


**CENTRE FOR  
ECONOMIC  
POLICY  
RESEARCH**

**CEPR PRESS**



**COVID ECONOMICS**  
VETTED AND REAL-TIME PAPERS

**ISSUE 54**  
29 OCTOBER 2020

**WORK, CARE AND GENDER**

Claudia Hupkau and Barbara Petrongolo

**UNCERTAINTY, HYSTERESIS AND  
LOCKDOWNS**

Charles Sims and David Finnoff

**SOCIAL LEARNING**

Yuan Tian, Maria Esther Caballero  
and Brian K. Kovak

**URBAN FLIGHT IS SPREADING  
INFECTIONS**

Joshua Coven, Arpit Gupta and Iris Yao

**THE FUTURE OF US FERTILITY**

Joshua Wilde, Wei Chen  
and Sophie Lohmann

---

# 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](mailto:covidecon@cepr.org).

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

### **The Centre for Economic Policy Research (CEPR)**

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

Chair of the Board

Sir Charlie Bean

Founder and Honorary President

Richard Portes

President

Beatrice Weder di Mauro

Vice Presidents

Maristella Botticini

Ugo Panizza

Philippe Martin

Hélène Rey

Chief Executive Officer

Tessa Ogden

---

# Editorial Board

**Beatrice Weder di Mauro**, CEPR

**Charles Wyplosz**, Graduate Institute Geneva and CEPR

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

**Guido Alfani**, Bocconi University and CEPR

**Franklin Allen**, Imperial College Business School and CEPR

**Michele Belot**, European University Institute and CEPR

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

**Nick Bloom**, Stanford University and CEPR

**Tito Boeri**, Bocconi University and CEPR

**Alison Booth**, University of Essex and CEPR

**Markus K Brunnermeier**, Princeton University and CEPR

**Michael C Burda**, Humboldt Universitaet zu Berlin and CEPR

**Aline Bütikofer**, Norwegian School of Economics

**Luis Cabral**, New York University and CEPR

**Paola Conconi**, ECARES, Universite Libre de Bruxelles and CEPR

**Giancarlo Corsetti**, University of Cambridge and CEPR

**Fiorella De Fiore**, Bank for International Settlements and CEPR

**Mathias Dewatripont**, ECARES, Universite Libre de Bruxelles and CEPR

**Jonathan Dingel**, University of Chicago Booth School and CEPR

**Barry Eichengreen**, University of California, Berkeley and CEPR

**Simon J Evenett**, University of St Gallen and CEPR

**Maryam Farboodi**, MIT and CEPR

**Antonio Fatás**, INSEAD Singapore and CEPR

**Francesco Giavazzi**, Bocconi University and CEPR

**Christian Gollier**, Toulouse School of Economics and CEPR

**Timothy J. Hatton**, University of Essex and CEPR

**Ethan Ilzetzki**, London School of Economics and CEPR

**Beata Javorcik**, EBRD and CEPR

**Simon Johnson**, MIT and CEPR

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

**Tom Kompas**, University of Melbourne and CEBRA

**Miklós Koren**, Central European University and CEPR

**Anton Korinek**, University of Virginia and CEPR

**Michael Kuhn**, Vienna Institute of Demography

**Maarten Lindeboom**, Vrije Universiteit Amsterdam

**Philippe Martin**, Sciences Po and CEPR

**Warwick McKibbin**, ANU College of Asia and the Pacific

**Kevin Hjortshøj O'Rourke**, NYU Abu Dhabi and CEPR

**Evi Pappa**, European University Institute and CEPR

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

**Richard Portes**, London Business School and CEPR

**Carol Propper**, Imperial College London and CEPR

**Lucrezia Reichlin**, London Business School and CEPR

**Ricardo Reis**, London School of Economics and CEPR

**Hélène Rey**, London Business School and CEPR

**Dominic Rohner**, University of Lausanne and CEPR

**Paola Sapienza**, Northwestern University and CEPR

**Moritz Schularick**, University of Bonn and CEPR

**Flavio Toxvaerd**, University of Cambridge  
**Christoph Trebesch**, Christian-Albrechts-Universitaet zu Kiel and CEPR

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

**Jan C. van Ours**, Erasmus University Rotterdam and CEPR

**Thierry Verdier**, Paris School of Economics and CEPR

---

# Ethics

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

## Submission to professional journals

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

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

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

---

# Covid Economics

## Vetted and Real-Time Papers

Issue 54, 29 October 2020

### Contents

Work, care and gender during the covid-19 crisis <i>Claudia Hupkau and Barbara Petrongolo</i>	1
Uncertainty, hysteresis and lockdowns <i>Charles Sims and David Finnoff</i>	29
Social learning along international migrant networks <i>Yuan Tian, Maria Esther Caballero and Brian K. Kovak</i>	62
Urban flight seeded the COVID-19 pandemic across the United States <i>Joshua Coven, Arpit Gupta and Iris Yao</i>	121
COVID-19 and the future of US fertility: What can we learn from Google? <i>Joshua Wilde, Wei Chen and Sophie Lohmann</i>	158

# Work, care and gender during the covid-19 crisis<sup>1</sup>

Claudia Hupkau<sup>2</sup> and Barbara Petrongolo<sup>3</sup>

Date submitted: 22 October 2020; Date accepted: 22 October 2020

*We explore impacts of the pandemic crisis and associated restrictions to economic activity on paid and unpaid work for men and women in the UK. Using data from the Covid-19 supplement of Understanding Society, we find evidence that labour market outcomes of men and women were roughly equally affected at the extensive margin, as measured by the incidence of job loss or furloughing, but if anything women suffered smaller losses at the intensive margin, experiencing slightly smaller changes in hours and earnings. Within the household, women provided on average a larger share of increased childcare needs, but in an important share of households fathers became the primary childcare providers. These distributional consequences of the pandemic may be important to understand its inequality legacy over the longer term.*

1 The authors thank Abi Adams-Prassl and Monica Costa Dias for helpful comments, and gratefully acknowledge financial support from the Nuffield Foundation (grant number WEL/34603).

2 Assistant Professor of Economics, CUNEF, and Research Associate, Centre for Economic Performance, LSE.

3 Professor of Economics, University of Oxford, and Research Associate, Centre for Economic Performance, LSE.

Copyright: Claudia Hupkau and Barbara Petrongolo

# 1 Introduction

Covid-19 is hitting most economies as hard as the deepest recessions, but given the exceptional nature of this crisis, the distribution of jobs and workers affected is quite different from previous downturns. While sectors like construction and manufacturing are typically most affected in regular recessions, including the Great Recession, the social distancing and lockdown measures implemented in response to the Covid-19 crisis have naturally hit service sectors with frequent interactions among consumers or between consumers and providers, such as retail, hotels, restaurants and travel. But even among workers whose activities are or were not directly subject to lockdowns, many are or have been unable to work as normal, as their work would not comply with social distancing (e.g. in construction, repairs and home services), and can be hardly performed from home.

Another distinctive feature of the pandemic crisis has been its impact on the volume of home production, reversing by decree a secular process of marketisation of childcare and home keeping. During lockdowns, virtually none of the typical components of home production could be outsourced to the market, and the closure of schools and nurseries meant that all education and childcare services were added to pre-existing home production needs.

The impact of the pandemic on the labour market as well as the volume of home production is likely to have consequences for the gender distribution of work. On the one hand, women tend to be over-represented in service industries that have been subject to lockdowns or social distancing measures. On the other hand, they are also over-represented in sectors that have been defined as critical to the Covid-19 response, as well as in occupations that can be performed from home. It is therefore ex-ante unclear whether one should expect women's labour market prospects to be more severely affected than men's. Another key aspect is that women on average perform the best part of home production tasks, most notably childcare, and more in general they bear almost the entirety of the earning penalty associated with childbearing (Kleven et al., 2019). Thus, increased care responsibilities while Covid-19 restrictions last could negatively impact gender inequality in earnings in the longer run.

This chapter contributes to a recent but growing economic literature investigating unequal socio-economic impacts of Covid-19 across a number of dimensions (see, among others, Adams-Prassl et al. 2020 and Blundell et al. 2020). A strand of this literature has

devoted special attention to unequal gender impacts. For the US, [Alon et al. \(2020\)](#) document larger employment losses for women than for men and explore their consequences for macroeconomic adjustment and the household division of labour. For the UK, [Andrew et al. \(2020\)](#) find that, in households with dependent children, mothers are more likely than fathers to be out of work or furloughed during the crisis and that the substantial increase in childcare for both parents has on average enlarged fathers' share of total childcare. [Sevilla and Smith \(2020\)](#) detect a larger increase in the overall childcare burden for mothers, but, as for [Andrew et al. \(2020\)](#), this is associated with a slight increase in father's share of total childcare, simply because fathers' contribution to childcare was on average much lower before the crisis. Finally, [Oreffice and Quintana-Domeque \(2020\)](#) find that poorer female employment outcomes during the crisis are also accompanied by a higher incidence of mental health issues. For other countries, [Farré et al. \(2020\)](#) and [Del Boca et al. \(2020\)](#) look into impacts on both paid and domestic work in Spain and Italy, respectively. In both countries, women take over most of the increased childcare burden, but evidence on their labour market outcomes is less clear-cut.

The majority of papers in this literature draw on evidence from ad hoc, real-time surveys carried out during the pandemic. These typically contain rich information on Covid-related aspects of work and family life, but they may not be linked to pre-pandemic outcomes. Only more recently have regular household and labour market surveys started to release waves of data that cover the pandemic period, with larger sample sizes and richer information on work and employment patterns at baseline.


Our work contributes to the literature on the Covid-19 impacts on the gender division of work in the labour market and the household, using data from the Covid-19 supplement to the Understanding Society longitudinal study. In contrast to results from independent surveys, we find evidence of roughly equal furloughing (and job loss) incidence across genders, but women on average experience slightly smaller hours' and earnings' losses, whether unconditional or controlling for a rich set of individual and job characteristics. Within the household, women on average take over the majority of increased childcare hours during the pandemic, but in a sizeable share of households fathers become the primary providers of childcare. These distributional consequences of the pandemic are important to understand its inequality legacy over the longer term.



## 2 Work patterns at baseline

We start by showing a snapshot of male and female work patterns at baseline under the lens of the Covid-19 incidence. This is done using data from the UK Quarterly Labour Force Survey (LFS) for April-June 2019, whose large sample size and detailed occupation and industry classifications allow us to precisely identify jobs subject to shut-downs and those that have been defined as critical to the Covid-19 response.<sup>1</sup>

Figure 1: The composition of jobs according to Covid-19 incidence



Notes. The bars show the incidence of critical jobs and shut-down jobs, as well as the incidence of working from home among those not in critical or shut-down jobs. For completeness, the percentage of critical jobs that can be done from home is 44 for men and 41 for women, and the percentage of shut-down jobs that can be done from home is 22 for men and 24 for women. Sample: employed men and women aged 16-64. Source: UK LFS, April-June 2019.

We classify jobs into four categories. The first group includes jobs in critical industries (mostly health care, public services and security). The second group includes jobs in shut-down industries (mostly non-essential retail, hospitality, accommodation and food services).<sup>2</sup> We categorise all remaining jobs into those that can be done from home and those that cannot, which is plausibly the relevant distinction to predict employment

<sup>1</sup>This evidence was previously shown in [Hupkau and Petrongolo \(2020\)](#).

