the county-level leave-out mean of total expenditures (i.e. $\overline{E}_{it} = \sum_{j \neq i} \omega_{ij} E_{jt}$) and its lagged value. The leave-out instrument is correlated with the total expenditures of people living in county i, thus correlated with the total cash expenditures of those households, and it does not include the hoarding behavior of households in county i responding to the pandemic. Column (6) shows that in this case our results are similar to those shown in column (5).

In the last column, we instrument both the COVID index using our leave-out instrument and cash expenditures using total expenditures and its lagged value. We instrument both variables in order to address the measurement error of the COVID-19 variables and the endogeneity of cash expenditures simultaneously. Column (7) shows that if the intensity of the pandemic doubles, the ratio W/N increases by approximately 2%. Overall, we find very consistent results throughout all the specifications. The current pandemic, viewed through the lens of cash management models, has increased the transaction cost of using cash.[21]

Next, we use the average size of withdrawals as a dependent variable. The results are presented in Table A2, which shows that when we control for total cash expenditures, the COVID index has a positive and significant coefficient on the average withdrawal size. Column (2) shows that doubling the intensity of COVID-19 increases the average size of withdrawals by approximately 1.5%. Column (8) shows that when we instrument the COVID index and total cash expenditures, we obtain a similar estimate, an increase of approximately 2.1%. Unsurprisingly, the COVID-19 pandemic also had a significant negative impact on the number of transactions, N. These results are presented in Table A3. Thus, the COVID-19 pandemic has led to an increase in the transaction cost of using cash which, consistent with the prediction of cash management models, has increased the overall size of withdrawals and has decreased the frequency of ATM transactions.[22]

4.5 Cash-Credit Substitution

Lastly, we show that COVID-19 has disrupted households' choices of payment methods in the US. We use as a dependent variable the logarithm of the ratio of expenditures paid in cash and those paid in card, including debit and credit payments. The coefficient of the independent variables in this case combines: i) the impact of changes in the cost of obtaining cash, ii) the impact of changes in the cost of using cash relative to cards, and iii)

[21] Table A5 shows that our results are similar when we include face-to-face cash disbursements in our measures of cash transactions and withdrawals.

[22] In the Appendix we show that the response of ATM transactions to the pandemic is quantitatively similar in Mexico (Table C2). We also use data at the bank-municipality level to show that, consistent with cash management models, branch closures due to COVID-19 have a substantial impact on ATM transactions even after controlling for municipality-time, bank-time, and bank-municipality effects (Table C3).

Covid Economics 68, 10 February 2021: 139-171

CEPR PRESS

the elasticity of substitution between cash and cards. Column (1) in Table 3 shows that an increase in the prevalence of the virus has led to a decrease in cash payments relative to card payments, which is consistent with increases in the cost of both obtaining and using cash. Columns (2) and (3) show that this result holds if we consider total cases and total deaths separately. Columns (4) and (5) instrument the COVID index with its lagged value and with the leave-out COVID index, respectively, in order to ameliorate measurement error concerns. These columns indicate that if the prevalence of the virus doubles in a county, total cash expenditures relative to card expenditures decrease approximately 2.3-2.5%.

Table 3: COVID-19 and the Use of Cash: Cash-Credit Response

Note: The table reports the estimates of equation (6). The dependent variable is the logarithm of the ratio of expenditures paid in cash and those paid in card. The independent variable in columns (1), (4) and (5) is the logarithm of the COVID index$_{it}$ = $(Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ are the total confirmed cases in the county over the last 14 days and $Deaths_{it}$ are the total confirmed deaths over the last 14 days in county i and period t. In column (2) the independent variable is the logarithm of $Cases_{it}$ and in column (3) is the logarithm of $Deaths_{it}$. In column (4) we instrument the logarithm of the COVID index$_{it}$ with its one-period lagged value. In column (5) we instrument the logarithm of the COVID index$_{it}$ with a leave-out instrument as described in the main text. We consider county-two-week pairs with at least 5 ATM transactions and use Driscoll and Kraay standard errors with four lags. All the specifications include county and time effects.

	(1)	(2)	(3)	(4)	(5)
Log COVID index(t)	-0.011***			-0.025*	-0.023**
	(0.003)			(0.013)	(0.008)
Log Cases(t)		-0.007***			
		(0.002)			
Log Deaths(t)			-0.008***		
			(0.002)		
Observations	21,008	21,008	21,008	17,698	20,914
R-squared				-0.001	-0.000
County	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y

In order to interpret this coefficient as the impact of changes in the cost of obtaining or using cash on the relative share of cash and credit expenditures, we must divide the coefficient by the elasticity of substitution between cash and card payments. If cash and credit are perfect substitutes, an increase in the total cost of obtaining and using cash would disrupt households' choice of means of payments drastically. Alternatively, if the elasticity of substitution is very low, changes in the cost of using cash would have only a slight effect on the ratio of cash to card payments. Unfortunately, very few estimates of the elasticity of substitution between cash and cards exist. Alvarez and Argente (2020) estimate the

Covid Economics 68, 10 February 2021: 139-171

elasticity of substitution to be between 3 and 5 using field experiments in Mexico. Borrowing these estimates, if the intensity of the pandemic doubles, cash expenditures relative to credit expenditures decrease between 0.4-0.8 % due to increases in the cost of obtaining and using cash.

5 Conclusion

During the COVID-19 pandemic, households are using less cash because spending has decreased and because they are using alternative payment methods instead. Using a simple generalization of Baumol-Tobin, we show that this behavior is consistent with an increase in the transaction cost of using cash.

We then test these predictions using detailed data about ATM transactions in the US and Argentina. The data confirms the predictions of the model, in that changes in households' cash management decisions are explained by the higher transaction cost of obtaining cash. Our results indicate that the outbreak of the COVID-19 virus along with policies to fight it have significantly disrupted households' decisions as reflected by the number trips to the ATM, the size of ATM withdrawals, and the use of cash relative to other payment methods.

Our findings imply an increasing welfare cost of inflation, given that inflation erodes the stock of cash held by households at a higher speed and forces them to go to the ATM more often. Furthermore, given the imperfect substitutability across payment methods, switching between payment methods carries some cost. Lastly, increases in transaction costs could also have distributional effects. Although low-income households hold a higher fraction of cash than their share in the population, high-income households often hold more cash. Ragot (2014) shows that transaction costs are crucial for an explanation of this large inequality in cash holdings. We leave a detailed study of the distributional effects of increases in the transaction cost of obtaining cash for future research.

Covid Economics 68, 10 February 2021: 139-171

References

Alvarez, F., Argente, D., 2020. Consumer surplus of alternative payment methods: Paying uber with cash .

Alvarez, F., Lippi, F., Robatto, R., 2019. Cost of inflation in inventory theoretical models. Review of Economic Dynamics 32, 206–226.

Alvarez, F.E., Lippi, F., 2009. Financial innovation and the transactions demand for cash. Econometrica 77, 363–402.

Bagnall, J., Bounie, D., Huynh, K.P., Kosse, A., Schmidt, T., Schuh, S., Stix, H., 2014. Consumer cash usage: a cross-country comparison with payment diary survey data. Working Paper Series 1685. European Central Bank.

Baumol, W.J., 1952. The transactions demand for cash: An inventory theoretic model. Quarterly Journal of Economics 66, 545–556.

Kargar, M., Lester, B., Lindsay, D., Liu, S., Weill, P.O., Zúñiga, D., 2020. Corporate bond liquidity during the covid-19 crisis. Covid Economics, Vetted and Real-Time Papers. 27, 31–47.

Kim, L., Kumar, R., O'Brien, S., 2020. Consumer payments and the covid-19 pandemic: A supplement to the 2020 findings from the diary of consumer payment choice .

Kumar, R., Maktabi, T., O'Brien, S., 2018. 2018 findings from the diary of consumer payment choice, in: Federal Reserve Bank of San Francisco.

Lucas, R.E., Nicolini, J.P., 2015. On the stability of money demand. Journal of Monetary Economics 73, 48–65.

Miller, M., Orr, D., 1966. A model of the demand for money by firms. Quarterly Journal of Economics 80, 413–435.

O'Hara, M., Zhou, X.A., 2020. Anatomy of a liquidity crisis: Corporate bonds in the covid-19 crisis. Available at SSRN 3615155 .

Ragot, X., 2014. The case for a financial approach to money demand. Journal of Monetary Economics 62, 94–107.

Tobin, J., 1956. The interest elasticity of transactions demand for money. Review of Economics and Statistics 38, 241–247.

Covid Economics 68, 10 February 2021: 139-171

Covid Economics 68, 10 February 2021: 139-171

APPENDIX

A United States

Figure A1: Share of Cash Expenditures by County

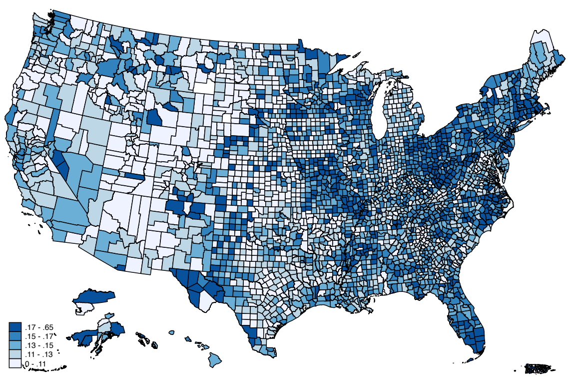

Note: The figure shows the share of cash expenditures (ATM disbursements) over the total expenditures, at the county level. The data include information of 3,199 counties.

Figure A2: Total Expenditures and Income

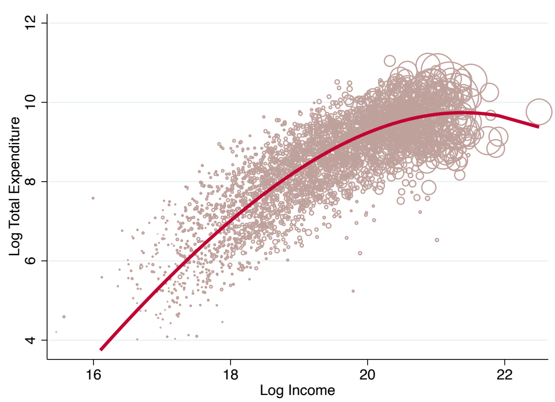

Note: The figure shows the relationship between total spending and total income in a county. Total spending is computed averaging across 2017-2019. Income is measured using individual income tax returns (Forms 1040) filed with the Internal Revenue Service (IRS) between January 1, 2017 and December 31, 2017. The size of the marker indicates the size of the population in each county obtained from the US Census.

CEPR PRESS

Covid Economics 68, 10 February 2021: 139-171

Table A1: Summary Statistics - County Level (US)

Note: The table shows descriptive statistics of the variables of interest at the county level (mean, standard deviation, 25th percentile, median, and 75th percentile) in the year 2020. The share of cash expenditures (Expenditures) is the total cash expenditures over the total expenditures, including ATM disbursements and card transactions. The share of cash expenditures (Transactions) indicates the total ATM transactions over the total transactions. The variables presented are daily averages, except those that relate to the COVID-19 pandemic. "New COVID-19 Cases" indicates the changes in the confirmed cases in a 14-day period at the county level. "New COVID-19 Deaths" indicates the changes in the confirmed deaths in a 14-day period at the county level. The average of these variables is taken after the first case was confirmed on January 21st, 2020.

	(1) Mean	(2) Std. Dev.	(3) Pct. 25	(4) Median	(5) Pct. 75
ATM Transactions	8.64	10.75	1.51	4.73	11.94
ATM Disbursements	1252.66	1485.06	232.92	713.08	1747.41
ATM Disbursements per Transaction	149.97	26.30	136.50	147.08	159.53
Share of cash expenditures (Expenditures)	0.13	0.04	0.11	0.13	0.15
Share of cash expenditures (Transactions)	0.03	0.01	0.03	0.03	0.04
Total Expenditures	9311.48	10460.00	1859.56	5732.00	13260.34
Total Transactions	249.28	278.01	50.46	155.48	352.65
Total Expenditures per Transaction	37.78	6.40	34.58	36.84	39.41
New COVID-19 Cases (two-week)	64.25	300.27	1.71	6.69	27.72
New COVID-19 Deaths (two-week)	2.57	13.92	0.00	0.08	0.69

Table A2: COVID-19 and the Use of Cash: Withdrawals (W)

Note: The table reports the estimates of equation (6). The dependent variable is the daily average size of withdrawals. The independent variable is the logarithm of the COVID index$_{it} = (Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ are the total confirmed cases in the county over the last 14 days and $Deaths_{it}$ are the total confirmed deaths over the last 14 days in county i and period t. In column (2) we control for the logarithm of total expenditures paid in cash. In column (3) we instrument the logarithm of the COVID index$_{it}$ with its one-period lagged value. In column (4) we instrument the logarithm of the COVID index$_{it}$ with a leave-out instrument as described in the main text. In column (5) we instrument the logarithm of total cash expenditures with the logarithm of total expenditures and its lagged value. In column (6) we instrument the logarithm of total cash expenditures with a leave-out instrument of the logarithm of total expenditures and its lagged value. In column (7) we instrument both the COVID index$_{it}$ and the logarithm of total cash expenditures using the leave out instrument and the logarithm of total expenditures and its lagged value. We consider county-two-week pairs with at least 5 ATM transactions and use Driscoll and Kraay standard errors with four lags. All the specifications include county and time effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log COVID index(t)	-0.006***	0.008***	0.013**	0.017***	0.004**	0.003*	0.010**
	(0.002)	(0.001)	(0.005)	(0.004)	(0.001)	(0.002)	(0.004)
Log C(t)		0.328***	0.326***	0.331***	0.244***	0.218***	0.248***
		(0.010)	(0.012)	(0.013)	(0.019)	(0.026)	(0.021)
Observations	21,009	21,009	17,698	20,914	20,863	20,863	20,856
R-squared			0.401	0.395	0.372	0.353	0.373
County	Y	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y	Y

Covid Economics 68, 10 February 2021: 139-171

Table A3: COVID-19 and the Use of Cash: Transactions (N)

Note: The table reports the estimates of equation (6). The dependent variable is the daily average of the total ATM transactions for each county at the bi-weekly level. The independent variable is the logarithm of the COVID index$_{it}$ = $(Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ are the total confirmed cases in the county over the last 14 days and $Deaths_{it}$ are the total confirmed deaths over the last 14 days in county i and period t. In column (2) we control for the logarithm of total expenditures paid in cash. In column (3) we instrument the logarithm of the COVID index$_{it}$ with its one-period lagged value. In column (4) we instrument the logarithm of the COVID index$_{it}$ with a leave-out instrument as described in the main text. In column (5) we instrument the logarithm of total expenditures paid in cash with the logarithm of total expenditures and its lagged value. In column (6) we instrument the logarithm of total expenditures paid in cash with a leave-out instrument of the logarithm of total expenditures and its lagged value. In column (7) we instrument both the COVID index$_{it}$ and the logarithm of total expenditures paid in cash using the leave out instrument and the logarithm of total expenditures and its lagged value. We consider county-two-week pairs with at least 5 ATM transactions and use Driscoll and Kraay standard errors with four lags. All the specifications include county and time effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log COVID index(t)	-0.036***	-0.008***	-0.013**	-0.017***	-0.004**	-0.003*	-0.010**
	(0.005)	(0.001)	(0.005)	(0.004)	(0.001)	(0.002)	(0.004)
Log C(t)		0.672***	0.674***	0.669***	0.756***	0.782***	0.752***
		(0.010)	(0.012)	(0.013)	(0.019)	(0.026)	(0.021)
Observations	21,009	21,009	17,698	20,914	20,863	20,863	20,856
R-squared			0.748	0.736	0.727	0.719	0.728
County	Y	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y	Y

Covid Economics 68, 10 February 2021: 139-171

CEPR PRESS

Table A4: Cases, Deaths, and the Use of Cash: Transaction Cost (W/N)

Note: The table reports the estimates of equation (6). The dependent variable is the transaction cost of adjusting the stock of cash, which is approximated using the ratio of the daily average size of withdrawals and the daily average of the total ATM transactions for each county at the bi-weekly level. The independent variable in columns (1) to (4) is the total confirmed cases in the county over the last 14 days and in columns (5) to (8) is the total confirmed deaths over the last 14 days in county i and period t. In columns (2) and (6) we control for the logarithm of total expenditures paid in cash. In columns (3) and (7) we instrument the logarithm of the total confirmed cases and the total deaths with their respective lag variables. In columns (4) and (8) we instrument the logarithm of the total confirmed cases and the total deaths with a leave-out instrument as described in the main text. We consider county-biweek pairs with at least 5 ATM transactions and use Driscoll and Kraay standard errors with four lags. All the specifications include county and time effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Cases(t)	0.016***	0.008***	0.010*	0.010				
	(0.003)	(0.002)	(0.005)	(0.009)				
Log Deaths(t)					0.021***	0.016***	0.030**	0.049***
					(0.003)	(0.003)	(0.012)	(0.011)
Log C(t)		-0.264***	-0.268***	-0.262***		-0.265***	-0.268***	-0.261***
		(0.032)	(0.040)	(0.039)		(0.032)	(0.040)	(0.038)
Observations	21,165	21,165	17,820	21,067	21,165	21,165	17,820	21,069
County	Y	Y	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y	Y	Y

Covid Economics 68, 10 February 2021: 139-171

CEPR PRESS

Covid Economics 68, 10 February 2021: 139-171

Table A5: COVID-19 and the Use of Cash: Transaction Cost (W/N) - All Cash Transactions

Note: The table reports the estimates of equation (6). The estimates include MCC 6010 ("Manual Cash Disbursements") , which includes face-to-face cash disbursements at financial institutions. The dependent variable is the transaction cost of adjusting the stock of cash, which is approximated using the ratio of the daily average of withdrawals and the daily average of the total ATM transactions for each county at the bi-weekly level. The independent variable is the logarithm of the COVID index$_{it}$ = $(Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ are the total confirmed cases in the county over the last 14 days and $Deaths_{it}$ are the total confirmed deaths over the last 14 days in county i and period t. In column (2) we control for the logarithm of total expenditures paid in cash. In column (3) we instrument the logarithm of the COVID index$_{it}$ with its one-period-lagged value. In column (4) we instrument the logarithm of the COVID index$_{it}$ with a leave-out instrument as described in the main text. In column (5) we instrument the logarithm of total expenditures paid in cash with the logarithm of total expenditures and its lagged value. In column (6) we instrument the logarithm of total expenditures paid in cash with a leave-out instrument of the logarithm of total expenditures and its lagged value. In column (7) we instrument both the COVID index$_{it}$ and the logarithm of total expenditures paid in cash using the leave out instrument and the logarithm of total expenditures and its lagged value. We consider county-two-week pairs with at least 5 ATM transactions and use Driscoll and Kraay standard errors with four lags. All the specifications include county and time effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log COVID index(t)	0.027***	0.016***	0.022*	0.031***	0.008**	0.005	0.017*
	(0.004)	(0.003)	(0.010)	(0.009)	(0.003)	(0.004)	(0.009)
Log C(t)		-0.263***	-0.266***	-0.258***	-0.460***	-0.539***	-0.454***
		(0.032)	(0.040)	(0.039)	(0.058)	(0.075)	(0.060)
Observations	21,165	21,165	17,820	21,067	21,016	21,016	21,009
R-squared			0.089	0.083	0.041	-0.003	0.042
County	Y	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y	Y

CEPR PRESS

Table A6: COVID-19 and the Use of Cash: Transaction Cost (W/N) - Alternative Standard Errors

Note: The table reports the estimates of equation (6). The dependent variable is the transaction cost of adjusting the stock of cash, which is approximated using the ratio of the daily average size of withdrawals and the daily average of the total ATM transactions for each county at the bi-weekly level. The independent variable is the logarithm of the COVID index$_{it}$ = $(Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ are the total confirmed cases in the county over the last 14 days and $Deaths_{it}$ are the total confirmed deaths over the last 14 days in county i and period t. In column (2) we control for the logarithm of total expenditures paid in cash. In column (3) we instrument the logarithm of the COVID index$_{it}$ with its one-period lagged value. In column (4) we instrument the logarithm of the COVID index$_{it}$ with a leave-out instrument as described in the main text. In column (5) we instrument the logarithm of total expenditures paid in cash with the logarithm of total expenditures and its lagged value. In column (6) we instrument the logarithm of total expenditures paid in cash with a leave-out instrument of the logarithm of total expenditures and its lagged value. In column (7) we instrument both the COVID index$_{it}$ and the logarithm of total expenditures paid in cash using the leave out instrument and the logarithm of total expenditures and its lagged value. We consider county-two-week pairs with at least 5 ATM transactions and cluster the standard errors at the county level. All the specifications include county and time effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log COVID index(t)	0.027***	0.016***	0.022***	0.031***	0.008***	0.005*	0.017***
	(0.003)	(0.003)	(0.006)	(0.006)	(0.003)	(0.003)	(0.006)
Log C(t)		-0.263***	-0.266***	-0.258***	-0.460***	-0.539***	-0.454***
		(0.012)	(0.013)	(0.012)	(0.016)	(0.020)	(0.016)
Observations	21,074	21,074	17,820	21,067	21,016	21,016	21,009
County	Y	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y	Y

Covid Economics 68, 10 February 2021: 139-171

CEPR PRESS

Covid Economics 68, 10 February 2021: 139-171

Table A7: COVID-19 and the Use of Cash: Transaction Cost (W/N) - Monthly

Note: The table reports the estimates of equation (6). The dependent variable is the transaction cost of adjusting the stock of cash, which is approximated using the ratio of the daily average size of withdrawals and the daily average of the total ATM transactions for each county at the monthly level. The independent variable is the logarithm of the COVID index$_{it}$ = $(Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ are the total confirmed cases in the county over the last month and $Deaths_{it}$ are the total confirmed deaths over the last month in county i and period t. In column (2) we control for the logarithm of total expenditures paid in cash. In column (3) we instrument the logarithm of the COVID index$_{it}$ with its one-period lagged value. In column (4) we instrument the logarithm of the COVID index$_{it}$ with a leave-out instrument as described in the main text. In column (5) we instrument the logarithm of total expenditures paid in cash with the logarithm of total expenditures and its lagged value. In column (6) we instrument the logarithm of total expenditures paid in cash with a leave-out instrument of the logarithm of total expenditures and its lagged value. In column (7) we instrument both the COVID index$_{it}$ and the logarithm of total expenditures paid in cash using the leave out instrument and the logarithm of total expenditures and its lagged value. We consider county-monthly pairs with at least 5 ATM transactions and cluster the standard errors at the county level. All the specifications include county and time effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log COVID index(t)	0.025***	0.013***	0.088***	0.013**	0.009***	0.008***	0.008
	(0.003)	(0.003)	(0.018)	(0.005)	(0.003)	(0.003)	(0.005)
Log C(t)		-0.435***	-0.419***	-0.435***	-0.584***	-0.636***	-0.584***
		(0.013)	(0.015)	(0.013)	(0.019)	(0.024)	(0.020)
Observations	11,860	11,860	8,976	11,860	11,853	11,853	11,853
R-squared	0.975	0.981	0.167	0.228	0.201	0.180	0.201
County	Y	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y	Y

Covid Economics 68, 10 February 2021: 139-171

Table A8: COVID-19 and the Use of Cash: Transaction Cost (W/N) - First Stage

Note: The tables shows the first-stage regressions of the instrumented specifications in Table 2. In columns (1) and (2) the instrumented variable is the logarithm of the COVID index$_{it}$ = $(Cases_{it})^{1/2}(Deaths_{it})^{1/2}$, where $Cases_{it}$ are the total confirmed cases in the county over the last 14 days and $Deaths_{it}$ are the total confirmed deaths over the last 14 days in county i and period t. In column (1) the instrument is the lagged value of COVID index. In column (2) the instrument is a leave-out instrument of the mean of COVID index at the county level where we use the commuting flows as weights. In both columns we control for total cash expenditures. In columns (3) and (4) the instrumented variable is total cash expenditures. In column (3) the instruments are the logarithm of total expenditures and its lagged value. In column (4) the instrument is a leave-out instrument of the logarithm of total expenditures and its lagged value. In both columns we control for COVID index. The estimates in columns (5) and (6) correspond to those presented in column (8) of Table 2, where we instrument both Log COVID index and Log C(t). In column (5) we instrument COVID index with a leave-out instrument of the mean of COVID index at the county level. In column (6) we instrument Log C(t) with the logarithm of total expenditures and its lagged value. We consider county-two-week pairs with at least 5 ATM transactions and use Driscoll and Kraay standard errors with four lags. All the specifications include county and time effects.