<sup>2</sup>We classify industries as critical if they are mentioned in [Cabinet Office and Department for Education \(19 March 2020\)](#) and as shut-down if they are mentioned in [Cabinet Office and Ministry of Housing, Communities and Local Government \(9 April 2020\)](#).


and earning losses outside critical and shut-down sectors. This classification is done by matching [Dingel and Neiman \(2020\)](#)'s classification of teleworkable occupations – based on task descriptions in O\*NET – with the UK classification of occupations in the LFS.<sup>3</sup>

The distribution of employment across these four categories is shown in [Figure 1](#). More women than men are employed in critical sectors (about 46% and 39% of working women and men, respectively). Offsetting this, more women than men are employed in shut-down sectors (about 19% and 13%, respectively). For the remaining 48% of men and 35% of women, the incidence of earnings losses is closely linked to their ability to work from home (WFH). WFH is largely possible in female-dominated sectors like education, where teachers support distance learning for many children and young people. In contrast, WFH is not possible in many male-dominated sectors like construction, repairs, and large parts of manufacturing. Indeed, about 24% of women and 19% of men are in jobs that can be done from home – having excluded critical and shut-down sectors. Taking these factors into account, it is ex-ante unclear whether women's employment and earning prospects should be more or less severely affected than men's.

The other relevant aspect of the pandemic regards the gender distribution of home production, including (most notably) childcare. One important factor behind changes in childcare needs is marital or cohabitation status. Women are more likely than men to raise children as single parents. Using LFS data, we estimate that 20.3% of households with dependent children (aged 15 or below) are headed by single mothers, against 3.3% headed by single fathers. Hence, for single parent households, women are far more likely than men to be the sole providers of the sharp increase in childcare during the lockdown. Second, the distribution of home production depends on the working status of partners (if any), which is itself affected by the crisis.

<sup>3</sup>[Dingel and Neiman \(2020\)](#) use responses to O\*NET surveys on work context and activities to classify 6-digit occupations into those that can be performed from home and those that cannot (binary classification). We map the resulting 6-digit O-NET-SOC2010 classification into the 4-digit UK SOC2010 classification available in the UK LFS based on a cross-walk from 6-digit O-NET-SOC2010 to 4-digit US-SOC2010 and finally to 4-digit UK-SOC2010 occupations (369 categories). When a few 6-digit occupations feed into one 4-digit occupation, we classify the 4-digit occupation as doable from home if the majority of 6-digit occupations associated with it are classified as such. We manually re-classify as doable from home a handful of managerial and technical 4-digit occupations (e.g. elected officers and representatives, financial administrative occupations); and manually re-classify as not doable from home about 30 miscellaneous occupations (a few occupations in public transport, a few care and service occupations, and a few technician occupations associated with workplaces, e.g. lab technicians). Overall, we estimate that 43% of jobs in the UK can be done from home (based on LFS data for April-June 2019). [Dingel and Neiman \(2020\)](#) perform a similar exercise and obtain an estimate of 43.5% for the UK, based on ILO data from 2018.

Figure 2: The distribution of partner's status, by women's status



Notes. The “other” status indicates women staying at home during COVID-19 (including: in shut-down jobs, in non-critical jobs, not employed). The sample does not include households with two or more family units or same-sex couples with children (representing, respectively, 2.33% and 0.23% of households with children). Sample: Women with dependent children aged 15 or below. Source: UK LFS, April-June 2019.

Figure 2 shows the distribution of partner status for women with dependent children. Around one third of all women with dependent children work in critical jobs (as opposed to 46% of all working women). Of these, 57% have either no partner or a partner who also works in a critical job, and are likely to rely on basic childcare services guaranteed by the education system to parents in critical jobs. The remaining 43% has a partner who is staying at home – whether he is employed in a shut-down sector (6%), or cannot go to work due to social distancing (33%), or does not work at all (4%). In these households, we would expect a reversal of the home production gap, with men taking over the bulk of increased childcare and housekeeping needs. Among mothers who are not in critical jobs, and therefore stay at home during the lockdown, 21% have no partner and 26% have a partner in a critical job, and hence are likely fully in charge of home production. The other 53% has a stay-at-home partner, and home production is somehow shared between spouses.

There is plenty of pre-Covid-19 evidence on the contribution of men and women to home production from time use data. According to the 2014-15 Time Use Survey for the

UK, women do 27 hours per week of home production on average, while men do 16 hours on average. Among households with dependent children, weekly home production hours are 40 for mothers and 20 for fathers, of which 17 and 8, respectively, represent childcare. The key question is therefore whether the additional home production falls on men and women according to baseline specialization patterns, in which case women would be at the receiving end of the best part of increased home production requirements.

Below we address questions on impacts of Covid-19 on the gender division of both paid and unpaid work using the Covid-19 supplement of the Understanding Society (USoc) study. Relative to the UK LFS, the Covid-19 study has the advantage of surveying participants at the monthly frequency, linking their answers to regular USoc waves, and providing detailed information on domestic work, including childcare and home schooling. The disadvantage of the Covid-19 study, however, is that it does not contain fine-grained information on occupation or industry at baseline, hence we cannot identify jobs that are subject to shut-downs or jobs that have been defined as critical. The next section will give details on this dataset.

### 3 Data

With the introduction of the Covid-19 Study in April 2020 (ISER, 2020a,b),<sup>4</sup> participants from the main USoc sample have been asked to complete a short web-based survey each month, eliciting information on the impact of the pandemic on their work and family lives. These data have some clear strengths. First, they record retrospective information on outcomes of interest at baseline, i.e. before the onset of the pandemic, as of January-February 2020,<sup>5</sup> as well as contemporaneous information from April onwards. Second, individual records can be linked to past and future waves of the annual USoc survey, facilitating long-run analyses of Covid-19 impacts. Third, selective non-response can be tracked down to a rich set of individual characteristics (available from the earlier USoc annual waves) and accounted for using the weights provided.

We use information from the first two Covid-19 monthly surveys, which were carried out between 24-30 April and between 27 May and 2 June, respectively, among all USoc participants who had responded in at least one of the two previous annual surveys (wave

---

<sup>4</sup>This is available through the UK Data Service (SN8644)

<sup>5</sup>For simplicity we will refer to the baseline period as January 2020.

Table 1: Characteristics of survey respondents

% share of respondents	(1)	(2)	(3)	(4)
	Wave 9 respondents	Covid-19 Study respondents	Difference (1)-(2)	(p-value)
Female	55.71	57.96	-2.25	(0.00)
16-19	5.81	3.94	1.87	(0.00)
20-29	12.13	10.26	1.87	(0.00)
30-29	14.21	14.67	-0.46	(0.17)
40-49	17.57	19.02	-1.45	(0.00)
50-59	18.77	21.63	-2.86	(0.00)
60+	31.50	30.47	1.03	(0.02)
College and above	29.84	37.50	-7.66	(0.00)
Non-British ethnicity	22.50	16.49	6.01	(0.00)
Married	54.04	59.54	-5.50	(0.00)
Living as couple	9.85	10.96	-1.11	(0.00)
Never married	22.42	18.02	4.40	(0.00)
Working at wave nine	57.91	63.71	-5.79	(0.00)
Human health and social work	6.51	7.59	-1.08	(0.00)
Public administration and defence	3.92	4.95	-1.02	(0.00)
Accommodation and food service	2.84	2.44	0.40	(0.01)
HH income quintile 1	13.36	9.48	3.87	(0.00)
HH income quintile 2	16.73	15.01	1.72	(0.00)
HH income quintile 3	20.06	20.75	-0.68	(0.07)
HH income quintile 4	22.66	24.54	-1.87	(0.00)
HH income quintile 5	24.76	28.20	-3.43	(0.00)
HH income quintile unknown	2.43	2.03	0.40	(0.01)
Children aged 15 and below	25.43	26.05	-0.62	(0.13)
Age of youngest child	7.14	7.21	-0.07	(0.36)
<i>N</i>	32,596	16,934		

Notes: The table compares the characteristics of individuals who gave a full adult interview in USoc wave nine with the subset who also gave a full or partial interview in wave one or two of the Covid-19 Study. All individual and household characteristics are measured in wave nine. P-value of two-sample *t*-test for equal means in parenthesis (column 4). Source: USoc wave nine and Covid-19 Study, waves one and two.

nine and ten, carried out in 2017-18 and 2018-19, respectively). The response rates are 46% and 48.5% in the April and May waves, respectively (slightly rising to 48.6% and 49.1% if one includes partial responses), among those who responded in wave nine. Compared to the 86% response rate in wave nine, relative to wave eight participants, retention in the Covid-19 study is considerably lower.<sup>6</sup>

To get a sense of selective attrition in the data, Table 1 compares descriptive statistics among USoc wave nine and Covid-19 waves one and two respondents (data from

<sup>6</sup>Because the Covid-19 study is treated as an instrument of the wave nine annual interview, respondents who did not complete a wave nine interview are assigned a zero weight. More information on how weights were developed can be found in ISER (2020b).

wave ten will only be released in November 2020). Covid-19 respondents are on average slightly older, more likely to be female, college educated, British, married, employed at wave nine and higher earners. To facilitate population inferences, the Covid-19 Study provides weights to account for differential selection probabilities and nonresponse. These are based on information from wave nine, allowing to estimate differential response conditional on a very rich set of individual and household characteristics. All descriptive evidence and regression results based on Covid-19 data below are obtained using such weights.<sup>7</sup>

For validation, we compare retrospective information on labour market outcomes in the Covid-19 Study with information from the January-March 2020 UK Labour Force Survey. Table 2 reports descriptive statistics on employment rates and working hours for the overall population and for men and women separately. Figures on employment rates are remarkably close across the two data sources, but there are some slight differences in working hours. A potential reason for small divergences is that the weighted Covid-19 data provide estimates that are representative of the UK adult population as of the USoc wave nine, which was conducted during 2017 and 2018. This implies that the weights provided might not be exactly representative of the adult population in 2020.

To describe Covid-19 impacts on labour market outcomes, we select all individuals aged 16 to 64 who participated in at least one of the Covid-19 waves and who had previously participated in USoc wave nine ( $N = 10,703$ ). We restrict further to those who report being employed as of January 2020 ( $N = 8,362$ ), and drop individuals with missing information on age, education, region, or basic employment characteristics, leaving us with a sample of 8,073 individuals. Descriptive statistics for this sample are reported in Table 3.

The analysis of outcomes relating to domestic work and childcare combines the Covid-19 Study with pre-Covid data sources. For information on hours of housework, we select individuals living as a couple, who participated in at least one of the Covid-19 waves. Having dropped individuals with missing information on basic individual or employment controls, or missing information on household composition, we are left with a sample of 10,643 individuals (17,614 observations). Pre-Covid information on housework is obtained from USoc wave eight (2016-17,  $N = 17,610$ ).

<sup>7</sup>Indeed, we found some of the results on gender differences to be sensitive to the use of weights (unweighted estimates not reported).