	(1)	(2)	(3)	(4)	(5)	(6)
	Log COVID index(t)		Log C(t)		Log COVID index(t)	Log C(t)
Log COVID index(t-1)	0.474***					
	(0.039)					
Log COVID index(t-1) - IV		0.734***			0.718***	-0.006
		(0.006)			(0.009)	(0.005)
Log C(t)	-0.214***	-0.117**				
	(0.036)	(0.040)				
Log E(t)			1.117***		-0.230***	1.087***
			(0.014)		(0.022)	(0.031)
Log E(t-1)			-0.104***		-0.236**	-0.072
			(0.024)		(0.090)	(0.043)
Log E(t) - IV				0.968***		
				(0.035)		
Log E(t-1) - IV				0.013		
				(0.026)		
Log COVID index(t)			-0.009***	-0.019***		
			(0.002)	(0.004)		
Observations	17,819	21,002	20,951	20,951	20,944	28,068
F-Statistic	105.5	7422.7	2257.8	333.6	1577.6	1413.4
County	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y

B Argentina

Table B1: Summary Statistics - Locality Level (Argentina)

Note: The table shows descriptive statistics of the variables of interest at the locality level (mean, standard deviation, percentile 25th, median, and percentile 75th) in the year 2020. The exchange rate used is the one that prevailed on January 1, 2020 (i.e. 1 Argentine Peso equals 0.01671 United States Dollar); all amounts are expressed in real pesos. The variables presented are daily averages, except those that relate to the COVID-19 pandemic. "New COVID-19 Cases" indicates the changes in the confirmed cases in a 14-day period at the locality level. "New COVID-19 Deaths" indicates the changes in the confirmed deaths in a 14-day period at the locality level. The average of these variables is taken after the first case was confirmed on March 3rd, 2020.

	(1) Mean	(2) Std. Dev.	(3) Pct. 25	(4) Median	(5) Pct. 75
ATM Transactions	627.95	467.02	317.27	467.52	841.22
ATM Disbursements	32903.67	24411.19	18997.03	25111.96	44440.37
ATM Disbursements per Transaction	54.16	8.22	48.65	52.84	60.15
New COVID-19 Cases (two-week)	126.87	193.67	2.88	30.97	164.23
New COVID-19 Deaths (two-week)	3.80	6.02	0.06	0.64	6.47

Table B2: Cases and the Use of Cash: Argentina

Note: The table reports the estimates of equation (6) for Argentina. The dependent variable in columns (1) and (4) is the transaction cost of adjusting the stock of cash, which is approximated using the ratio of the daily average size of withdrawals and the daily average of the total ATM transactions for each county at the bi-weekly level. The dependent variable in columns (2) and (5) is the average size of withdrawals and in columns (3) and (6) the dependent variable is the total ATM transactions. The independent variable is the logarithm of the total confirmed cases over the last 14 days in locality i and period t. In columns (4)-(6) we control for the logarithm of total expenditures paid in cash. We use Driscoll and Kraay standard errors with four lags. All the specifications include locality and time effects.

	(1) Log $\frac{W}{N}$	(2) Log W	(3) Log N	(4) Log $\frac{W}{N}$	(5) Log W	(6) Log N
Log Cases(t)	0.028***	0.004	-0.023***	0.014***	0.007***	-0.007***
	(0.003)	(0.002)	(0.002)	(0.004)	(0.002)	(0.002)
Log C(t)				-0.721***	0.140***	0.860***
				(0.035)	(0.018)	(0.018)
Observations	691	691	691	691	691	691
Locality	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y

Covid Economics 68, 10 February 2021: 139-171

Covid Economics 68, 10 February 2021: 139-171

Table B3: Deaths and the Use of Cash: Argentina

Note: The table reports the estimates of equation (6) for Argentina. The dependent variable in columns (1) and (4) is the transaction cost of adjusting the stock of cash, which is approximated using the ratio of the daily average size of withdrawals and the daily average of the total ATM transactions for each county at the bi-weekly level. The dependent variable in columns (2) and (5) is the average size of withdrawals and in columns (3) and (6) the dependent variable is the total ATM transactions. The independent variable is the logarithm of the total confirmed deaths over the last 14 days in locality i and period t. In columns (4)-(6) we control for the logarithm of total cash expenditures. We use Driscoll and Kraay standard errors with four lags. All the specifications include locality and time effects.

	(1) Log $\frac{W}{N}$	(2) Log W	(3) Log N	(4) Log $\frac{W}{N}$	(5) Log W	(6) Log N
Log Deaths(t)	0.010**	0.006**	-0.004	0.012**	0.006**	-0.006**
	(0.004)	(0.003)	(0.004)	(0.005)	(0.002)	(0.002)
Log C(t)				-0.752***	0.124***	0.876***
				(0.067)	(0.034)	(0.034)
Observations	1,867	1,867	1,867	1,867	1,867	1,867
Number of groups	169	169	169	169	169	169
County	Y	Y	Y	Y	Y	Y
Time	Y	Y	Y	Y	Y	Y

CEPR PRESS

C Mexico

We use the Financial Inclusion Database (BDIF) from the National Banking and Securities Commission (CNBV). The data consist of monthly data gathered from commercial banks and other financial entities related to financial inclusion. The databases include variables such as bank branches, ATMs, ATM transactions, and debit contracts. Data set is disaggregated at the bank and municipality level and contains information on the number of bank branches that have closed due to the pandemic each time period. The data gathered for this paper corresponds to the period 2011-2020. Since we study the pandemic period, we focus on data from January to August 2020.

The average municipality in our data has 65386 (std. 255850) ATM transactions per month. It also has 4 banks, 5 bank branches, 26 ATMs. The table also reports the average changes in the confirmed cases and deaths in a month. Over our sample period, the average municipality suffered an increase of approximately 64 new confirmed cases per month.

Table C1: Summary Statistics - Municipality Level (Mexico)

Note: The table shows descriptive statistics of the variables of interest at the municipality level (mean, standard deviation, 25th percentile, median, and 75th percentile) in the year 2020. The variables presented are daily averages, except those that relate to the COVID-19 pandemic. "New COVID-19 Cases" indicates the changes in the confirmed cases in a month at the municipality level. "New COVID-19 Deaths" indicates the changes in the confirmed deaths in a month at the municipality level. The average of these variables is taken after the first case was confirmed on February 28th, 2020.

	(1) Mean	(2) Std. Dev.	(3) Pct. 25	(4) Median	(5) Pct. 75
ATM Transactions	65386.14	255850.70	0.00	2635.67	21929.00
Banks	4.00	4.35	1.00	2.00	5.00
ATMs	26.28	117.95	0.00	1.00	7.00
Branches	5.07	19.49	0.00	0.00	2.00
Branches Closed	0.75	3.63	0.00	0.00	0.00
New COVID-19 Cases (Monthly)	63.99	211.57	3.50	9.17	30.05
New COVID-19 Deaths (Monthly)	7.68	26.08	0.50	1.25	3.76

Covid Economics 68, 10 February 2021: 139-171

CEPR PRESS

Covid Economics 68, 10 February 2021: 139-171

Figure C1: COVID-19 and the Use of Cash: Mexico

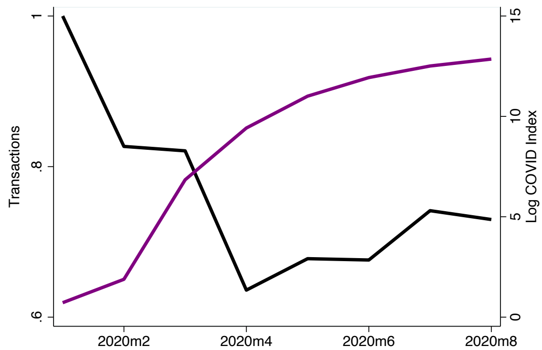

Note: The figure shows the evolution of ATM transactions normalized to 1 on January 2020 (black line) and the logarithm of COVID index (i.e. COVID index $= (Cases)^{1/2}(Deaths)^{1/2}$) in Mexico (purple line).

Table C2: COVID-19 and the Use of Cash: Transactions (N) - Mexico

Note: The table reports the estimates of equation (6). The dependent variable is the logarithm of the total ATM transactions for each bank-municipality at the monthly level. The independent variable in column (1) is the logarithm of the total confirmed cases over the last month in a given municipality and period. In column (2) the independent variable is the total confirmed deaths over the last month. In column (3) the independent variable is COVID index $= (Cases)^{1/2}(Deaths)^{1/2}$, where *Cases* are the total confirmed cases in the municipality over the last month and *Deaths* are the total confirmed deaths over the last month in a given municipality and period. In column (4) the independent variable is the total branches closed due to COVID-19 for a given bank-municipality and period. We use Driscoll and Kraay standard errors. All the specifications include bank-municipality and time effects.

	(1)	(2)	(3)	(4)
Log Cases(t)	-0.013*			
	(0.005)			
Log Deaths(t)		-0.008*		
		(0.004)		
Log COVID index(t)			-0.012*	
			(0.005)	
Log Branches Closed(t)				-0.100*
				(0.043)
Observations	32,167	32,167	32,167	41,629
Bank-Municipality	Y	Y	Y	Y
Time	Y	Y	Y	Y

Covid Economics 68, 10 February 2021: 139-171

Table C3: Branches Closed and the Use of Cash: Transactions (N) - Mexico

Note: The table reports the estimates of equation (6). The dependent variable is the logarithm of the total ATM transactions for each bank-municipality at the monthly level. The independent variable is the total branches closed due to COVID-19 for a given bank-municipality and period. The data is monthly at the bank-municipality level and comes from the National Banking and Securities Commission (CNBV). The standard errors are clustered at the municipality-time level.

	(1)	(2)	(3)	(4)
Log Branches Closed(t)	-0.097***	-0.086***	-0.099***	-0.133***
	(0.015)	(0.015)	(0.016)	(0.019)
Observations	40,935	39,110	40,830	39,005
R-squared	0.976	0.979	0.979	0.982
Bank-Municipality	Y	Y	Y	Y
Time	Y	N	N	N
Municipality-Time	N	Y	N	Y
Bank-Time	N	N	Y	Y

COVID-19 and precautionary corporate cash holdings: Evidence from Japan[1]

Tomohito Honda[2] and Iichiro Uesugi[3]

Date submitted: 2 February 2021; Date accepted: 7 February 2021

This study examines how listed firms have managed their cash holdings since the outbreak of the COVID-19 crisis, using quarterly data on publicly-traded firms in Japan. After providing an overview of developments in cash holdings since the start of the crisis, we focus on the precautionary motive for corporate cash holdings and examine the role of firms' cash flow and volatility therein in firms' cash holdings to find the following: (1) corporate cash holdings have increased rather than decreased since the start of the crisis; (2) an increase in firms' cash flow has a positive impact on their cash holdings during normal times, and the sensitivity of cash holdings to cash flows was more pronounced during the first three months of the crisis; (3) firms facing higher sales volatility held more cash in the second three-month period following the start of the crisis; and (4) the cash flow sensitivity of financially constrained firms' cash holdings during the crisis period increased more than that of unconstrained firms. Overall, the COVID-19 crisis has had a substantial impact on corporate cash management strategies and the results are consistent with the precautionary motive theory for cash holdings.

Covid Economics 68, 10 February 2021: 172-204

1 We are grateful to Ralph Paprzycki for valuable comments.
2 Post-doctoral researcher, Institute of Economic Research, Hitotsubashi University.
3 Professor, Institute of Economic Research, Hitotsubashi University.