Table 2: Employment statistics validation

	(1) LFS Jan-Mar 2020	(2) USoc Jan-Feb 2020	(3) Difference (1)-(3)
<i>All</i>			
Employment rate (%)	76.60 [42.34]	77.38 [41.84]	-0.78 (0.06)
Weekly working hours	31.77 [17.09]	34.44 [12.51]	-2.67 (0.00)
<i>Men</i>			
Employment rate (%)	79.87 [40.09]	80.80 [39.39]	-0.92 (0.14)
Weekly working hours	35.95 [16.54]	38.22 [11.20]	-2.27 (0.00)
<i>Women</i>			
Employment rate (%)	73.35 [44.21]	74.17 [43.77]	-0.82 (0.14)
Weekly working hours	27.36 [16.54]	30.56 [12.60]	-3.20 (0.00)

Notes: Standard deviations in brackets. All figures are obtained using weights. P-value of two-sample  $t$ -test for equal means in parenthesis (column 3). Weekly hours in the LFS correspond to total actual hours worked in the reference week. Weekly hours in the Covid-19 Study correspond to usual hours worked in January and February 2020. All samples include individuals aged 16-64. Source: UK Labour Force Survey January-March 2020 for column 1; USoc Covid-19 Study, waves one and two, for column 2.

For the analysis of childcare we further select individuals with children aged 15 and below, leaving us with a sample of 3,384 individuals and 5,384 observations across the two waves of the Covid-19 study. For the pre-Covid period, we only have limited information on childcare provision, as individuals are only asked in USoc wave eight about the person mainly in charge of childcare in their household ( $N=5,892$ ), and no information on childcare hours is provided. We therefore compare descriptive statistics on childcare hours during Covid-19 to corresponding statistics from the latest UK Time Use Survey (2014-15). Table 4 provides detailed summary statistics for variables regarding non-market work.

## 4 Labour market outcomes

The (short-run) impact of the pandemic on the labour market can be assessed by comparing information on outcomes as of the April and May 2020 survey dates with retrospective

Table 3: Summary Statistics - Market work

	(1)	(2)	(3)
	All	Male	Female
Age (years)	42.06	42.47	41.64
College and above (%)	37.99	36.86	39.14
Female (%)	49.55	0.00	100.00
Children aged 0-15 (%)	37.49	37.71	37.28
<b>Labour market characteristics Jan/Feb 2020:</b>			
Working (%)	100.00	100.00	100.00
Weekly working hours	34.57	38.44	30.63
Weekly earnings (£)	418.38	489.48	344.85
<i>Type of employment:</i>			
Employed (%)	85.94	83.24	88.69
Self-employed (%)	11.09	13.58	8.55
Both employed and self-employed (%)	2.97	3.18	2.76
<i>Worked from home:</i>			
Always (%)	5.18	4.75	5.63
Often (%)	5.88	6.49	5.26
Sometimes (%)	17.54	18.05	17.01
Never (%)	71.40	70.71	72.10
<i>Contract type:</i>			
Fixed hours (%)	67.69	63.50	71.95
Fixed salary (%)	60.14	59.34	60.95
Paid by hours worked (%)	24.44	22.01	26.90
<b>Labour market outcomes April/May 2020:</b>			
Ever job loss since baseline (%)	4.27	4.41	4.13
Ever furloughed since baseline (%)	28.59	30.14	27.10
Reduced hours (%)	49.62	48.87	50.35
Reduced earnings (%)	36.40	38.60	34.24
Weekly working hours	23.25	25.82	20.76
Change in working hours	-11.20	-12.62	-9.83
Weekly earnings (£)	382.54	443.23	323.45
Change in weekly earnings (£)	-36.35	-50.31	-22.67
<i>No. Individuals</i>	8,073	3,389	4,684

The sample includes individuals aged 16-64 who were employed in January-February 2020 and have no missing control variables. Summary statistics are derived using cross-sectional weights. Source: USoc, wave nine, and Covid-19 Study, waves one and two.

information referring to January 2020. Among individuals who report being employed in January 2020 (including employees and the self-employed), about 4.3% report being out of work by May 2020, including involuntary separations and quits.<sup>8</sup> Most of the

<sup>8</sup>For comparison, ONS (2020b) estimates of UK employment based on the Quarterly Labour Force Survey show only slight variations in employment and unemployment rates between the first and second quarter of the year. Despite relatively flat unemployment figures, the number of people claiming benefits roughly doubled from 1.3 to 2.7 million between March and May 2020, corresponding to 3.5 and 7.4 of the workforce, respectively, in large part due to enhancements to Universal Credit coverage, which made a higher share of workers eligible for unemployment-related benefits while still in work.



Table 4: Summary Statistics - Non-market work

	(1)	(2)	(3)	(4)	(5)	(6)
	All couples			Couples w children aged 0-15		
	All	Male	Female	All	Male	Female
<b>Household work:</b>						
Weekly hours housework (Covid-19)	12.77	9.78	15.83	13.83	10.89	16.80
<i>N</i>	17,614	7,939	9,675	5,402	2,266	3,136
Weekly hours housework (USoc wave 8)	10.25	6.46	14.11	10.90	6.63	15.14
<i>N</i>	17,610	8,365	9,245	5,821	2,681	3,140
<b>Childcare hours:</b>						
Weekly hours childcare (Covid-19)				20.58	14.76	26.50
<i>N</i>				5,384	2,262	3,122
Weekly hours childcare (UK TUS 2014-15)				12.43	7.81	16.99
<b>Main responsible childcare (Covid-19):</b>						
Mainly self				34.85	18.60	57.49
Mainly partner				39.44	55.41	17.20
Shared				20.82	21.03	20.54
Couple reports 0 hours childcare in total				4.88	4.96	4.77
<i>N</i>				3,147	1,557	1,590
<b>Main responsible childcare (USoc wave 8):</b>						
Mainly self				29.70	2.60	56.81
Mainly partner				22.10	41.34	2.86
Shared				47.94	55.68	40.20
Someone else				0.26	0.39	0.13
<i>N</i>				5,892	2,718	3,174

Source: Understanding Society wave eight and nine and COVID-19 Study, waves one and two. UK Time Use Survey 2014/15.

Notes: Columns (1)-(3) refer to the sample of individuals who are living in couple, columns (4)-(6) refer to the subsample of those with children aged 15 and below. Summary statistics are derived using cross-sectional weights.

adjustment in working hours during the downturn has taken place via furloughing under the Coronavirus Job Retention Scheme introduced on 20 March 2020<sup>9</sup> and, according to Government guidance, those on furlough are classified as being employed. By the end of May, about 29% of employees report having ever been furloughed in the Covid-19 Study, in line with evidence collected by ONS (2020a) in a survey of businesses.<sup>10</sup> Over the same period, working hours among those employed in January 2020 fell by 11.2 hours weekly

<sup>9</sup>This provides grants to employers to pay 80% wages to furloughed employees, up to a cap of £2,500 per person per month.

<sup>10</sup>This overall picture is in contrast with corresponding evidence for the US, where furloughing was much less prevalent and the overall employment rate fell by about 11 percentage points according to the Current Population Survey and 19 percentage points percentage points according to the Real-Time Population Survey (Bick and Blandin, 2020), where the difference between the two figures in large part accounts for the number of individuals who are employed but not at work in the reference week, and hence akin to being furloughed.

on average and earnings fell by 9.5% on average (or £36).

Table 5: Job loss during COVID-19

	(1)	(2)	(3)	(4)	(5)	(6)
	All	All	All	All	Males	Females
Female	-0.00355 (0.0100)	-0.00525 (0.00946)	-0.00457 (0.00887)	-0.00528 (0.0101)		
Living as a couple		-0.0165 (0.0110)	-0.0124 (0.0108)	-0.0151 (0.00905)	-0.0141 (0.0159)	-0.0155 (0.0109)
Has children age 0-4		-0.00106 (0.0124)	-0.000859 (0.0122)	0.00506 (0.0116)	-0.00139 (0.0166)	0.0129 (0.0174)
Has children age 5-15		0.0120 (0.0142)	0.0112 (0.0137)	0.00753 (0.0119)	0.00779 (0.0118)	0.00507 (0.0196)
Always WFH			-0.00106 (0.0190)	-0.00164 (0.0150)	0.00891 (0.0298)	-0.00788 (0.0134)
Often WFH			-0.00886 (0.0125)	-0.00629 (0.0131)	0.00303 (0.0214)	-0.00985 (0.0136)
Sometimes WFH			-0.0136* (0.00629)	-0.0109 (0.00682)	-0.00837 (0.00957)	-0.00653 (0.00976)
Constant	0.0473*** (0.00860)	0.108* (0.0463)	0.130* (0.0572)	0.0508 (0.0524)	0.101 (0.122)	0.0354 (0.0602)
Observations	8073	8073	8073	8073	3389	4684
Age and Education	No	Yes	Yes	Yes	Yes	Yes
Region FE	No	Yes	Yes	Yes	Yes	Yes
Job characteristics	No	No	Yes	Yes	Yes	Yes
Occupation FE	No	No	No	Yes	Yes	Yes
Industry FE	No	No	No	Yes	Yes	Yes

Notes. The dependent variable is equal to one if the individual reports to be non-employed in April or May 2020, and zero otherwise. Non-employment in April is treated as an absorbing state. Age controls are dummy variables for ages 20-29, 30-39, 40-49, 50-59, 60+ (16-19 is the excluded category); education controls are dummy variables for GCSEs or equiv., A-levels or equiv., and college education or higher; job characteristics are indicators for self employment, fixed hours, fixed salary and paid by the hour; occupation and industry fixed-effects are at the 2-digit level. All covariates refer to January 2020, except education, occupation and industry, which are imported from USoc wave nine (2017-18). All specifications control for an April wave dummy and use cross-sectional weights. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Sample: all employed individuals in January 2020, aged 16-64. Source: Understanding Society (wave nine) and Covid-19 Study (waves one and two).

Table 5 shows results from linear probability models for the incidence of job loss. The dependent variable is equal to one for all individuals who report being out of work in either April or May 2020, having reported to be in work in January. We treat job loss as an absorbing state, and the sample contains one observation per individual. Regressions control for a set of individual and job-related characteristics. Most characteristics are recorded in the Covid-19 Study and refer to January 2020. Whenever relevant controls are not available in the Covid-19 Study, as is the case for education, industry and occupation, we use information recorded in wave nine.

The specification in column 1 only controls for a female dummy (and a dummy for

survey wave), and shows evidence of virtually no gender differences in the likelihood of job loss. In column 2, we control for household composition (as well as age, education and region dummies), in column 3 we introduce job controls, including WFH habits, and in column 4 we additionally control for two-digit industry and occupation, the finest classification available in USoc.<sup>11</sup> While there is some indication that WFH at baseline reduces the probability of job loss (from column 3), the incidence of job loss is very similar across genders. This result is in contrast with evidence for the US reported by [Alon et al. \(2020\)](#), who find much larger (and unprecedented) unemployment increases among women than men. It also somewhat differs from evidence based on real-time data for the UK analysed by [Adams-Prassl et al. \(2020\)](#), who find that women are about 2-3 percentage points more likely to report job losses than men, having controlled for individual and job characteristics. Columns 5 and 6 report results from separate regressions for men and women and show no evidence of gender differences in the impacts of household and job characteristics.

Table 6 reports corresponding evidence for the incidence of furloughing among employees. The dependent variable is equal to one for individuals who report having ever been furloughed by May 2020. The raw gender differential reported in column 1 implies that women are nearly three percentage points less likely to be furloughed than men, but this effect is imprecisely estimated. The gender differential turns positive when including occupation and industry controls in column 4, consistent with lower furlough incidence in female-dominated jobs, but again the associated coefficient does not reach standard significance levels. As one would expect, the likelihood of furloughing is negatively and strongly correlated with the incidence of WFH before Covid-19, both in the whole sample (columns 3 and 4) and for each gender taken separately (columns 5 and 6).

We next present evidence on working hours in Table 7. As changes in hours may not be absorbing states, we exploit information on hours contained in each Covid-19 wave, and the sample includes repeated observations for individuals who responded in both waves. Panel A estimates a linear probability model for reduced working hours

---

<sup>11</sup>WFH variables refer to how often an individual was working from home in January 2020. While the frequency of WFH is directly related to the share of job tasks that can be performed remotely, it is also affected by other personal and workplace factors, thus it is not directly comparable to the WFH definition that we have used to classify jobs in Figure 1. Unfortunately, the coarser occupational classification available in the USoc Covid-19 Study does not allow us to implement the [Dingel and Neiman \(2020\)](#) classification of jobs that can be performed from home.

Table 6: Ever furloughed during Covid-19

	(1)	(2)	(3)	(4)	(5)	(6)
	All	All	All	All	Males	Females
Female	-0.0291 (0.0168)	-0.0283 (0.0164)	-0.0277 (0.0164)	0.0174 (0.0163)		
Living as a couple		-0.0252 (0.0197)	-0.00844 (0.0199)	-0.0103 (0.0176)	-0.0224 (0.0289)	0.00259 (0.0219)
Has children age 0-4		0.00236 (0.0298)	0.00231 (0.0288)	0.0216 (0.0256)	0.0435 (0.0342)	-0.00154 (0.0362)
Has children age 5-15		0.0209 (0.0198)	0.0226 (0.0200)	0.0265 (0.0172)	0.0176 (0.0253)	0.0362 (0.0220)
Always WFH			-0.121*** (0.0310)	-0.118*** (0.0315)	-0.137** (0.0428)	-0.0930* (0.0433)
Often WFH			-0.0790* (0.0309)	-0.0586* (0.0271)	-0.0303 (0.0424)	-0.0760* (0.0319)
Sometimes WFH			-0.107*** (0.0175)	-0.0923*** (0.0169)	-0.0676* (0.0272)	-0.109*** (0.0203)
Constant	0.295*** (0.0142)	0.577*** (0.0833)	0.663*** (0.0922)	0.182 (0.114)	0.167 (0.192)	0.540* (0.230)
Observations	7118	7118	7118	7118	2878	4240
Age and Education	No	Yes	Yes	Yes	Yes	Yes
Region FE	No	Yes	Yes	Yes	Yes	Yes
Job characteristics	No	No	Yes	Yes	Yes	Yes
Occupation FE	No	No	No	Yes	Yes	Yes
Industry FE	No	No	No	Yes	Yes	Yes

Notes. The dependent variable is equal to one if the individual reports to be furloughed in April or May 2020, and zero otherwise. Furloughing in April is treated as an absorbing state. Age controls are dummy variables for ages 20-29, 30-39, 40-49, 50-59, 60+; education controls are dummy variables for GCSEs or equiv., A-levels or equiv., and college education or higher; job characteristics are indicators for self employment, fixed hours, fixed salary and paid by the hour; occupation and industry fixed-effects are at the 2-digit level. All covariates refer to January 2020, except education, occupation and industry, which are imported from USoc wave nine (2017-18). All specifications control for an April wave dummy and use cross-sectional weights. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Sample: all employees in January 2020, aged 16-64. Source: Understanding Society (wave nine) and Covid-19 Study (waves one and two).

among those employed in January 2020, whether they are fully employed, furloughed or nonemployed. In May 2020, About 51% of men and 52% of women report reduced weekly hours (from column 1), but the gender differential is not statistically significant. Only when controlling for job characteristics in columns 3 and 4 does the gender differential become significant (and larger in magnitude). As expected, the likelihood of working shorter hours increases with the presence of young kids in the household, and decreases with WFH habits. While women are more likely to experience hours losses, their average hours reduction is smaller than for men, as shown in Panel B. The raw differential is about 2.8 weekly hours (column 1) and shrinks by about a half when controlling for the full set of job characteristics (column 4).

Evidence on changes in earnings is shown in Table 8. Differently from hours losses, raw

Table 7: Working hours during Covid-19

	Panel A: Incidence of reduced weekly hours					
	(1)	(2)	(3)	(4)	(5)	(6)
	All	All	All	All	Males	Females
Female	0.0148 (0.0153)	0.0168 (0.0150)	0.0367* (0.0146)	0.0310* (0.0157)		
Living as a couple		-0.0319 (0.0190)	-0.0219 (0.0180)	-0.0194 (0.0158)	-0.0261 (0.0257)	-0.0108 (0.0194)
Has children age 0-4		0.0605* (0.0254)	0.0615** (0.0236)	0.0705** (0.0216)	0.0767* (0.0306)	0.0607* (0.0301)
Has children age 5-15		0.0280 (0.0177)	0.0228 (0.0166)	0.00963 (0.0153)	-0.0166 (0.0224)	0.0378 (0.0204)
Always WFH			-0.0836** (0.0267)	-0.0730** (0.0270)	-0.0678 (0.0427)	-0.0593 (0.0350)
Often WFH			-0.0784** (0.0260)	-0.0531* (0.0246)	-0.0395 (0.0359)	-0.0618 (0.0333)
Sometimes WFH			-0.0595*** (0.0172)	-0.0380* (0.0165)	-0.0290 (0.0247)	-0.0444* (0.0214)
Constant	0.507*** (0.0133)	0.605*** (0.0712)	0.711*** (0.0806)	0.100 (0.147)	0.133 (0.230)	0.0729 (0.195)
	Panel B: Change in weekly hours					
	(1)	(2)	(3)	(4)	(5)	(6)
Female	2.782*** (0.594)	2.643*** (0.583)	2.108*** (0.584)	1.371* (0.613)		
Living as a couple		1.375 (0.740)	0.985 (0.723)	1.147 (0.633)	1.300 (1.095)	0.672 (0.737)
Has children age 0-4		-0.583 (0.957)	-0.594 (0.933)	-0.970 (0.814)	-1.189 (1.345)	-0.537 (0.871)
Has children age 5-15		-0.570 (0.669)	-0.467 (0.633)	0.189 (0.576)	0.484 (0.921)	-0.254 (0.667)
Always WFH			5.648*** (1.131)	4.927*** (1.119)	5.204** (1.894)	4.215** (1.355)
Often WFH			2.896** (1.046)	1.944* (0.970)	0.730 (1.534)	2.948** (1.118)
Sometimes WFH			3.279*** (0.633)	2.524*** (0.597)	2.318* (0.929)	2.460*** (0.716)
Constant	-11.89*** (0.535)	-12.06*** (2.981)	-16.42*** (3.306)	-1.020 (6.273)	2.880 (9.833)	4.514 (7.416)
Observations	14133	14133	14133	14133	5860	8273
Age and Education	No	Yes	Yes	Yes	Yes	Yes
Region FE	No	Yes	Yes	Yes	Yes	Yes
Job characteristics	No	No	Yes	Yes	Yes	Yes
Occupation FE	No	No	No	Yes	Yes	Yes
Industry FE	No	No	No	Yes	Yes	Yes

Notes. The dependent variable in Panel A is equal to one if the individual reports fewer working hours in April-May than in January 2020 and zero otherwise; in Panel B it is equal to the change in weekly hours between January and April-May 2020. Each individual contributes a number of observations equal to the number of Covid-19 waves s/he participated to. Age controls are dummy variables for ages 20-29, 30-39, 40-49, 50-59, 60+; education controls are dummy variables for GCSEs or equiv., A-levels or equiv., and college education or higher; job characteristics are indicators for self employment, fixed hours, fixed salary and paid by the hour; occupation and industry fixed-effects are at the 2-digit level. All covariates refer to January 2020, except education, occupation and industry, which are imported from USoc wave nine (2017-18). All specifications control for an April wave dummy and use cross-sectional weights. Standard errors are clustered at the individual level. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Sample: all employed individuals in January 2020, aged 16-64. Source: Understanding Society (wave nine) and Covid-19 Study (waves one and two).

Table 8: Earnings losses during Covid-19

Panel A: Incidence of earnings losses						
	(1)	(2)	(3)	(4)	(5)	(6)
	All	All	All	All	Males	Females
Female	-0.0437** (0.0156)	-0.0387** (0.0149)	-0.0188 (0.0142)	0.00677 (0.0152)		
Living as a couple		-0.0119 (0.0195)	-0.00107 (0.0177)	-0.0110 (0.0159)	0.00474 (0.0259)	-0.0135 (0.0196)
Has children age 0-4		0.0448 (0.0276)	0.0441 (0.0248)	0.0478* (0.0224)	0.0478 (0.0327)	0.0435 (0.0281)
Has children age 5-15		0.0194 (0.0184)	0.0119 (0.0175)	0.00969 (0.0163)	0.00658 (0.0234)	0.0166 (0.0216)
Always WFH			-0.0524 (0.0295)	-0.0417 (0.0291)	-0.0509 (0.0462)	-0.0197 (0.0340)
Often WFH			-0.0212 (0.0257)	-0.00886 (0.0254)	0.0316 (0.0372)	-0.0363 (0.0333)
Sometimes WFH			-0.0295 (0.0166)	-0.0197 (0.0160)	-0.0286 (0.0241)	-0.00310 (0.0210)
Constant	0.417*** (0.0137)	0.466*** (0.0734)	0.669*** (0.0839)	0.246 (0.178)	0.253 (0.289)	0.173 (0.150)
Panel B: Change in weekly earnings						
	(1)	(2)	(3)	(4)	(5)	(6)
Female	27.62** (8.552)	27.79** (8.666)	21.97* (8.590)	16.05* (7.337)		
Living as a couple		13.16 (11.79)	13.38 (11.61)	14.56 (10.40)	20.16 (21.95)	9.879 (9.509)
Has children age 0-4		-4.685 (13.89)	-5.270 (13.77)	-11.27 (12.07)	-2.700 (21.85)	-22.81* (9.589)
Has children age 5-15		-11.19 (10.14)	-8.053 (9.445)	-6.640 (8.002)	-12.95 (12.66)	-3.377 (8.597)
Always WFH			11.54 (17.13)	16.54 (19.29)	36.21 (31.04)	-8.794 (23.29)
Often WFH			-14.13 (14.39)	-0.693 (20.14)	-12.51 (37.77)	12.85 (12.11)
Sometimes WFH			0.187 (10.55)	7.982 (13.00)	0.795 (24.32)	6.740 (8.307)
Constant	-37.63*** (10.77)	-24.88 (30.79)	-86.61* (34.80)	238.2 (219.8)	167.4 (280.9)	86.66 (45.56)
Observations	12813	12813	12813	12813	5337	7476
Age and Education	No	Yes	Yes	Yes	Yes	Yes
Region FE	No	Yes	Yes	Yes	Yes	Yes
Job characteristics	No	No	Yes	Yes	Yes	Yes
Occupation FE	No	No	No	Yes	Yes	Yes
Industry FE	No	No	No	Yes	Yes	Yes