Covid Economics 68, 10 February 2021: 172-204

1. Introduction

Cash holdings play a critical role in corporate liquidity management and have been the focus of many previous studies. A number of researchers have studied the determinants of firms' cash holdings and have identified two major factors: the transaction motive and the precautionary motive.[1] The transaction motive refers to firms' desire to hold a sufficient amount of cash to pay for transactions without having to incur the cost of converting fixed assets into liquid ones. Meanwhile, the precautionary motive refers to the desire to hold sufficient cash for unexpected contingencies. Firms tend to hold substantial amounts of cash for unexpected funding demands (see, e.g., Almeida et al., 2004; Riddick and Whited, 2009; Duchin et al. 2010). Numerous studies examine situations in which the precautionary motive for cash holdings plays an important role. Opler et al. (1999), for example, argue that the precautionary motive becomes more important when firms' cash flow is subject to greater risk or firms have limited access to external financing. Similarly, modeling firms' demand for liquidity, Almeida et al. (2004) show that financially constrained firms are likely to save a larger amount of their cash flow for precautionary cash holdings than unconstrained firms.

The precautionary motive for corporate cash holdings becomes even more important during times of financial or economic crisis, such as the current crisis brought about by the COVID-19 pandemic. Several studies examine whether the propensity to build up precautionary cash reserves is greater during financial crises than during normal times and find that this is indeed the case. For instance, Sun and Wang (2015), focusing on the impact of the global financial crisis in 2008, find that the cash flow sensitivity of cash was significantly greater during the crisis period. Similarly, investigating the long-term effect of the Asian financial crisis on corporate cash holdings in eight East Asian countries, Song and Lee (2012) find that firms in these countries built up cash holdings

[1] In the literature on corporate finance, there are other motives such as the tax motive and the agency motive. See Bates et al. (2009), who describe these motives in more detail.

following the crisis by decreasing investment. They also show that firms' increased sensitivity to cash flow volatility was one of the main factors for their higher level of cash holdings. Meanwhile, examining the 2008 European financial crisis, Lozano and Yaman (2020) find that the crisis had a positive impact on corporate cash holdings for three years following the crisis. Further, investigating the link between cash flow volatility and cash holdings for constrained firms, they observe that the positive correlation was larger during the crisis than before the crisis.

While these studies focus on the impact of financial crises on corporate cash holdings, to the best of the authors' knowledge, there are no studies to date that examine the impact of the outbreak of the COVID-19 pandemic on precautionary corporate cash holdings.[2] The crisis triggered by the pandemic has several unique features that warrant further research on cash holdings. First, in contrast with periods of financial crisis in the past, the financial sector has remained quite stable since the emergence of the current crisis. For example, in Japan, credit spreads on corporate bonds, which had jumped during the global financial crisis, increased only marginally at the onset of the crisis and have leveled off since then (Bank of Japan, 2020: Chart II-2-11). Various policy measures introduced by the Japanese government and the central bank have contributed to the stability of the financial system thus far. Second, the shock to the real economy has been unprecedented not only in its sheer size but also in the way industries are affected. The initial drop in aggregate output in Japan, for example, was the largest in the past 70 years. Moreover, due to the nature of the shock, the damage was distributed unevenly across industries. Some industries, such as transportation, accommodations and restaurants, and services for individuals incurred massive losses due to the state of emergency declared by the Japanese government made people refrain from going out and eating out, while for other industries,

[2] While there are no studies so far focusing on the impact of the pandemic on precautionary cash holdings, there already are a considerable number of studies on the impact of the COVID-19 pandemic on various other aspects of corporate finance. For example, Francis et al. (2020) examine the impact on firms' capital structure across 31 countries, Acharya and Steffen (2020) examine firm financing through the corporate bond market and existing credit lines, Li et al. (2020) analyze firms' demand for bank liquidity and banks' capacity to supply the liquidity in the first four months of 2020, and De Vito and Gomez (2020) examine how the COVID-19 health crisis could affect the liquidity of listed firms across 26 countries.

such as construction, telecommunications, and business services, the damage was relatively minor.[3] Third, the degree of uncertainty in a variety of areas has risen sharply since the outbreak of the pandemic. For example, looking at various measures for economic uncertainty for Japan, we find that the macroeconomic uncertainty index and the economic policy uncertainty index have reached the highest value in two decades.[4] This increase in economic uncertainty potentially may have led to an increase in demand for precautionary corporate cash holdings.

Against this background, this study examines how the precautionary demand for cash has been affected by the outbreak of the pandemic by focusing on the period from January to June 2020. Given that the outbreak of COVID-19 and government restrictions to restrain it represent a massive external shock to the economy, the pandemic provides an excellent natural experiment to examine whether and how the precautionary motive affects corporate cash holdings.

For our analysis, we employ a sample of 1,773 listed Japanese firms for the period up to the end of the second quarter of 2020. Our observation period includes not only the quarter from April to June, when the economy was massively affected by the state of emergency declared by the government for about two months, but also the preceding quarter from January to March.[5] Using quarterly data allows us to identify in a timely manner how the shock affected corporate cash holdings.

We obtain the following four findings. First, corporate cash holdings have increased rather than decreased since the onset of the COVID-19 crisis. Second, an increase in firms' cash flow has a positive impact on their cash holdings during normal times, and this positive cash flow sensitivity of cash was more pronounced in January–March 2020. Third, firms facing higher sales volatility held more cash in April–June 2020. Fourth, the increase in the cash flow sensitivity of cash during the crisis

---

[3] See, for example, the result of the Bank of Japan's September 2020 Tankan Survey for the heterogeneous impact on business conditions across industries.

[4] Other uncertainty measures include the economic surprise index and the market volatility index. For developments in each of these indices, see Shinohara et al. (2020).

[5] The declaration asked people to refrain from going out, but it was not legally enforceable. In addition, the Japanese government required elementary school to close temporarily on February 27. See, e.g., Watanabe and Yabu (2020) for more details.

was larger for financially constrained firms.

Overall, our findings suggest that the precautionary motive is the primary cause for the increase in Japanese publicly traded firms' cash holdings during the COVID-19 crisis, and the increase in precautionary cash holdings is more pronounced for firms that are likely to be financially constrained. Our study differs from extant research on corporate cash holdings during times of crisis in that it focuses on the COVID-19 crisis, which differs substantially in nature from the financial crises involving an increase in firms' external financing costs that previous studies focus on. In contrast, the current crisis caused direct damage to the real economy rather than via the financial sector. Our findings show that despite the absence of a rise in external financing costs, a significant increase in corporate cash holdings can be observed, which is a novel finding in the literature.

The remainder of the study proceeds as follows. Section 2 provides a more detailed overview of the related literature and posits our empirical hypotheses. Section 3 then explains the data we use and our empirical approach. Next, Section 4 presents the results, while Section 5 offers concluding remarks.

2. Related Literature and Hypotheses

Previous research has highlighted four motives for firms to hold cash: the transaction motive, the tax motive, the agency motive, and the precautionary motive. Among these, the motive that has received the most research attention is the precautionary motive. Several studies provide theoretical models that explain the role of the precautionary motive for corporate cash holdings and present supportive empirical evidence. Studies on the precautionary motive for cash holdings can be divided into two broad strands, which differ in terms of the variables they employ to examine to what extent cash holding is precautionary.

The first strand of the literature focuses on the level of firms' cash flow. For instance, Almeida

Covid Economics 68, 10 February 2021: 172-204

et al. (2004) construct a model for the precautionary demand for cash and posit that financially constrained firms save cash out of their cash flow, while the cash savings of unconstrained firms should not be systematically related to cash flows. Almeida et al. (2004) regard firms' cash flow sensitivity of cash, that is, the propensity to save cash flow for cash, as an indicator of the extent to which firms are financially constrained and hence save for precautionary reasons. Estimating cash ratio equations, they find that the correlation between cash and cash flow is indeed positive. Another study focusing on firms' cash flow but employing a different theoretical setup from that of Almeida et al. (2004) is Acharya et al. (2007). Specifically, in their setup, firms have only a limited capacity to hedge future investment opportunities against income shortfalls. In their model, firms expecting a large number of investment opportunities are more likely to save cash from their cash flow than to reduce their current debt. In their empirical analysis based on this model, the coefficients on firms' cash flow are significantly positive, which is consistent with their theoretical prediction.

A number of studies have followed up on these studies employing the methodology introduced by Almeida et al. (2004) and Acharya et al. (2007). A notable example is the study by Sun and Wang (2015), who examine precautionary corporate savings during the 2008 global financial crisis to find that the cash flow sensitivity of cash was significantly larger during the crisis than normal times.

Based on these studies, we posit our empirical hypothesis for corporate cash holdings during the COVID-19 crisis. Specifically, we argue that firms' cash flow is an important determinant of their cash holdings in normal times, since firms require a precautionary cash buffer for their day-to-day activities. In addition, the COVID-19 pandemic likely increased firms' desire for precautionary cash holdings. Our first empirical hypothesis, therefore, is as follows:

Hypothesis 1: The cash flow sensitivity of cash is positive in normal times and became more pronounced during the COVID-19 crisis.

Covid Economics 68, 10 February 2021: 172-204

Covid Economics 68, 10 February 2021: 172-204

The second strand of the literature focuses on the volatility of firms' cash flow as a determinant of precautionary cash holdings. In the literature, the volatility of a firm's cash flow is regarded as reflecting the degree of uncertainty regarding its future income and affects the amount of cash it holds. For instance, Opler et al. (1999) find that firms that face greater cash flow uncertainty hold a larger amount of cash than those with less uncertainty. Similarly, Han and Qiu (2007) theoretically show that financially constrained firms with higher cash flow volatility tend to hold a larger amount of cash for precautionary purposes. Bates et al. (2009) report that the average amount of cash held by firms in the US increased during the period 1980–2006 and conclude that the precautionary motive plays an important role in explaining the increase in the cash ratio. Finally, Riddick and Whited (2009) find that income uncertainty affects cash holdings more than do external finance constraints.

Among the follow-up studies employing the methodology introduced in this strand of the literature, several studies investigate the impact of financial crises on precautionary corporate cash holdings. For instance, Song and Lee (2012) investigate the long-term effect of the Asian financial crisis on corporate cash holdings in eight East Asian countries to show an increased sensitivity to cash flow volatility. They observe that this is one of the main factors explaining the higher level of firms' cash holdings after the crisis. Meanwhile, examining the 2008 European financial crisis, Lozano and Yaman (2020) find that for financially constrained firms the cash sensitivity to cash flow volatility was higher in the three years after the onset of the crisis than before the crisis.

Based on the above literature, we now posit our empirical hypothesis about the impact of firms' cash flow volatility on their cash holdings during the COVID-19 crisis. Specifically, we assume that firms' cash flow volatility is an important determinant of their cash holdings and that the crisis caused by the pandemic further increased the relevance of this determinant. Therefore, our second hypothesis regarding corporate cash holdings is as follows:

Hypothesis 2: There is a positive link between firms' cash flow volatility and their cash holdings in normal times, and this link became more pronounced during the COVID-19 crisis.

In the following section, we first describe the data and empirical strategy employed in our analysis and then examine the overall developments in firms' cash holdings before and since the outbreak of the pandemic, before empirically testing our hypotheses in Section 4.

3. Data and Empirical Strategy

3.1 Data and sample selection

For the analysis, we employ firm-level data from Nikkei NEEDS Financial QUEST provided by Nikkei Incorporated. Our sample comprises publicly-traded non-financial firms in Japan during the period March 2019–June 2020 and is limited to firms whose fiscal year ends in March. We drop firms that newly listed in 2020, leaving us with a final sample of 10,638 firm-quarter observations for 1,773 firms. All the variables used in our analysis are winsorized at the 1st and 99th percentiles to avoid problems caused by extreme outliers.

3.2 Empirical approach

To examine the two empirical hypotheses on the impact of cash flows and their volatility on cash holdings, we employ the following the conventional specification in the empirical literature on the determinants of cash holdings:

$$Y_{i,t} = \alpha + \beta_1 EBITDA_{i,t} + \beta_2 EBITDA_{i,t} * Crisis_t + \beta_3 Sales\ \ Volatility_{i,t} + \beta_4 Sales\ \ Volatility_{i,t} * Crisis_t + \boldsymbol{\beta X} + \delta_i + \varepsilon_{i,t} \quad (1)$$

Covid Economics 68, 10 February 2021: 172-204

CEPR PRESS

Covid Economics 68, 10 February 2021: 172-204

As dependent variable $Y_{i,t}$ we employ four different variables. The first is $Cash$, which is defined as the sum of cash and deposits outstanding divided by book assets. The second, $Liquidity$, is calculated as the sum of cash, deposits, and marketable securities outstanding divided by book assets. The third, $\Delta Cash$, is defined as the quarter-on-quarter change in $Cash$. Finally, the fourth, $\Delta Liquidity$, is the quarter-on-quarter change in $Liquidity$.

Turning to the explanatory variables, $EBITDA$, a proxy of cash flow, is the ratio of earnings before interest, taxes, and depreciation and amortization to book assets. Based on Hypothesis 1, we expect the coefficient on $EBITDA$ to be positive. $Sales\ Volatility$ is the standard deviation of a firm's sales over the five years preceding the current period standardized by the average amount of assets during the same five years. The reason that we employ the standard deviation of sales rather than that of $EBITDA$ for cash flow volatility is that a firm's cash flow consists of sales and costs, and firms likely regard volatility in sales as more exogenous and difficult to control than volatility in costs. Hypothesis 2 expects that firms with higher sales volatility are likely to be more at risk of becoming financially distressed and therefore have a greater demand for precautionary cash holdings. We therefore expect the coefficient on $Sales\ Volatility$ to be positive as well. Next, $Crisis$ represents the period of the coronavirus crisis. Specifically, we use two dummies: $Crisis\ I$ is a dummy for the first quarter (January–March) of 2020, while $Crisis\ II$ is a dummy for the second quarter, i.e., April–June 2020. The purpose of using these two different dummies is to capture how Japanese firms' cash management changed in each of two periods.

We are also interested in the interaction terms between $EBITDA$ or $Sales\ Volatility$ and $Crisis$, since both hypotheses predict that the coefficients on $EBITDA$ and $Sales\ Volatility$ should be larger during the COVID-19 crisis. We therefore expect the coefficients on these interaction terms, i.e., $EBITDA \times Crisis$ and $Sales\ Volatility \times Crisis$, to be positive.

We add various other explanatory variables as controls. $Size$ is the natural logarithm of a firm's book assets. The larger firms are, the more easily they can access external finance, and the smaller their demand for cash will be due to the smaller information asymmetry between the firm and lenders. We therefore expect the coefficient on $Size$ to be negative. Next, Q is the ratio of a firm's market to book value of assets. Firms with more growth opportunities prefer cash to external finance because of the greater extent of information asymmetry for high growth firms. Consequently, we expect a positive coefficient on Q. Further, $Capex$ is the ratio of the sum of the quarterly change in tangible assets, depreciation, and amortization to the book value of assets. We expect the coefficient on $Capex$ to be negative because capital investment increases a firm's stock of collateralizable assets and enhances its debt capacity, which leads to a smaller demand for cash. Meanwhile, NWC is the ratio of net working capital to book assets, and we predict a negative coefficient because firms use working capital as an alternative source of cash. Finally, $Debt$ is the ratio of the sum of short- and long-term debt to book assets. Different theories yield different predictions, so the sign of the coefficient on $Debt$ could be positive or negative. On the one hand, theories suggesting that debt and cash act as substitutes in terms of firms' funding sources suggest that the sign should be negative (Opler et al. 1999, Kim et al. 1998, Bates et al. 2009). On the other hand, if firms are limited in their hedging capacity and debt and cash are imperfect substitutes, as suggested by the theoretical models in Acharya et al. (2007) and Guney et al. (2007), the sign will be positive.

Covid Economics 68, 10 February 2021: 172-204

Table 1: Descriptive statistics for the variables employed in the estimations

This table reports summary statistics for the sample. Definitions of variables are provided in Section 3.2. The unit for Size is million yen.

Variables	N	Mean	Sd	Min	Median	Max
Cash	10,637	0.200	0.147	0.013	0.164	0.749
Liquidity	10,637	0.208	0.152	0.013	0.171	0.767
ΔCash	10,519	0.004	0.033	-0.097	0.002	0.126
ΔLiquidity	10,519	0.004	0.034	-0.100	0.002	0.126
EBITDA	9,771	0.019	0.022	-0.076	0.018	0.088
Sales Volatility	10,438	0.045	0.040	0.004	0.032	0.215
Size	10,637	444,339	2,079,013	284	58,046	55,900,000
Q	10,637	0.840	1.154	0.072	0.460	7.914
Capex	9,756	0.009	0.013	-0.029	0.007	0.064
NWC	10,056	0.138	0.145	-0.205	0.133	0.598
Debt	10,637	0.146	0.151	0.000	0.102	0.630

3.3 Summary Statistics

In this subsection, we describe the characteristics of the variables that we employ in our analysis. We start with descriptive statistics, which are provided in Table 1.

The mean of $Cash$ is 0.200 and that of $Liquidity$ is 0.208, indicating that in Japan, firms' cash holdings amount to about 20% of their total assets. This is above the average cash holding ratio of 16.8% reported by De Vito and Gomez (2020) for 26 mostly developed countries in 2018. The averages of $\Delta Cash$ and $\Delta Liquidity$ are identical at 0.004, indicating that the cash ratio and the liquidity ratio slightly increased during the period.

Covid Economics 68, 10 February 2021: 172-204

CEPR PRESS

Table 2: Means of variables for different subperiods

This table reports the averages of *Cash*, *Liquidity*, *Δ Cash*, *Δ Liquidity*, *EBITDA*, and *Sales Volatility* for subperiods and measures differences between periods. Definitions of variables are provided in Section 3.2. *Pre Crisis* is from 2019Q1 to 2019Q4. *Crisis I* is 2020Q1 and *Crisis II* is 2020Q2. *** and * indicate that the difference is statistically significant at the 1% or 10% level, respectively.

Period	*Cash*	*Liquidity*	*ΔCash*	*ΔLiquidity*	*EBITDA*	*Sales Volatility*
Pre Crisis	0.195	0.204	0.000	-0.001	0.021	0.045
Crisis I	0.202	0.211	0.013	0.013	0.017	0.044
Crisis II	0.214	0.222	0.011	0.011	0.011	0.046
Difference (*Crisis I* - *Pre Crisis*)	0.007*	0.007*	0.013***	0.013***	-0.003***	-0.001
Difference (*Crisis II* - *Pre Crisis*)	0.019***	0.018***	0.012***	0.012***	-0.010***	0.001

Next, we examine how the means of these variables differ across subperiods, that is, the period before and the period after the outbreak of the COVID-19 pandemic. Table 2 shows the means for various variables in each subperiod. We also test the statistical significance of differences between these subperiods. There are four notable findings. First, the cash and liquidity ratios increased significantly after the outbreak of the pandemic. *Cash* increased from 0.195 in the pre-crisis period to 0.202 in the first quarter of 2020 and 0.214 in the second quarter of 2020. *Liquidity* increased to a similar extent. The increase in these variables between the periods is statistically significant. Second, the growth in the cash and liquidity variables accelerated during the crisis period. While both $\Delta Cash$ and $\Delta Liquidity$ were around zero before the crisis, they increased to around 0.011 to 0.013 during the crisis. Third, due to the deterioration in business conditions during the crisis period, $EBITDA$, our measure of cash flow, decreased significantly after the outbreak of the crisis. While the profit rate was 0.021 before the crisis, it fell to 0.017 in the first quarter of 2020 and 0.011 in the second quarter of 2020. Fourth, there was little change in *Sales Volatility* after the start of the crisis. This is mostly due to the way we construct the variable, since we calculate the standard deviation of a firm's sales over the preceding five years.

Further, we examine detailed information on the distribution of differences in variables

Covid Economics 68, 10 February 2021: 172-204

Table 3: Percentile statistics on the differences between the pre-crisis and crisis periods

This table reports percentile statistics of differences for *Cash* , *Liquidity* , and *EBITDA* between *Crisis I* and *Pre Crisis* and between *Crisis II* and *Pre Crisis* . Definitions of variables are provided in Section 3.2. *Pre Crisis* refers to the period from 2019Q1 to 2019Q4, *Crisis I* to 2020Q1, and *Crisis II* to 2020Q2.

Difference between *Crisis I* and *Pre Crisis*								
	p1	p5	p10	p25	p50	p75	p90	p99
Cash	-0.114	-0.049	-0.032	-0.009	0.006	0.023	0.046	0.143
Liquidity	-0.112	-0.051	-0.033	-0.010	0.006	0.024	0.045	0.149
EBITDA	-0.074	-0.040	-0.024	-0.011	-0.002	0.005	0.017	0.063
Difference between *Crisis II* and *Pre Crisis*								
	p1	p5	p10	p25	p50	p75	p90	p99
Cash	-0.122	-0.049	-0.028	-0.004	0.014	0.040	0.075	0.176
Liquidity	-0.131	-0.053	-0.030	-0.004	0.013	0.040	0.074	0.176
EBITDA	-0.094	-0.048	-0.032	-0.017	-0.007	0.000	0.009	0.040

between periods. Specifically, for $Cash$, $Liquidity$, and $EBITDA$, we produce percentile statistics for the differences between $Crisis\ I$ or $Crisis\ II$ on the one hand and $Pre\ Crisis$ on the other. These statistics allow us to observe the share of firms that experienced a drop in their profitability and the share of firms that saw an increase or decrease in their cash balance.

Table 3 shows the results. The distributions of $Cash$ and $Liquidity$ shift toward the right, indicating that these variables increased in the crisis period, while that of $EBITDA$ moves toward the left, meaning that firms' profitability substantially dropped during the crisis. The increase in the cash holding ratio and the decline in profitability became more pronounced as the crisis deepened from the first quarter to the second quarter of 2020. It should be noted that in the second quarter of 2020, the cash holding ratio of the majority of firms increased even though three-quarters of the firms experienced a decline in profitability.

To summarize, there was a substantial increase in corporate cash holdings during the first half of 2020, when the economy was affected by the COVID-19 shock and firms' profitability substantially dropped. Moreover, the increase in the average cash holding ratio was due not to an increase in the

Covid Economics 68, 10 February 2021: 172-204

cash ratio of a small number of large firms but reflects an increase in the cash ratio for the majority of firms.

4. Results

In the summary statistics in the previous section, we observed an increase in corporate cash holdings and a decrease in cash flows in 2020 when the COVID-19 crisis unfolded. We also found that there was little change in sales volatility during the observation period. However, the descriptive statistics tell us little about the link between cash holdings and cash flows and, moreover, do not control for other factors. Therefore, in this section, we first present our estimation results controlling for other factors that are important determinants of cash holdings as well as firm fixed effects. Next, given that financial constraints have been highlighted as a key reason for precautionary cash holdings, we conduct various subsample analyses to examine how financial constraints affect corporate cash holdings.

4.1 Baseline Results

Table 4 presents our baseline results on the determinants of cash holdings and changes therein. Columns (1) and (2) show the results when we employ the level of cash and liquid asset holdings as the dependent variable. There are several notable findings. First, the coefficients on $EBITDA$ are positive and significant. This result indicates that the cash flow sensitivity of cash is positive in normal times. Turning to the interaction terms between cash flow and the crisis dummies, the coefficient on the interaction term between $EBITDA$ and $Crisis\ I$ is positive and significant in column (2), while it is positive but insignificant in column (1). On the other hand, the coefficients on $EBITDA *$ $Crisis\ II$ are insignificant in both columns. These results indicate that the cash flow sensitivity of cash was more pronounced at the onset of the crisis in the first quarter of 2020, but this was not

Covid Economics 68, 10 February 2021: 172-204

Table 4: Baseline results

This table shows the estimation results for cash and liquidity holdings. The dependent variables are *Cash*, *Liquidity*, *Δ Cash*, and *Δ Liquidity*. All estimations include a constant term and firm fixed effects. Definitions of variables are provided in Section 3.2. The estimations employ standard errors that are clustered by firms and reported in parentheses. *** and ** denote statistical significance at the 1% and 5% level, respectively.