Notes. The dependent variable in Panel A is equal to one if the individual reports lower weekly earnings in April/May than in January 2020 and zero otherwise; in Panel B it is equal to the change in weekly earnings between January and April/May 2020. Each individual contributes a number of observations equal to the number of Covid-19 waves s/he participated to. Age controls are dummy variables for ages 20-29, 30-39, 40-49, 50-59, 60+; education controls are dummy variables for GCSEs or equiv., A-levels or equiv., and college education or higher; job characteristics are indicators for self employment, fixed hours, fixed salary and paid by the hour; occupation and industry fixed-effects are at the 2-digit level. All covariates refer to January 2020, except education, occupation and industry, which are imported from USoc wave nine (2017/18). All specifications control for a April wave dummy and use cross-sectional weights. Standard errors are clustered at the individual level. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Sample: all employed individuals in January 2020, aged 16-64. Source: Understanding Society (wave nine) and Covid-19 Study (waves one and two).


earnings losses are less frequent among women (Panel A, column 1), but this differential is fully explained by job characteristics (columns 3 and 4). Moreover, earnings' losses for women are on average smaller than for men (Panel B). Using estimates for raw differences in column 1 of Panel B, men lose on average about 38 £per week, corresponding to a 7.7% fall with respect to their January 2020 earnings. Women lose on average 10 £per week, corresponding to about 3% of their January 2020 earnings. More than 40% of this differential is explained by job characteristics (column 4).

In summary, we find evidence that labour market outcomes of men and women were roughly equally affected at the extensive margin, as measured by the incidence of job loss or furloughing, but if anything, women suffered smaller losses at the intensive margin, experiencing slightly smaller changes in both working hours and earnings. This finding is broadly in line with evidence from administrative data on the claimant count, which includes both those out of work and those working on low earnings or hours, and thus represents an indicator of overall economic disadvantage for those who participate in the labour force. Between March and May 2020, the proportion of the male labour force in the claimant count rose from 3.9% to 8.6%, while the corresponding figure for women rose from 3.1% to 6.1%, thus showing a slightly higher increase for men than for women, both in absolute and relative terms. These gender differentials are also echoed by information on welfare receipt from the Covid-19 Study: 4.5% and 3.3% of men and women, respectively, report to have applied for Universal Credit since January 2020, with 3.5% and 2.5%, respectively, already in receipt by May 2020.

## 5 Home production

Measuring changes in home production (and childcare in particular) during the pandemic is complicated by the fact that the Covid-19 Study does not contain retrospective information on these variables. We thus benchmark information provided in the Covid-19 Study to comparable information from previous USoc surveys and the UK Time Use Survey. As we are primarily interested in the gender division of work within the household, we restrict our working sample to heterosexual couples, whether married or cohabiting.

Figure 3: Gender gaps in housework hours, before and during Covid-19



Notes. The bars show usual weekly hours spent on housework before Covid-19 (2016-17) and during Covid-19 (April-May 2020). Sample: men and women living in couple. Source: USoc wave nine and Covid-19 Study, waves one and two.

The Covid-19 questionnaire covers several aspects of domestic work, including hours spent on housework (cooking, cleaning and doing the laundry) and hours spent on child-care (including home schooling). Figure 3 gives a snapshot of the gender division of housework before and during Covid-19. Information on housework for the earlier period is available from USoc wave eight. This shows that, in 2016-17, women were doing just over 14 hours of housework weekly, while men were doing about 6.5 hours. Corresponding figures for the Covid-19 period have risen to about 16 and 10 hours for women and men, respectively. The overall amount of housework for the average 2-adult household has thus increased by about 25%, with a higher absolute and proportional increase for men, and a reduction of the corresponding gender gap from 7.6 to 6 hours.

Regression results reported in Table 9 show that the gender gap in housework hours is only slightly affected by individual and job characteristics, whether before or during Covid-19 (see columns 1-3 in Panels A and B, respectively). While the overall gender gap in housework hours has fallen during Covid-19, it remains more sensitive to the presence of children, own employment status and partner’s employment status for women than for men, as shown in columns 4 and 5.




Table 9: Hours spent on housework before and during Covid-19

Panel A: Hours of housework in 2016-17					
	(1)	(2)	(3)	(4)	(5)
Female	7.651*** (0.131)	7.808*** (0.130)	7.320*** (0.135)		
Has children age 0-4		1.981*** (0.241)	1.828*** (0.241)	0.854** (0.267)	2.836*** (0.411)
Has children age 5-15		2.418*** (0.181)	2.255*** (0.182)	0.755*** (0.193)	3.962*** (0.318)
Not working (wave 8)			4.049*** (0.227)	2.727*** (0.273)	4.849*** (0.343)
Not working (wave 8) (partner)			-1.099*** (0.189)	-0.846*** (0.203)	-1.293*** (0.351)
Constant	6.461*** (0.0730)	0.564 (1.196)	0.440 (1.379)	2.589 (2.018)	6.468*** (1.609)
Observations	17610	17610	16035	7829	8206
Panel B: Hours of housework in April/May 2020					
	(1) All	(2) All	(3) All	(4) Males	(5) Females
Female	6.057*** (0.258)	6.226*** (0.253)	6.531*** (0.287)		
Has children age 0-4		2.626*** (0.500)	2.351*** (0.582)	1.727* (0.747)	3.221*** (0.893)
Has children age 5-15		2.239*** (0.349)	2.275*** (0.389)	1.715*** (0.496)	2.889*** (0.617)
Furloughed			1.864*** (0.521)	1.112 (0.689)	3.004*** (0.766)
Not working (Jan 2020)			2.807*** (0.445)	2.079** (0.648)	3.477*** (0.594)
Furloughed (partner)			-1.392** (0.499)	-1.315* (0.602)	-1.417 (0.764)
Not working (Jan 2020) (partner)			-0.820 (0.444)	-0.821 (0.588)	-1.271* (0.607)
Constant	9.986*** (0.205)	0.385 (1.339)	6.357*** (0.920)	8.162*** (1.314)	11.23*** (1.201)
Observations	17614	17614	11628	5712	5916
Individual controls	No	Yes	Yes	Yes	Yes
Own job characteristics	No	No	Yes	Yes	Yes
Partner job characteristics	No	No	Yes	Yes	Yes

Notes. The dependent variable is the number of weekly hours spent on housework, measured in USoc wave eight (Panel A) and in Covid-19 waves one and two (Panel B). In Panel B, each individual contributes a number of observations equal to the number of Covid-19 waves s/he participated to. Age controls are dummy variables for ages 20-29, 30-39, 40-49, 50-59, 60+ (16-19 is the excluded category); education controls are dummy variables for no qualifications, GCSEs or equiv., A-levels or equiv., and college education or higher. Own employment controls are indicators for being employed, frequency of working from home dummies, 2-digit industry and occupation dummies. All job-related dummies have an extra category for non-employed individuals. Partner's employment controls are only available for those whose partners gave a full interview. In Panel B, all covariates refer to January 2020, except education, occupation and industry, which are imported from USoc wave nine (2017-18). All specifications control for an April wave dummy and use cross-sectional weights. Standard errors are clustered at the individual level. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Sample: all individuals living in couple. Source: Understanding Society (wave nine) and Covid-19 Study (waves one and two).

Figure 4: Gender gaps in childcare hours, before and during Covid-19



Notes. The bars show usual weekly hours spent on childcare and home schooling before Covid-19 (2014-15) and during Covid-19 (April-May 2020). Sample: men and women living in couple, with children aged 15 and below. Source: UK Time Use Survey 2014-15 and Covid-19 Study, waves one and two.

To show evidence on changes in childcare hours (including home schooling), we combine information in the Covid-19 study with comparable information from the UK Time Use Data from 2014-15. Figure 4 shows average weekly childcare hours by gender, before and after Covid-19. The sample refers to couples with children aged 15 and under. The first salient fact is the sharp increase in total childcare time, from nearly 25 to over 41 hours weekly. In 2014-15, mothers were doing on average 17 hours of childcare, while fathers were doing just under 8 hours. In 2020, mothers’ hours have risen to 26.5, and fathers’ hours have risen to 14.8. Women take on board a higher share of increased childcare needs than men (9.5 extra hours as opposed to 6.9 extra hours), with a corresponding increase in the gender differential from 9.2 hours in 2014-15 to 11.7 hours in 2020 (more than offsetting the fall in the gender differential in housework time).

Table 10 shows evidence on the determinants of the childcare differential. Controlling for individual characteristics and own job characteristics in column 2 explains about one hour of the overall differential, and controlling for own and partner’s employment status in column 3 explains nearly another hour.

While there is no earlier information on childcare hours in the USoc, wave eight respondents are asked about the main provider of childcare in their household (with possible answers being: (a) mainly self, (b) mainly partner, (c) shared, (d) someone

Table 10: Hours spent on childcare and home schooling during Covid-19

	(1)	(2)	(3)	(4)	(5)
	All	All	All	Males	Females
Female	11.72*** (1.118)	10.73*** (1.101)	9.978*** (1.438)		
Has children age 0-4		13.89*** (1.265)	12.48*** (1.498)	9.199*** (1.838)	16.77*** (2.485)
Furloughed			4.357* (1.860)	7.665** (2.473)	-0.557 (2.467)
Not working (Jan 2020)			6.225* (2.580)	7.727 (5.387)	6.082* (3.066)
Furloughed (partner)			0.778 (2.048)	-1.361 (1.969)	3.781 (3.703)
Not working (Jan 2020) (partner)			-2.397 (2.185)	-2.559 (2.418)	-8.095* (3.380)
Constant	14.08*** (0.940)	1.400 (14.25)	9.054 (5.077)	18.19* (7.113)	4.341 (5.867)
Observations	5384	5384	3348	1647	1701
Individual controls	No	Yes	Yes	Yes	Yes
Own job characteristics	No	No	Yes	Yes	Yes
Partner job characteristics	No	No	Yes	Yes	Yes

Notes. The dependent variable is the number of weekly hours spent on childcare and home schooling, measured in April-May 2020. Each individual contributes a number of observations equal to the number of Covid-19 waves s/he participated to. Age controls are dummy variables for ages 20-29, 30-39, 40-49, 50-59, 60+ (16-19 is the excluded category); education controls are dummy variables for no qualifications, GCSEs or equiv., A-levels or equiv., and college education or higher. Own employment controls are indicators for being employed, frequency of working from home dummies, 2-digit industry and occupation dummies. All job-related dummies have an extra category for non-employed individuals. Partner's employment controls are only available for those whose partners gave a full interview. All covariates refer to January 2020, except education, occupation and industry, which are imported from USoc wave nine (2017-18). All specifications control for a April wave dummy and use cross-sectional weights. Standard errors are clustered at the individual level. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Sample: individuals living in couple, with children aged  $\leq 15$ . Source: Understanding Society (wave nine) and Covid-19 Study (waves one and two).

else.) We create comparable information in the Covid-19 survey for households in which both partners answer the question on total childcare time. We define the main provider of childcare in April 2020 as (a) mainly self, if the respondent does 60% or more of the total reported childcare hours for the couple; (b) mainly partner, if the respondent does less than 40% of the total; (c) shared, if the respondent does between 40% and 60% of the total.