	(1)	(2)	(3)	(4)
Dep. Var.	*Cash*	*Liquidity*	*ΔCash*	*ΔLiquidity*
EBITDA	0.116***	0.096**	0.117***	0.120***
	(0.044)	(0.043)	(0.040)	(0.039)
*EBITDA*Crisis I*	0.109	0.134**	0.106**	0.114**
	(0.070)	(0.062)	(0.053)	(0.051)
*EBITDA*Crisis II*	0.094	0.127	0.006	-0.002
	(0.086)	(0.081)	(0.065)	(0.066)
Sales Volatility	0.016	0.009	-0.059	-0.069
	(0.144)	(0.145)	(0.100)	(0.102)
*Sales Volatility*Crisis I*	-0.002	0.006	0.029	0.046
	(0.029)	(0.028)	(0.028)	(0.029)
*Sales Volatility*Crisis II*	0.071**	0.085***	0.070**	0.079**
	(0.032)	(0.032)	(0.034)	(0.034)
Crisis I	0.005***	0.004***	0.008***	0.008***
	(0.002)	(0.002)	(0.001)	(0.001)
Crisis II	0.010***	0.008***	0.002	0.002
	(0.002)	(0.002)	(0.002)	(0.002)
Size	-0.026	-0.033*	0.001	0.002
	(0.019)	(0.019)	(0.014)	(0.014)
Q	0.007	0.007	-0.002	-0.001
	(0.005)	(0.004)	(0.002)	(0.002)
Capex	-0.340***	-0.349***	-0.558***	-0.593***
	(0.036)	(0.036)	(0.039)	(0.039)
NWC	-0.666***	-0.690***	-0.543***	-0.561***
	(0.030)	(0.030)	(0.023)	(0.023)
Debt	0.225***	0.236***	0.289***	0.298***
	(0.047)	(0.047)	(0.030)	(0.031)
	Firm FE	Firm FE	Firm FE	Firm FE
Observations	9,723	9,723	9,723	9,723
R^2	0.379	0.397	0.269	0.286

Covid Economics 68, 10 February 2021: 172-204

necessarily the case later in the crisis (in the second quarter of 2020), when firms' cash flow declined substantially. A possible explanation for the insignificant coefficients on $EBITDA * Crisis\ II$ is that $EBITDA$ declined substantially in the $Crisis\ II$ period, so that firms may not have been able to afford to save cash out of their cash flow.

Second, we find that the coefficients on $Sales\ Volatility$ in columns (1) and (2) are insignificant. This indicates that in normal times firms do not hoard cash in response to higher sales volatility. In contrast, the coefficients on the interaction terms between sales volatility and the crisis dummies turn significantly positive in the second quarter of 2020, i.e., a few months into the crisis, while this is not the case for the first quarter of 2020, the onset of the crisis. These results suggest that firms that faced higher sales volatility started to prepare for the liquidity shortage once they realized the substantial impact of the crisis and began to expect that it would last for a long time.

Third, the coefficients on the crisis dummies indicate that firms began to hoard more cash as the crisis deepened. The result in column (1) indicates that in the first quarter of 2020 firms' cash ratio was 0.5 percentage points higher than in the same quarter of 2019. The year-on-year increase was even larger in the second quarter, reaching 1.0 percentage point. In column (2), we find a similar pattern for the liquid asset ratio.

Fourth, there are several other control variables whose coefficients are statistically significant. $Capex$ has negative coefficients, presumably because capital investment increases firms' collateralizable assets and debt capacity, thus reducing the need for them to hold cash. NWC also has negative coefficients, which suggests that working capital substitutes for cash. The coefficients on $Debt$ are positive, indicating that cash and debt are imperfect substitutes, as predicted by Acharya et al. (2007).

Next, we turn to the results in columns (3) and (4), where we employ the changes in cash and liquid asset holdings as the dependent variable. While the results are by and large similar to those

Covid Economics 68, 10 February 2021: 172-204

in columns (1) and (2), there are a few things to note. First, the coefficients on $EBITDA$ and its interaction terms with the crisis period dummies have generally the same sign as in columns (1) and (2), but some coefficients are more statistically significant. In particular, the coefficients on the interaction term between $EBITDA$ and $Crisis\ I$ are positive and significant in both columns (3) and (4). Second, the coefficients on the crisis dummies are positive and significant for the first quarter of 2020 but insignificant for the second quarter of 2020.

Overall, these results in Table 4 indicate that the increase in the cash ratio was larger at the onset of the crisis in the first quarter of 2020 and then leveled off in the second quarter of 2020, and the propensity to save cash out of cash flow was more pronounced during the crisis, which is consistent with our hypotheses.

4.2 Results for financially constrained firms

In this subsection, we conduct a set of subsample analyses and examine how financially constrained firms manage their cash holdings. As discussed in Section 2, previous theoretical studies on the precautionary motive for corporate cash holdings suggest that financially constrained firms are more likely to hoard precautionary cash than financially unconstrained firms.

Against this background, we expect the predictions of the two empirical hypotheses to apply more to financially constrained firms than to unconstrained ones. We therefore employ several variables to identify financially constrained firms, and by comparing them with unconstrained firms, we examine if financially constrained firms tended to save cash more during the crisis. The variables we use for defining constrained firms include firms' size, payout ratio, leverage, cash ratio, access to the bond market, and access to credit lines. Employing each of these variables in turn, we divide the total sample of firms into two groups (for example, small firms and large firms), define one group (small firms in this example) as financially constrained, and compare the estimation results with those

Covid Economics 68, 10 February 2021: 172-204

for unconstrainted firms (large firms in this example). Note that in the analysis that follows, we limit the dependent variables to $Cash$ and $Liquidity$, that is, variables that represent levels rather than changes. The reason is that the baseline estimation results for the coefficients on $EBITDA$, $Sales\ Volatility$, and their interaction terms with the $Crisis$ dummies in Section 4.1 were qualitatively similar regardless of whether we used the dependent variables in levels or we used changes.

4.2.1 Small versus large firms

We start by using firms' asset size to identify financially constrained firms, based on Almeida et al. (2004) and Acharya et al.'s (2007) argument that the degree of firms' external financial frictions is related to their size. Specifically, we divide the sample into two groups based on the average amount of book assets throughout the year 2019 (from the first to the fourth quarter of the year). We then regard firms with assets below the median as financially constrained and those with assets above the median as unconstrained.

Covid Economics 68, 10 February 2021: 172-204

CEPR PRESS

Covid Economics 68, 10 February 2021: 172-204

Table 5: Regression results for small versus large firms

This table shows the estimation results for cash and liquidity holdings by firms' size (measured in terms of their assets). The dependent variables are *Cash* and *Liquidity*. Columns (1) and (2) show the results for small and financially constrained firms, while columns (3) and (4) show the results for large and financially unconstrained firms. All estimations include a constant term and firm fixed effects. Definitions of variables are provided in Section 3.2. The estimations employ standard errors that are clustered by firms and reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

	(1)	(2)	(3)	(4)
	(A) Small Firms (constrained)		(B) Large Firms (unconstrained)	
Dep. Var.	*Cash*	*Liquidity*	*Cash*	*Liquidity*
EBITDA	0.131**	0.106*	0.071	0.060
	(0.058)	(0.057)	(0.046)	(0.046)
*EBITDA*Crisis I*	0.123	0.161*	0.078	0.077
	(0.101)	(0.088)	(0.065)	(0.064)
*EBITDA*Crisis II*	0.080	0.123	0.120	0.141
	(0.114)	(0.104)	(0.085)	(0.089)
Sales Volatility	0.031	0.030	-0.038	-0.073
	(0.197)	(0.199)	(0.144)	(0.144)
*Sales Volatility*Crisis I*	-0.000	0.013	0.021	0.023
	(0.043)	(0.042)	(0.033)	(0.032)
*Sales Volatility*Crisis II*	0.083*	0.104**	0.082*	0.085*
	(0.047)	(0.047)	(0.048)	(0.048)
Crisis I	0.005**	0.004	0.004**	0.003
	(0.003)	(0.002)	(0.002)	(0.002)
Crisis II	0.010***	0.008***	0.008***	0.006**
	(0.003)	(0.003)	(0.003)	(0.003)
Size	-0.034	-0.047	-0.014	-0.016
	(0.034)	(0.033)	(0.014)	(0.014)
Q	0.007	0.008*	-0.001	-0.004
	(0.005)	(0.005)	(0.006)	(0.004)
Capex	-0.394***	-0.394***	-0.272***	-0.296***
	(0.056)	(0.056)	(0.035)	(0.036)
NWC	-0.672***	-0.685***	-0.671***	-0.716***
	(0.042)	(0.041)	(0.029)	(0.027)
Debt	0.156**	0.167**	0.345***	0.363***
	(0.069)	(0.069)	(0.045)	(0.045)
	Firm FE	Firm FE	Firm FE	Firm FE
Observations	5,045	5,045	4,678	4,678
R^2	0.351	0.362	0.472	0.508

Table 5 shows the results. The two columns under (A) are for small, constrained firms, while those under (B) are for large, unconstrained firms. The coefficients on $EBITDA$ in the $Cash$ and $Liquidity$ estimations are significantly positive only for small firms. Moreover, the coefficients on the interaction terms between $EBITDA$ and the $Crisis$ dummies are marginally significant in the $Liquidity$ estimation for small firms but insignificant in all the other estimations. In sum, we find a positive cash flow sensitivity of cash only for small firms. Moreover, the $Liquidity$ estimation for small firms suggests that the extent of the cash flow sensitivity of cash increased, albeit marginally, during the onset of the crisis in the first quarter of 2020.

Meanwhile, the results for $Sales\ Volatility$ show no substantial differences between small and large firms in the way their cash holdings responded to sales volatility not only during normal times but also during the crisis period. Specifically, for both small and large firms, the coefficients on $Sales\ Volatility$ and the interaction term between $Sales\ Volatility$ and *Crisis I* are insignificant, while those on the interaction term between $Sales\ Volatility$ and *Crisis II* have the same positive sign and are of a similar magnitude.

4.2.2 Low versus high payout firms

Next, based on Fazzari et al. (1988) and Almeida et al.'s (2004) argument that firms facing high external financing costs have an incentive to reserve cash instead of paying out cash flows to shareholders, we divide the sample into two groups based on firms' payout ratio (annual payouts/total assets) in March 2020. We then regard firms below the median as constrained and those above as unconstrained.

Covid Economics 68, 10 February 2021: 172-204

Covid Economics 68, 10 February 2021: 172-204

Table 6: Regression results for low versus high payout firms

This table shows the estimation results for cash and liquidity holdings by firms' payout ratio. The dependent variables are *Cash* and *Liquidity*. Columns (1) and (2) show the results for firms with a low payout ratio (i.e., financially constrained firms), while columns (3) and (4) show the results for firms with a high payout ratio (i.e., financially unconstrained firms). All estimations include a constant term and firm fixed effects. Definitions of variables are provided in Section 3.2. The estimations employ standard errors that are clustered by firms and reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

	(1)	(2)	(3)	(4)
	(A) Low Payout Firms (constrained)		(B) High Payout Firms (unconstrained)	
Dep. Var.	*Cash*	*Liquidity*	*Cash*	*Liquidity*
EBITDA	0.053	0.066	0.170***	0.132**
	(0.067)	(0.063)	(0.056)	(0.056)
*EBITDA*Crisis I*	0.224*	0.203*	0.024	0.071
	(0.114)	(0.113)	(0.087)	(0.077)
*EBITDA*Crisis II*	0.289**	0.262**	-0.032	0.029
	(0.129)	(0.125)	(0.105)	(0.098)
Sales Volatility	0.108	0.082	-0.021	-0.018
	(0.163)	(0.165)	(0.191)	(0.192)
*Sales Volatility*Crisis I*	0.021	0.032	-0.018	-0.008
	(0.042)	(0.039)	(0.040)	(0.040)
*Sales Volatility*Crisis II*	0.121***	0.128***	0.024	0.046
	(0.036)	(0.035)	(0.051)	(0.050)
Crisis I	0.004	0.004	0.006***	0.004**
	(0.003)	(0.003)	(0.002)	(0.002)
Crisis II	0.008***	0.007**	0.011***	0.007***
	(0.003)	(0.003)	(0.003)	(0.003)
Size	-0.061**	-0.069***	-0.013	-0.021
	(0.024)	(0.025)	(0.025)	(0.024)
Q	0.009	0.011	0.004	0.005
	(0.009)	(0.008)	(0.004)	(0.004)
Capex	-0.385***	-0.396***	-0.303***	-0.311***
	(0.057)	(0.058)	(0.045)	(0.045)
NWC	-0.707***	-0.735***	-0.631***	-0.651***
	(0.033)	(0.033)	(0.049)	(0.048)
Debt	0.213**	0.230**	0.252***	0.260***
	(0.089)	(0.092)	(0.048)	(0.047)
	Firm FE	Firm FE	Firm FE	Firm FE
Observations	4,236	4,236	5,487	5,487
R^2	0.465	0.486	0.322	0.335

Table 6 presents the results. They show that, on the one hand, the coefficients on $EBITDA$ are insignificant for low payout firms, while they are significant and positive for high payout firms. On the other hand, the coefficients on the interaction terms between $EBITDA$ and the $Crisis$ dummies are positive and significant for low payout firms, while they are insignificant for their high payout counterparts. The results imply that for constrained firms the cash flow sensitivity of cash was insignificant during normal times but became positive in the first quarter of 2020 and stayed positive in the second quarter of 2020. In contrast, for unconstrained firms, the cash flow sensitivity of cash was already positive in normal times and the pandemic did not significantly change this sensitivity. Thus, Hypothesis 1, which predicts a larger cash flow sensitivity of cash during the crisis period, applies more to low payout (constrained) firms than to firms with a high payout ratio (unconstrained firms).