Column 1 in Panel A of Table 11 shows that about 57% of women were the main providers of childcare in 2016-17, against about 2.6% of men. Just over 10% of this

Table 11: Parent mainly in charge of childcare before and during Covid-19

Panel A: Main childcare provider in 2016-17					
	(1)	(2)	(3)	(4)	(5)
Female	0.542*** (0.011)	0.540*** (0.011)	0.484*** (0.012)		
Has children age 0-4		0.052*** (0.013)	0.053*** (0.013)	0.008 (0.010)	0.097*** (0.027)
Not working (wave 8)			0.189*** (0.018)	0.114*** (0.027)	0.217*** (0.025)
Not working (wave 8) (partner)			-0.104*** (0.013)	-0.025** (0.008)	-0.261*** (0.039)
Constant	0.026*** (0.003)	0.373*** (0.028)	0.223*** (0.033)	0.003 (0.015)	0.574*** (0.054)
Observations	5892	5892	5397	2596	2801
Panel B: Main childcare provider in April-May 2020					
	(1) All	(2) All	(3) All	(4) Males	(5) Females
Female	0.389*** (0.024)	0.390*** (0.023)	0.345*** (0.025)		
Has children age 0-4		0.011 (0.026)	-0.000 (0.026)	-0.004 (0.031)	0.023 (0.040)
Furloughed			0.105*** (0.031)	0.101* (0.040)	0.120** (0.043)
Not working (Jan 2020)			0.155*** (0.043)	0.297** (0.090)	0.120* (0.047)
Furloughed (partner)			-0.139*** (0.030)	-0.061 (0.038)	-0.226*** (0.046)
Not working (Jan 2020) (partner)			-0.133*** (0.034)	-0.105** (0.039)	-0.208* (0.088)
Constant	0.198*** (0.021)	0.177 (0.091)	0.232** (0.087)	0.293** (0.100)	0.487*** (0.118)
Observations	3147	3147	3147	1557	1590
Individual controls	No	Yes	Yes	Yes	Yes
Own job characteristics	No	No	Yes	Yes	Yes
Partner job characteristics	No	No	Yes	Yes	Yes

Notes. The dependent variable is equal to one if the respondent is the main childcare provider and zero otherwise, measured in USoc wave eight (Panel A) and in Covid-19 waves one and two (Panel B). This information is elicited directly in USoc wave 8, while it is obtained from reported hours of childcare and home schooling in the Covid-19 Study. We define the main provider as (a) mainly self, if the respondent does 60% or more of the total reported childcare hours for the couple; (b) mainly partner, if the respondent does less than 40%. In panel B each individual contributes a number of observations equal to the number of Covid-19 waves s/he participated to. Age controls are dummy variables for ages 20-29, 30-39, 40-49, 50-59, 60+ (16-19 is the excluded category); education controls are dummy variables for GCSEs or equiv., A-levels or equiv., and college education or higher. Own employment controls are indicators for being employed, frequency of working from home dummies, 2-digit industry and occupation dummies. All job-related dummies have an extra category for non-employed individuals. Partner's employment controls are only available for those whose partners gave a full interview. In Panel B, all covariates refer to January 2020, except education, occupation and industry, which are imported from USoc wave nine (2017-18). All specifications control for an April wave dummy and use cross-sectional weights. Standard errors are clustered at the individual level. \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Sample: individuals living in couple, with children aged  $\leq 15$ . Source: Understanding Society (waves eight and nine) and Covid-19 Study (waves one and two).

differential is explained by differences in the employment status of parents (column 3), while individual characteristics hardly make a difference (column 2). The change brought about by the pandemic is striking, with about one fifth of fathers mainly in charge of childcare in May 2020, against a roughly unchanged proportion of mothers (column 1, Panel B). With the adjustment to the lockdown, fathers seem to have taken over some of the childcare previously outsourced to the market or to extended family members, without directly biting into mothers' exclusive share of childcare.

While the best part of the additional childcare load has on average been taken over by mothers – largely according to pre-existing specialization patterns of spouses – the share of households in which the father is the main childcare provider has risen by nearly 8 times from 2.6% in 2016-17 to about 20% in May 2020. Distributional aspects of increased childcare needs are thus important to understand changes in gender roles during Covid-19. Gender differences in the role played by observable characteristics are also noteworthy. At baseline, both the presence of young children and the employment status of spouses has a much stronger impact on mothers' likelihood of being mainly in charge of childcare, as opposed to fathers'. Such gender differences are milder during the Covid-19 period, and in particular, being out of work is nudging fathers to be in charge of childcare more than mothers.

## 6 Conclusion

The recession caused by the pandemic has produced unprecedented economic losses and it has become clear that its effect has exacerbated existing inequalities along a number of dimensions, most notably socio-economic status and ethnicity, and have created new divides, for example between those who can work from home and those who cannot. Evidence on the gender dimension is somehow mixed. We find that, in what concerns the labour market, men and women experience similar employment losses or furloughing in the UK, although women suffer slightly smaller hour and earning losses overall. In the household, however, women provide on average for about 60% of increased childcare needs, implying a widening of pre-existing inequalities of parental roles. As school and nursery closures are ongoing in a number of countries around the world, including the US, women's increased care burden may build into longer-lasting inequalities via reduced labour market involvement. In this respect, prioritizing school openings over other sectors

subject to restrictions and introducing subsidies for individuals with care responsibilities could help alleviate some of the gendered effects of Covid-19.

Several of the impacts discussed are temporary in nature and can in principle be reversed with the end of the restrictions and the restart of usual economic activity. But given the recent radical changes to the organisation of work and family life, it is natural to reflect on potentially legacies of the crisis, via learning, habit formation and the evolution of social norms. First, the massive increase in the incidence of WFH has suddenly accelerated a pre-existing but slowly evolving tendency towards smart working and flexible work arrangements. The number of people working from home in the UK has risen from 2.9 million in 1998 to 4.2 million in 2014, representing 14% of employment, and an additional 1.8 million people report they would prefer to work from home if they were given the chance. According to a recent survey (CIPD (2020)), 86% of UK managers foresee organisational barriers to the adoption of flexible working in their workplaces. But it is possible that some of the perceived barriers will be eventually cracked by actual remote work patterns implemented during Covid-19. The demand for remote work varies across genders, with 48% of women employed in jobs that can be done from home in the UK, compared with 39% of men. Due to heavier household responsibilities, women also value flexible work schedules and shorter commutes more than men (Mas and Pallais, 2017; LeBarbanchon et al., 2020), and thus may be more beneficially affected by remote work opportunities. But while WFH may provide women with the flexibility to combine market work and family commitments, it may also dilute employee presence and attachment to the workplace, with possibly detrimental effects on career progression.

Second, the pandemic crisis has witnessed the reversal of traditional gender roles in a sizeable share of UK households, in which fathers took the role of primary childcare providers. There is evidence that the spousal allocation of childcare is shaped in large part by social norms on gender roles, and that gender identity norms are only slowly evolving (Bertrand, 2018). But evidence has also shown that “forced” changes in gender roles may have permanent consequences beyond short-term circumstances, by accelerating the evolution of norms and eroding gender comparative advantages. For example, the mobilisation of men during the Second World War in the United States induced more women to enter the labour market, and thereby shaped the norms and preferences of younger generations who were exposed to those early labour market entrants (Fernández

et al., 2004). Relatedly, there is evidence that the introduction of fathers' quotas of parental leave has induced them to spend more time with their children in the longer run in some (though not all) contexts (see [Farré et al. \(2020\)](#); [Farré and González \(2019\)](#); [Patnaik \(2019\)](#) and references therein). One may therefore expect that the substantial redistribution of childcare involvement in nearly a fifth of UK households during the crisis may ease the breakdown of traditional gender roles come the recovery. We leave this analysis to future research.

## References

- Adams-Prassl, Abigail, Teodora Boneva, Marta Golin, and Christopher Rauh**, “Inequality in the Impact of the Coronavirus Shock: Evidence from Real-Time Surveys,” *Journal of Public Economics*, 2020, 189.
- Alon, Titan, Mattias Doepke, Jane Olmstead-Rumsey, and Michèle Tertilt**, “This Time It’s Different: The Role of Women’s Employment in a Pandemic Recession,” *CEPR Discussion Paper 15149*, 2020.
- Andrew, Alison, Sarah Cattan, Monica Costa Dias, Christine Farquharson, Lucy Kraftman, Sonia Krutikova, Angus Phimister, and Almudena Sevilla**, “The Gendered Division of Paid and Domestic Work under Lockdown,” *Covid Economics*, 2020, pp. 109–138.
- Bertrand, Marianne**, “The Glass ceiling,” *Economica*, 2018, 85, 205–231.
- Bick, Alexander and Adam Blandin**, “Real-Time Labor Market Estimates During the 2020 Coronavirus Outbreak,” *Mimeo*, 2020.
- Blundell, Richard, Monica Costa-Dias, Robert Joyce, and Xiaowei Xu**, “COVID-19 and Inequalities,” *Fiscal Studies*, 2020, 41, 291,319.
- Boca, Daniela Del, Noemi Oggero, Paola Profeta, and Mariacristina Rossi**, “Women’s Work, Housework and Childcare, before and during COVID-19,” *Review of Economics of the Household*, 2020.
- Cabinet Office and Department for Education**, “Guidance for schools, childcare providers, colleges and local authorities in England on maintaining educational provision,” <https://www.gov.uk/government/publications/coronavirus-covid-19-maintaining-educational-provision/guidance-for-schools-colleges-and-local-authorities-on-maintaining-educational-provision>, 19 March 2020.
- **and Ministry of Housing, Communities and Local Government**, “Closing certain businesses and venues,” <https://www.gov.uk/government/publications/further-businesses-and-premises-to-close/further-businesses-and-premises-to-close-guidance-work-carried-out-in-peoples-homes>, 9 April 2020.
- CIPD**, “Megatrends: Flexible working,” Technical Report 2020.
- Dingel, Jonathan and Brent Neiman**, “How Many Jobs Can Be Done at Home?,” *Journal of Public Economics*, 2020, 189.
- Farré, Lidia and Libertad González**, “Does paternity leave reduce fertility?,” *Journal of Public Economics*, 2019, 172, 52,66.
- **, Yarine Fawaz, Libertad González, and Jennifer Graves**, “How the COVID-19 Lockdown Affected Gender Inequality in Paid and Unpaid Work in Spain,” *IZA Discussion Paper 13434*, 2020.



- Fernández, Raquel, Alessandra Fogli, and Claudia Olivetti**, “Mothers and Sons: Preference Formation and Female Labor Force Dynamics,” *Quarterly Journal of Economics*, 2004, 119 (4), 1249–1299.
- Hupkau, Claudia and Barbara Petrongolo**, “Work, care and gender during the Covid-19 crisis,” *CEP Covid-19 analysis No. 002*, 2020.
- ISER**, “Understanding Society: COVID-19 Study,” 2020, *SN: 8644, 10.5255/UKDA-SN-8644-1* (1st Edition).
- , “Understanding Society COVID-19 User Guide,” Technical Report 2020.
- Kleven, Henrik, Camille Landais, and Jakob Egholt Sogaard**, “Children and Gender Inequality: Evidence from Denmark,” *American Economic Journal: Applied Economics*, 2019, 37, 181–209.
- LeBarbanchon, Thomas, Roland Rathelot, and Alexandra Roulet**, “Gender Differences in Job Search: Trading off Commute Against Wage,” *CEPR Discussion Paper 15181*, 2020.
- Mas, Alex and Amanda Pallais**, “Valuing Alternative Work Arrangements,” *American Economic Review*, 2017, 107, 3722–3759.
- ONS**, “Coronavirus and the economic impacts on the UK: 18 June 2020,” Technical Report 2020.
- , “Labour market overview, UK: August 2020,” Technical Report 2020.
- Oreffice, Sonia and Quintana-Domeque**, “Gender Inequality in COVID-19 Times: Evidence from UK Prolific Participants,” *mimeo*, 2020.
- Patnaik, Ankita**, “Reserving Time for Daddy: The Consequences of Fathers’ Quotas,” *Journal of Labor Economics*, 2019, 11, 1009–1059.
- Sevilla, Almudena and Sarah Smith**, “Baby steps: The gender division of childcare during the COVID19 pandemic,” *mimeo*, 2020.

# Uncertainty, hysteresis and lockdowns

Charles Sims<sup>1</sup> and David Finnoff<sup>2</sup>

Date submitted: 23 October 2020; Date accepted: 25 October 2020

*Using currently available data we develop a benchmark scenario that finds that a publicly-imposed social distancing rule to curb the spread of COVID-19 ought to be implemented for 12 weeks. We also identify alternative scenarios where the shorter duration social distancing programs seen throughout the United States may be efficient. Our approach is novel in that it accounts for uncertainty in transmission of the disease and the potential for permanent economic effects of social distancing rules. The social distancing rule is treated as an asset whose benefit is uncertain due to the inability to predict the evolution of the disease. The novel features of our approach allow us to draw two conclusions about the efficient timing of public social distancing programs in response to COVID-19. First, uncertainty in transmission leads to a risk premium that creates a modest incentive to delay closing and reopening the economy. Second, hysteresis in economic impacts of social distancing leads to hysteresis in the efficient time to reopen the economy. Because reopening results in a second wave of infections that may be worse than expected and social distancing rules create permanent economic impacts, it is efficient to delay reopening the economy longer than suggested by benefit-cost analyses of social distancing programs. In our benchmark scenario, this bias results in reopening 27 days too soon.*

- 1 Faculty Fellow at the Howard H. Baker Jr. Center for Public Policy and Associate Professor in the Department of Economics at the University of Tennessee.
- 2 Professor in the Department of Economics, University of Wyoming.

Copyright: Charles Sims and David Finnoff

## 1 Introduction

To date, social distancing rules (e.g., stay at home orders, shelter in place orders, non-essential business closures) have been the primary public policy tool used to respond to the COVID-19 outbreak. These “one size fits all” policies require a severe reduction in economic activity where nonessential businesses are closed and individuals are encouraged to stay home to minimize contacts between susceptible and infected individuals and slow the spread of the disease. Governments are now grappling with a fundamental question: When should the economy reopen? Reopening too soon risks an increase in infected cases and a second round of social distancing. Reopening too late risks unnecessary economic impacts. Further, these decisions will have long-lasting effects and must be made in the face of considerable uncertainty.

We treat the social distancing program as a risky asset whose benefit is uncertain due to the inability to predict the evolution of the disease. We identify the economically efficient point in the COVID-19 outbreak to implement public social distancing rules and close the economy. In addition, we are able to find the point at which to suspend social distancing rules and reopen the economy. We extend the framework of Sims, et al. (2016) in three major steps. First, we use an SIR compartmental model instead of SIS which adds a second stochastic variable and greatly complicates the analysis and solution methodology. Second, instead of sunk costs of public health programs, permanent effects are caused by deferred consumption losses that never materialize after reopening. Third, we identify and estimate two risk adjustments that influence when to close and reopen the economy. Using currently available data, a benchmark scenario finds a social distancing program ought to be implemented for 86 days and results in control of the virus (defined to be less than 1% increase in cumulative cases; Tellis, et al. 2020).

Our results come with a warning. Following the implementation of social distancing, there will be an interval or window of hysteresis during which benefit-cost analysis will advocate the reopening of the economy. However, because reopening results in a second wave of infections that may be worse than expected and social distancing rules have permanent impacts on the economy, it will remain optimal to delay reopening the economy until the recovered population has grown to sufficient size that a second round of social distancing rules is unlikely. In our benchmark specification, the hysteresis window increases the expected duration of efficient social distancing rules by over 30 percent. By comparison, a 20 percent increase in the value of mortality risk reduction (i.e., value of statistical life) increases the duration of social

distancing rules by only 7 percent. Importantly, the higher the percentage of impacts that are permanent the longer the duration of the hysteresis window.

Unlike individual responses to infection risk that create endogenous social distancing (including but not limited to Eichenbaum, et al. 2020; Fenichel 2013; Fenichel, et al. 2011; Toxvaerd 2020), the mass quarantines of public social distancing rules are lumpy and ill-suited for daily adjustments. Firms respond to imposition of stay at home orders and non-essential business closures by laying off or furloughing workers and shuttering facilities. These responses lead to consumption losses as consumers are unable to patronize closed businesses and workers experience reductions in income. Some of the consumption losses triggered by the imposition of social distancing rules may be recouped after the economy is reopened as some businesses reopen and some employees return to work. However, like other rare macroeconomic disasters, there is a real threat that social distancing mandates meant to control COVID-19 may have permanent effects on the economy (Barro, et al. 2020; Bénassy-Quéré, et al. 2020; Fornaro and Wolf 2020). Nakamura, et al. (2013) and Barro and Jin (2020) find that about half of disaster-related declines in consumption since 1900 were permanent. Some businesses will not survive stay at home orders and non-essential business closures and will not reopen after these mandates are lifted.<sup>1</sup> Some workers will not return to the labor force after the economy is reopened.<sup>2</sup>

We show how this hysteresis in the economy bleeds back to create hysteresis in the optimal policy response to COVID-19. If social distancing mandates have permanent effects on the economy, part of the economic impacts of these mandates is sunk. The sunk costs of social distancing imply there will be a “window of inaction” in which a social distancing program should not be implemented but an existing social distancing program should not be suspended (Dixit 1992). The window of inaction, which we term a hysteresis window, highlights how the social distancing mandates have lasting effects on the economy and the optimal policy response to COVID-19 depends on past decisions to implement (or not) social distancing. A critical implication of the hysteresis window is that reopening should be delayed longer than suggested by a discounted cash flow, benefit-cost analysis of the social distancing program.

<sup>1</sup> From March 1 through June 25<sup>th</sup>, of the businesses registered on “yelp.com” that closed, 41% were permanently shut down: <https://www.yelpeconomicaverage.com/yelp-coronavirus-economic-impact-report.html>

<sup>2</sup> Labor force participation has been shown to decline by 7 percentage points due to COVID-19 (Coibion, et al. 2020a).

In addition, we realize there are tremendous uncertainties, problems with the identification of epidemiological parameters, and challenges to accurately estimate the benefits and costs of policies (Gollier 2020; Manski 2020; Pindyck 2020). The inability to predict the benefits and costs of social distancing mandates will increase the window of inaction (Dixit 1992) but its effect on the duration of social distancing mandates is unclear. In response to concerns about the multitude of uncertainties facing model-based policy recommendations for COVID-19 and its implications for optimal policy response to COVID-19, we extend the deterministic approach to consider a setting when the evolution of the outbreak is unpredictable. In this setting, it pays to keep society's options open and delay reopening the economy since reopening results in a second wave of infections, the magnitude of this second wave is uncertain, and a second round of social distancing imposes more sunk costs. This hysteresis window, and what drives it to be longer or shorter, is important when providing policy advice regarding reopening of the economy to the risk of COVID-19.

Uncertainty, sunk costs, and inflexibility characterize most social distancing policies, yet are not the focus of most related work that identifies a top down planner's program of social distancing (e.g., Bolzoni, et al. 2019; Gonzalez-Eiras and Niepelt 2020; Hansen and Day 2011; Lee, et al. 2010; Rowthorn, et al. 2009). Throughout this work, and that of the Ramsey model extensions of Alvarez, et al. (2020), Kruse and Strack (2020), and Piguillem and Shi (2020), optimal time paths of instantaneously adjustable social distancing controls, infections, and deaths are generated for a deterministic outbreak. The results regarding social distancing measures are consistent – let infections rise until they are close to the medical system capacity, implement social distancing measures in an aggressive fashion to keep the number of infections below the medical system's capacity constraint, and then continuously relax social distancing rules resulting in a muted (if any) second wave of infections. This strategy of aggressive social distancing policies followed by frequent marginal adjustments effectively shaves off the top of the infection curve instead of flattening it.

The hysteresis window we highlight results in a departure from recent insights based on instantaneously adjustable policies. In our case of commonly observed, broad scale social distancing policies that are not instantaneously adjustable, the infection curve is flattened instead of being shaved. When considering policies that flatten the curve, decision makers' ought to implement social distancing mandates much earlier than deterministic policies that allow for

marginal adjustments. While the uncertainty in the evolution of the outbreak means it is impossible to avoid the medical system threshold with certainty, the implementation at low prevalence is necessary to flatten the curve out enough to lower the probability the medical system threshold is crossed.

## 2 Methods

We model the spread of COVID-19 through the U.S. population ( $N = 327$  million) using a stochastic version of the standard SIR framework parameterized for the current COVID-19 situation in the U.S.A. Compartmental models of this type track the numbers of susceptible, infected, and recovered individuals over the course of an infectious disease outbreak (Hethcote 2000; Kermack and McKendrick 1927). The transmission rate  $\beta$  is a lumped parameter that reflects the infectiousness of COVID-19 and the contact rate of individuals.  $\beta$  is a random variable, with variability due to environmental stochasticity (more infectious during cold weather) or stochasticity in contact rates (e.g., superspreader events).<sup>3</sup> Following Gray, et al. (2011), the transmission rate over  $dt$  is normally distributed with mean  $\tilde{\beta}dt$  and variance  $\sigma^2dt$ :  $\beta dt = \tilde{\beta}dt + \sigma^2 dz$  where  $dz$  is the increment of a standard Wiener process. This specification implies that the effective reproduction number,  $R_e$ , which is equivalent to the basic reproduction number,  $R_0$ , at the start of the epidemic, evolves stochastically over time.<sup>4</sup> When accounting for this inability to predict COVID-19 transmission, the SIR epidemiological system becomes:

$$dS(t) = -\tilde{\beta}I(t)S(t)dt - \sigma I(t)S(t)dz \quad (1)$$

$$dI(t) = \left[ \tilde{\beta}I(t)S(t) - \frac{\gamma}{1-m(I)}I(t) \right] dt + \sigma I(t)S(t)dz \quad (2)$$

$$dR(t) = \gamma I(t)dt \quad (3)$$

where  $I(t)$ ,  $S(t)$ , and  $R(t)$  are the proportion of the total population  $N$  that is infected, susceptible, and recovered respectively. The stochastic epidemiological system in (1)-(3)

<sup>3</sup> Specifically, the transmission of COVID-19 will vary over time because some cases transmit to many others (superspreader events), while many other cases transmit much less (Kucharski, et al. 2020; Tuite, et al. 2020).