Meanwhile, the results for $Sales\ Volatility$ show substantial differences between low and high payout firms. Specifically, while $Sales\ Volatility$ did not affect corporate cash holdings (i.e., $Cash$ and $Liquidity$) during normal times, the significant positive coefficient on the interaction term between $Sales\ Volatility$ and $Crisis\ II$ for low payout firms indicates that cash holdings did increase for constrained firms in the second quarter of 2020. This finding suggests that Hypothesis 2, which expects firms to respond more to volatility in their performance during a crisis than in normal times, holds for constrained firms but not for unconstrained ones.

4.2.3 Low versus high cash holding firms

Further, we use firms' amount of cash holdings to identify constrained firms, based on Duchin et al.'s (2010) finding that during the global financial crisis firms with low cash reserves reduced capital investment more than firms with high cash reserves. Specifically, we divide the sample into two groups based on firms' average cash ratio from the first to the fourth quarter of 2019. We regard firms with

Covid Economics 68, 10 February 2021: 172-204

Table 7: Estimation results for low versus high cash holding firms

This table shows the estimation results for cash and liquidity holdings by firms' cash holdings. The dependent variables are *Cash* and *Liquidity*. Columns (1) and (2) show the results for firms with a low cash ratio (i.e., financially constrained firms), while columns (3) and (4) show the results for firms with a high cash ratio (i.e., financially unconstrained firms). All estimations include a constant term and firm fixed effects. Definitions of variables are provided in Section 3.2. The estimations employ standard errors that are clustered by firms and reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

	(1)	(2)	(3)	(4)
	(A) Firms with Low Cash Holdings (constrained)		(B) Firms with High Cash Holdings (unconstrained)	
Dep. Var.	*Cash*	*Liquidity*	*Cash*	*Liquidity*
EBITDA	0.036	0.044	0.130**	0.099*
	(0.035)	(0.034)	(0.057)	(0.056)
*EBITDA*Crisis I*	0.224***	0.232***	0.086	0.118
	(0.051)	(0.051)	(0.099)	(0.087)
*EBITDA*Crisis II*	0.067	0.109	0.102	0.138
	(0.083)	(0.085)	(0.111)	(0.102)
Sales Volatility	-0.092	-0.080	0.028	0.016
	(0.114)	(0.115)	(0.181)	(0.182)
*Sales Volatility*Crisis I*	0.045	0.030	-0.014	0.005
	(0.031)	(0.031)	(0.039)	(0.038)
*Sales Volatility*Crisis II*	0.120***	0.140***	0.076*	0.094**
	(0.042)	(0.041)	(0.044)	(0.044)
Crisis I	0.003*	0.002	0.006**	0.004*
	(0.001)	(0.001)	(0.002)	(0.002)
Crisis II	0.009***	0.007***	0.009***	0.006**
	(0.002)	(0.002)	(0.003)	(0.003)
Size	0.001	-0.002	-0.040	-0.051
	(0.015)	(0.014)	(0.032)	(0.031)
Q	0.006	0.006	0.005	0.006
	(0.004)	(0.004)	(0.005)	(0.005)
Capex	-0.239***	-0.270***	-0.417***	-0.412***
	(0.031)	(0.033)	(0.056)	(0.056)
NWC	-0.528***	-0.573***	-0.719***	-0.737***
	(0.030)	(0.029)	(0.036)	(0.036)
Debt	0.325***	0.345***	0.164**	0.176***
	(0.031)	(0.030)	(0.068)	(0.068)
	Firm FE	Firm FE	Firm FE	Firm FE
Observations	3,974	3,974	5,749	5,749
R^2	0.484	0.513	0.378	0.392

an average cash ratio below the median as constrained and those above as unconstrained.

Table 7 reports the results. Similar to the results in Table 6, the coefficients on $EBITDA$ are insignificant but those on the interaction term between $EBITDA$ and $Crisis\ I$ are positive and

Covid Economics 68, 10 February 2021: 172-204

significant for constrained firms. For unconstrained firms, the coefficients on $EBITDA$ are significant but those on the interaction terms are insignificant. Hence, Hypothesis 1 applies more to firms with a low cash ratio than those with a high cash ratio.

The results for $Sales\ Volatility$ show no substantial differences between low and high cash ratio firms in terms of how their cash holdings respond to sales volatility. That is, for firms in both subsamples, the coefficients on the interaction term between $Sales\ Volatility$ and $Crisis\ II$ are positive and significant. Note, however, that the size of the coefficients is somewhat larger for constrained than for unconstrained firms, suggesting that our Hypothesis 2 applies more to low than high cash ratio firms.

4.2.4 High versus low leverage firms

Next, we focus on firms' leverage. Highly levered firms are often regarded as financially constrained. This is due to the debt overhang problem, which means that firms loaded with a large amount of debt are unable to find new funding sources. We therefore divide the sample into two groups based on firms' leverage, which we define as the average ratio of a firm's book value of liabilities to the total assets from the first to the fourth quarter of 2019. We regard firms that are above the median leverage as constrained and those that are below the median as unconstrained.

Covid Economics 68, 10 February 2021: 172-204

Covid Economics 68, 10 February 2021: 172-204

Table 8: Estimation results for high versus low leverage firms

This table shows the estimation results for cash and liquidity holdings by firms' leverage. The dependent variables are *Cash* and *Liquidity*. Columns (1) and (2) show the results for firms with a high leverage (i.e., financially constrained firms), while columns (3) and (4) show the results for firms with a low leverage (i.e., financially unconstrained firms). All estimations include a constant term and firm fixed effects. Definitions of variables are provided in Section 3.2. The estimations employ standard errors that are clustered by firms and reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

	(1)	(2)	(3)	(4)
	(A) High Leverage Firms (constrained)		(B) Low Leverage Firms (unconstrained)	
Dep. Var.	*Cash*	*Liquidity*	*Cash*	*Liquidity*
EBITDA	0.071	0.054	0.200***	0.182***
	(0.053)	(0.056)	(0.067)	(0.064)
*EBITDA*Crisis I*	0.223***	0.235***	-0.010	0.034
	(0.062)	(0.063)	(0.137)	(0.120)
*EBITDA*Crisis II*	0.199*	0.225**	0.004	0.058
	(0.111)	(0.113)	(0.133)	(0.116)
Sales Volatility	0.192	0.195	-0.184	-0.206
	(0.192)	(0.192)	(0.199)	(0.199)
*Sales Volatility*Crisis I*	0.048	0.048	-0.048	-0.019
	(0.033)	(0.033)	(0.053)	(0.052)
*Sales Volatility*Crisis II*	0.070	0.075*	0.087*	0.116**
	(0.045)	(0.045)	(0.049)	(0.048)
Crisis I	0.001	0.000	0.011***	0.008***
	(0.002)	(0.002)	(0.003)	(0.003)
Crisis II	0.010***	0.009***	0.012***	0.007**
	(0.002)	(0.003)	(0.004)	(0.004)
Size	-0.016	-0.015	-0.060	-0.091**
	(0.021)	(0.021)	(0.042)	(0.037)
Q	0.006	0.005	0.006	0.007
	(0.004)	(0.004)	(0.007)	(0.006)
Capex	-0.286***	-0.292***	-0.437***	-0.454***
	(0.042)	(0.042)	(0.064)	(0.064)
NWC	-0.569***	-0.583***	-0.804***	-0.841***
	(0.045)	(0.046)	(0.035)	(0.034)
Debt	0.229***	0.226***	0.173	0.212*
	(0.047)	(0.048)	(0.108)	(0.108)
	Firm FE	Firm FE	Firm FE	Firm FE
Observations	5,170	5,170	4,553	4,553
R^2	0.386	0.388	0.402	0.438

Covid Economics 68, 10 February 2021: 172-204

Table 8 shows the results. Similar to the results in Tables 6 and 7, the coefficients on $EBITDA$ are insignificant for constrained firms, while they are significant for unconstrained firms. The coefficients on the interaction terms between $EBITDA$ and the $Crisis$ dummies are positive and significant for high leverage firms, while none of the coefficients on the interaction terms are significant for low leverage firms. Therefore, we can say that our Hypothesis 1 applies more to high leverage than low leverage firms.

However, the results for the interaction terms between $Sales\ Volatility$ and the $Crisis$ dummies are not in line with our hypothesis that the cash holdings of financially constrained firms were likely to have increased more during the crisis than those of unconstrained firms. Specifically, we find that the coefficients on $Sales\ Volatility * Crisis\ II$ are no larger for constrained firms than for unconstrained firms.

4.2.5 Firms without versus firms with access to the bond market

In addition, we follow the literature focusing on bond market access in order to identify financially constrained firms. Firms that have access to the bond market are more creditworthy and face lower external financing costs (Almeida et al., 2004; Acharya et al., 2007). In contrast, firms that do not have access to the bond market are unable to tap this important source for financing and consequently need to rely on banks for funding. We therefore regard these firms as firms without bond market access and consider them to be more financially constrained than firms that have access to the bond market. To identify whether firms have access to the bond market, we employ the method introduced in Iwaki (2019). Specifically, we look at corporate bond (CB) and corporate paper (CP) issuance records as well as firms' balance sheet information on CBs and CP outstanding spanning the period from 2000 to 2019. We regard firms that issued CBs or CP or had a non-zero amount of CBs or CP outstanding on their balance sheet at least once during the period as firms with access to the bond market. On the

Covid Economics 68, 10 February 2021: 172-204

Table 9: Estimation results for firms without versus firms with access to the bond market

This table shows the estimation results for cash and liquidity holdings for firms with and without access to the bond market. The dependent variables are *Cash* and *Liquidity*. Columns (1) and (2) show the results for firms without access to the bond market (i.e., financially constrained firms), while columns (3) and (4) show the results for firms with access to the bond market (i.e., financially unconstrained firms). All estimations include a constant term and firm fixed effects. Definitions of variables are provided in Section 3.2. The estimations employ standard errors that are clustered by firms and reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

	(1)	(2)	(3)	(4)
	(A) Firms without Bond Market Access (constrained)		(B) Firms with Bond Market Access (unconstrained)	
Dep. Var.	*Cash*	*Liquidity*	*Cash*	*Liquidity*
EBITDA	0.116**	0.093**	0.184*	0.208**
	(0.046)	(0.045)	(0.099)	(0.096)
*EBITDA*Crisis I*	0.129*	0.156**	-0.183	-0.193
	(0.075)	(0.067)	(0.118)	(0.117)
*EBITDA*Crisis II*	0.110	0.145*	-0.167	-0.115
	(0.091)	(0.085)	(0.185)	(0.178)
Sales Volatility	0.004	-0.003	0.186	0.185
	(0.152)	(0.153)	(0.300)	(0.310)
*Sales Volatility*Crisis I*	-0.002	0.006	0.045	0.020
	(0.030)	(0.030)	(0.070)	(0.071)
*Sales Volatility*Crisis II*	0.075**	0.090***	0.240***	0.211**
	(0.034)	(0.034)	(0.091)	(0.095)
Crisis I	0.005***	0.004**	0.008***	0.006**
	(0.002)	(0.002)	(0.003)	(0.003)
Crisis II	0.009***	0.007***	0.010**	0.008**
	(0.002)	(0.002)	(0.004)	(0.004)
Size	-0.033	-0.041**	0.101***	0.101***
	(0.020)	(0.020)	(0.025)	(0.025)
Q	0.007	0.008*	-0.011	-0.018
	(0.005)	(0.005)	(0.013)	(0.013)
Capex	-0.350***	-0.356***	-0.317***	-0.348***
	(0.040)	(0.040)	(0.050)	(0.051)
NWC	-0.673***	-0.696***	-0.563***	-0.611***
	(0.032)	(0.031)	(0.049)	(0.045)
Debt	0.217***	0.232***	0.325***	0.303***
	(0.051)	(0.051)	(0.044)	(0.039)
	Firm FE	Firm FE	Firm FE	Firm FE
Observations	8,405	8,405	1,318	1,318
R^2	0.380	0.398	0.464	0.488

other hand, we define firms that have not issued any bonds or CP as firms without access to the bond market.

Table 9 reports the estimation results for the two groups of firms. The coefficients on

EBITDA are positive and significant in all estimations. In contrast, the coefficients on the interaction term between *EBITDA* and *Crisis I* are positive for firms without access to the bond market but not for those with. It can therefore be said that our Hypothesis 1 applies more to firms without access to the bond market than to firms with bond market access.

However, the results for the interaction terms between *Sales Volatility* and the *Crisis* dummies are not in line with our hypothesis that the cash holdings of financially constrained firms were likely to have increased more during the crisis than those of unconstrained firms. More specifically, the size of the coefficients on *Sales Volatility***Crisis II* for firms without access to the bond market is substantially smaller than that for firms with bond market access.

4.2.6 Firms that have access to credit lines and firms that do not

Finally, we focus on whether firms have access to credit lines and use this information to identify constrained firms. Studies examining firm financing during the 2007–2008 global financial crisis found that large firms in the United States massively drew down credit lines (e.g., Ivashina and Scharfstein, 2010). This underlines the importance of credit lines as a means for firms to access emergency funding during crisis times. In Japan, the use of credit lines has been on the rise since legal reforms at the end of the 1990s. Therefore, to take credit lines into account, we use information on credit lines as of the end of March 2020 to identify firms with credit lines, which we regard as unconstrained firms, and those without, which we regard as constrained firms.[6]

[6] Information on credit lines is taken from Nikkei NEEDS Financial QUEST. However, since some of the necessary information is missing in the database, we supplement it with information from another database called "eol" provided by PRONEXUS Incorporated. Meanwhile, the reason we focus on information at year-end regarding firms' credit lines is that firms usually disclose credit line information on an annual rather than a quarterly basis.

Covid Economics 68, 10 February 2021: 172-204

Table 10: Estimation results for firms that have access to credit lines and firms that do not

This table shows the estimation results of cash and liquidity holdings for firms without and with access to credit lines. The dependent variables are *Cash* and *Liquidity*. Columns (1) and (2) show the results for firms with no access to credit lines (i.e., financially constrained firms), while columns (3) and (4) show the results for firms with access to credit lines (i.e., financially unconstrained firms). All estimations include a constant term and firm fixed effects. Definitions of variables are provided in Section 3.2. The estimations employ standard errors that are clustered by firms and reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% level, respectively.

	(1)	(2)	(3)	(4)
	(A) Firms without Credit Lines (constrained)		(B) Firms with Credit Lines (unconstrained)	
Dep. Var.	*Cash*	*Liquidity*	*Cash*	*Liquidity*
EBITDA	0.104**	0.091**	0.202**	0.119
	(0.048)	(0.046)	(0.101)	(0.127)
*EBITDA*Crisis I*	0.099	0.128*	0.083	0.104
	(0.080)	(0.071)	(0.124)	(0.133)
*EBITDA*Crisis II*	0.101	0.132	-0.097	-0.034
	(0.095)	(0.088)	(0.189)	(0.195)
Sales Volatility	0.086	0.072	-0.305	-0.284
	(0.174)	(0.175)	(0.230)	(0.232)
*Sales Volatility*Crisis I*	-0.014	-0.013	0.046	0.054
	(0.034)	(0.033)	(0.048)	(0.050)
*Sales Volatility*Crisis II*	0.056	0.075**	0.131	0.135
	(0.034)	(0.034)	(0.118)	(0.116)
Crisis I	0.006***	0.005***	0.005*	0.003
	(0.002)	(0.002)	(0.003)	(0.003)
Crisis II	0.010***	0.008***	0.015**	0.012
	(0.002)	(0.002)	(0.007)	(0.007)
Size	-0.017	-0.026	-0.047	-0.043
	(0.020)	(0.019)	(0.033)	(0.034)
Q	0.008	0.008*	0.002	0.001
	(0.005)	(0.005)	(0.008)	(0.008)
Capex	-0.331***	-0.339***	-0.302***	-0.298***
	(0.040)	(0.041)	(0.083)	(0.086)
NWC	-0.687***	-0.709***	-0.601***	-0.635***
	(0.031)	(0.030)	(0.040)	(0.042)
Debt	0.225***	0.238***	0.213**	0.209**
	(0.054)	(0.054)	(0.087)	(0.089)
	Firm FE	Firm FE	Firm FE	Firm FE
Observations	8,338	8,338	1,385	1,385
R^2	0.381	0.399	0.452	0.456

Covid Economics 68, 10 February 2021: 172-204

Table 10 shows the results for both groups of firms. For firms with no credit lines, the coefficients on *EBITDA* are positive and significant in the estimations, while for firms that have credit lines that they can draw down if necessary the coefficient is positive and significant in the estimation for *Cash* but insignificant in the estimation for *Liquidity.* Further, the coefficient on the interaction term between *EBITDA* and *Crisis II* is positive and significant for firms without credit lines in the *Liquidity* estimation but insignificant in all the other estimations. These results provide another piece of evidence that Hypothesis 1 applies more to firms that are more financially constrained than to less constrained firms. In contrast, the coefficients on the interaction terms between *Sales Volatility* and the *Crisis* dummies are insignificant in all estimations, meaning that these estimations do not provide any insights with regard to Hypothesis 2.

In the above analyses, we used various definitions for financially constrained firms and examined if our Hypotheses 1 and 2 apply more to such constrained firms than unconstrained firms. Hypothesis 1, which predicts a substantial increase in the cash flow sensitivity of cash during the crisis, holds more for constrained firms of all types than for their unconstrained counterparts. For Hypothesis 2, which predicts a substantial increase in the response of cash holdings to volatility in cash flows for constrained firms, our findings are less clear-cut. The hypothesis holds for some definitions of constrained firms (namely, firms with a low payout ratio and firms with a low cash ratio) but not for others.

5. Concluding Remarks

COVID-19 started spreading around the world at the beginning of 2020 and has caused severe damage to the Japanese economy. This study focused on the firm sector in Japan and examined how firms' cash holdings have been affected during the crisis. We found the following: (1) corporate cash holdings have increased rather than decreased since the onset of the crisis; (2) an increase in firms' cash flow

Covid Economics 68, 10 February 2021: 172-204

had a positive impact on their cash holdings during normal times, and the positive cash flow sensitivity of cash was more pronounced during the first three months of the crisis; (3) firms facing higher sales volatility than other firms held more cash as the crisis unfolded; and (4) the increase in the cash flow sensitivity of cash during the crisis was larger for financially constrained firms.

Overall, the COVID-19 crisis has had a substantial impact on corporate cash management and the results are consistent with the precautionary motive theory for cash holdings. However, the present study only represents a first attempt at examining these issues using the COVID-19 crisis as an experiment. The observation period in our analysis covers only the first six months of 2020 and our sample consists only of listed firms. A task for the future, therefore, is to extend the observation period as more data become available and to expand the analysis to smaller, unlisted firms.

References

Acharya, V. V., Almeida, H., and Campello, M. (2007). “Is cash negative debt? A hedging perspective on corporate financial policies,” Journal of Financial Intermediation, 16(4): 515–554.

Acharya, V. V., and Steffen, S. (2020). “The risk of being a fallen angel and the corporate dash for cash in the midst of COVID,” The Review of Corporate Finance Studies , 9(3): 430–471

Almeida, H., Campello, M., and Weisbach, M. S. (2004). “The cash flow sensitivity of cash,” The Journal of Finance, 59(4): 1777 –1804.

Bates, T. W., Kahle, K. M., and Stulz, R. M. (2009). “Why do US firms hold so much more cash than they used to?” The Journal of Finance, 64(5): 1985–2021.

Bank of Japan (2020) “Financial Stability Report,” (October).

De Vito, A., and Gomez, J. P. (2020). “Estimating the COVID-19 cash crunch: Global evidence and policy,” Journal of Accounting and Public Policy, 39(2): 106741.

Duchin, R., Ozbas, O., and Sensoy, B. A. (2010). “Costly external finance, corporate investment, and the subprime mortgage credit crisis," Journal of Financial Economics, 97(3): 418–435.

Fazzari, S. M., Hubbard, R. G., and Petersen, B. C. (1988). "Financing constraints and corporate investment," Brookings Papers on Economic Activity, 1: 141–195

Francis, B. B., García, R. E., and Sharma, Z. (2020). "Impact of COVID-19 on corporate debt structure: Cross country evidence," available at SSRN 3666684.

Guney, Y., Ozkan, A., and Ozkan, N. (2007). "International evidence on the non-linear impact of leverage on corporate cash holdings," Journal of Multinational Financial Management, 17(1): 45–60.

Han, S., and Qiu, J. (2007). "Corporate precautionary cash holdings," Journal of Corporate Finance, 13(1): 43–57.

Iwaki, H. (2019). "The effect of debt market imperfection on capital structure and investment: Evidence from the 2008 global financial crisis in Japan," The Quarterly Review of Economics and Finance, 74: 251–266.

Ivashina, V., and Scharfstein, D. (2010). "Bank lending during the financial crisis of 2008," Journal of Financial Economics, 97(3): 319–338.

Kim, C. S., Mauer, D. C., and Sherman, A. E. (1998). "The determinants of corporate liquidity: Theory and evidence," Journal of Financial and Quantitative Analysis, 33(3): 335–359.

Li, L., Strahan, P. E., and Zhang, S. (2020). "Banks as lenders of first resort: Evidence from the COVID-19 crisis," The Review of Corporate Finance Studies, 9(3): 472–500.

Lozano, M. B., and Yaman, S. (2020). "The European financial crisis and firms' cash holding policy: An analysis of the precautionary motive," Global Policy, 11(S1): 84–94.

Opler, T., Pinkowitz, L., Stulz, R., and Williamson, R. (1999). "The determinants and implications of corporate cash holdings," Journal of Financial Economics, 52(1): 3–46.

Riddick, L. A., and Whited, T. M. (2009). "The corporate propensity to save," The Journal of Finance,

Covid Economics 68, 10 February 2021: 172-204

64(4): 1729–1766.

Shinohara, T., Okuda, T., and Nakajima, J. (2020). "Characteristics of uncertainty indices in the macroeconomy," Bank of Japan Working Paper Series 20-E-6, Bank of Japan.

Song, K., and Lee, Y. (2012). "Long-term effects of a financial crisis: Evidence from cash holdings of East Asian firms," Journal of Financial and Quantitative Analysis, 47(3): 617–641.

Sun, Z., and Wang, Y. (2015). "Corporate precautionary savings: Evidence from the recent financial crisis," The Quarterly Review of Economics and Finance, 56: 175–186.

Watanabe, T., and Yabu, T. (2020). "Japan's voluntary lockdown," Covid Economics: Vetted and Real-Time Papers, 46: 1–31.