<sup>4</sup> In this case  $R_0 = \beta(1-m)/\gamma$  where  $m$  is the initial probability of mortality in the infected state and  $\gamma$  the duration of infectiousness.

assumes the current proportion of the population infected, susceptible, and recovered is known but the future course of the outbreak is unknown.<sup>5</sup>

**Epidemiological Specification.** Our benchmark parameterization is given in Table 1. We follow the work of Thunstrom, et al. (2020) and assume  $R_0$  is 2.4 as seen in China and elsewhere (Ferguson, et al. 2020; Liu, et al. 2020). The average infectious period is taken to be 6.5 days:  $\gamma = 1/6.5$  (Lauer, et al. 2020; Liu, et al. 2020). The probability an infected individual dies before recovering is  $m(I)$ , and taken to be responsive to the evolution of infected cases. As the number of infected individuals requiring medical care grows, the health care system becomes stressed, health care resources become scarcer, and infected individuals receive a lower standard of care. We model this medical system feedback as a logistic function (Newbold, et al. 2020):

$$m(I) = m_l + \frac{m_h - m_l}{1 + e^{-g(I(t) - \bar{I})}} \quad (4)$$

where  $g$  is large so that the relationship is a step function with  $m_l$  the probability an infected individual dies before recovering when the number of infected cases is below the health care system capacity ( $I < \bar{I}$ ) and  $m_h > m_l$  when the number of infected cases is above the health care system capacity ( $I > \bar{I}$ ). Following Thunstrom, et al. (2020), we assume that the health care system has sufficient resources to provide adequate treatment for one half of the maximum number of individuals who would be infected at any one time with no social distancing to slow the spread of the virus. In the benchmark case, this threshold is 35.3 million infected people (approximately 11 percent of the population). When the health care system is not overwhelmed, the probability an infected individual dies before recovering is 0.5%, which is greater than seasonal influenza. When the health care system is overwhelmed, the probability an infected individual dies before recovering is 1.5%. Together, these parameter values imply an expected transmission rate of  $\tilde{\beta} = 0.37$ . We set the standard deviation of the transmission rate ( $\sigma$ ) to 0.15 which is consistent with how transmission varied over time in Wuhan, China between January, 2020, and February, 2020 (Kucharski, et al. 2020).

<sup>5</sup> Our focus on uncertainty in transmission abstracts from uncertainty in the current number of infected individuals and the mortality rate which we leave for future work.

Table 1. Parameter values for benchmark specification

Parameter	Definition	Parameter value
<b>Economics</b>		
$K$	Deferred consumption losses due to social distancing	\$2.1 trillion
$k_S$	Non-deferable consumption losses due to social distancing (\$ per susceptible individual per day)	\$35
$k_R$	Non-deferable consumption losses due to social distancing (\$ per recovered individual per day)	\$0
$c_I$	Daily consumption reduction per infected due to outbreak	\$55
$c_S$	Daily consumption reduction per susceptible due to outbreak	\$50
$\rho$	Discount rate	0.03
$\mu$	Percent of deferred consumption losses recouped after reopening the economy	0.75
$v$	Value of mortality risk reduction	\$7 million
$\tilde{\beta}_{sd}$	Mean transmission rate during social distancing	0.230*
$\sigma_{sd}$	Standard deviation of transmission rate during social distancing	0.093
<b>Epidemiology</b>		
$\bar{I}$	Infected threshold where health care system is stressed	0.108
$m_l$	Probability of mortality when health care system is not stressed	0.005
$m_h$	Probability of mortality when health care system is stressed	0.015
$\gamma$	Recovery rate	0.154
$\tilde{\beta}$	Mean transmission rate with no social distancing	0.371
$\sigma$	Standard deviation of transmission rate with no social distancing	0.15

\* For the benchmark, the basic reproduction number without social distancing would be  $R_0 = \frac{\beta(1-m)}{\gamma} = \frac{.371(1-0.005)}{.154} \approx 2.4$  and with social distancing it becomes  $R_0^{sd} = \frac{\beta_{sd}(1-m)}{\gamma} = \frac{.230(1-0.005)}{.154} \approx 1.49$

**Economic Specification.** The welfare implications of the virus can be lost consumption or lost lives as documented by Martin and Pindyck (2020). We extend this notion to account for the evolution of these elements over the course of the pandemic as well as with, and without, mandated social distancing. The stream of COVID-19 related impacts incurred by society is  $W(S, I, X)$  where  $X$  is an indicator of the current status of mandated social distancing. If  $X = 1$ , mandated social distancing is not currently in effect (i.e., the economy is open and the expected transmission rate is  $\tilde{\beta}$ ). The impact of COVID-19 is the value of lives lost plus reduced consumption due to susceptible individuals voluntarily social distancing and infected individuals self-quarantining:  $W(S, I, 1) = vN \frac{dD}{dt} + c_I NI(t) + c_S NS(t)$  where  $\frac{dD}{dt} = \frac{\gamma m}{1-m} I(t)$  is the



proportion of people killed by the disease over  $dt$ ,  $v$  is the value of mortality risk reductions, and  $c_I$  and  $c_S$  are the reductions in consumption per infected and susceptible per day respectively. The value of mortality risk reduction is taken to be at an intermediate level of \$7 million, which is lower than U.S. federal agency guidelines (Kniesner and Viscusi 2019) but higher than several recent papers that adjust the value of mortality risk reduction to account for the age profile of COVID-related deaths (e.g., Greenstone and Nigam 2020). The parameters  $c_S$  and  $c_I$  are calibrated to yield a 7% reduction in consumption in the first year of the epidemic with no social distancing mandates. This is comparable to other estimates in the literature (Eichenbaum, et al. 2020).<sup>6</sup>

Given the implications of the spreading virus, officials can mandate social distancing that effectively closes the economy when disease prevalence reaches a threshold level  $I_c$ . The program instantly lowers the expected transmission rate from  $\tilde{\beta}$  to  $\tilde{\beta}_{sd}$  and the standard deviation of the transmission rate from  $\sigma$  to  $\sigma_{sd}$ . Following Thunstrom, et al. (2020) social distancing measures are taken to lower the expected average contact rate among individuals by 38% which is comparable to effectiveness of social distancing during the Spanish flu outbreak in 1918 (Caley, et al. 2008).<sup>7</sup> Social distancing is also expected to lower the volatility of transmission (Kucharski, et al. 2020) and we assume the same 38% reduction in  $\sigma$ .

Mandated social distancing exacerbates consumption losses due to mobility restrictions and the effect of the shutdown on employment, income, and economic expectations (Andersen, et al. 2020; Aum, et al. 2020; Baker, et al. 2020; Coibion, et al. 2020b; Gupta, et al. 2020). Some of the additional consumption losses represent consumption that is deferred until after the economy reopens. For example, Coibion, et al. (2020b) find that consumers under lockdown are 3.5 percentage points less likely to purchase larger ticket items such as automobiles or appliances in the next 12 months. We define these deferred consumption losses as  $K$ . A portion of the deferred consumption losses,  $\mu(S)K$ , is immediately recovered when the economy is reopened

<sup>6</sup> The reduction in consumption from susceptible individuals voluntarily engaging in social distancing,  $c_S$ , is 28% of 2019 U.S. GDP per capita per day (\$50 per susceptible individual per day). The reduction in consumption from infected individuals self-quarantining,  $c_I$ , is 31% of 2019 U.S. GDP per capita per day (\$55 per infected individual per day) to account for the additional hospitalization costs and forced reductions in work hours due to COVID-19 infection.

<sup>7</sup> This is representative of the average reductions in the continuously controlled deterministic literature, for example see Alvarez, et al. (2020), and Gonzalez-Eiras and Niepelt (2020).

with  $\mu(S) \leq 1$  and  $\partial\mu/\partial S > 0$ . For example, after social distancing is suspended and the economy reopens, some (but not all) businesses will reopen, some previously laid off workers will be rehired, and some deferred purchases will be made. However, the longer reopening is delayed, the smaller the susceptible population and the more permanent these deferred consumption losses become. The remainder of the additional consumption losses cannot be deferred until after the economy reopens. For example, lost spending at bars and restaurants due to forced closures or lower incomes due to job losses elsewhere in the economy cannot be recouped by reopening the economy. Because this type of forgone consumption likely depends on the number of susceptible and recovered individuals, we define these non-deferable consumption losses as  $k_S NS + k_R NR(t)$ .

If indicator variable  $X = 2$ , mandated social distancing is currently in effect (the economy is closed and the expected transmission rate is  $\tilde{\beta}_{sd}$ ):  $W(S, I, 2) = \nu N \frac{dD}{dt} + c_I NI(t) + k_S NS + k_R NR(t)$ . The values selected for the three cost parameters ( $K$ ,  $k_S$ ,  $k_R$ ) imply consumption drops by 22% if social distancing mandates were left in place for an entire year which is consistent with the findings in Eichenbaum, et al. (2020) and implies the majority of consumption losses during the outbreak are attributable to social distancing mandates (Andersen, et al. 2020; Coibion, et al. 2020b). We assume that deferable consumption losses represent roughly half of this drop in consumption.<sup>8</sup> While consumer spending increased in May and June as the economy reopened, personal consumption expenditures remain below pre-COVID levels. Specifically, consumption expenditure data suggests  $\mu < 1$  since reopening the economy generated additional consumption that offset only 58% of the consumption losses experienced since the first of the year. This finding is consistent with Barro and Jin (2020) and Nakamura, et al. (2013) which finds that approximately half of disaster-related declines in consumption since 1900 were permanent. However, it remains unclear how the duration of the shutdown influences  $\mu$ . In the absence of data needed to accurately specify  $\mu(S)$ , we make the simplifying assumption that 25 percent of the consumption losses deferred until the economy is reopened

<sup>8</sup> In our benchmark specification, we assume  $K$  is equal to 10 percent of 2019 U.S. GDP (\$2.14 trillion) and  $k_S$  is 20% of 2019 U.S. GDP per capita per day (\$35 reduction in consumption per susceptible individual per day). Our benchmark scenario assumes that those that recover go back to work, so that they are not subject to the lockdown ( $k_R = 0$ ).

become permanent ( $\mu = 0.75$ ). Because these parameters are not well documented, we consider the sensitivity of our primary results to a range of parameter values.

The objective of a social planner is to determine if and when to implement a social distancing rule (close the economy) and suspend the social distancing rule (reopen the economy) to minimize the expected present value of lives lost to COVID-19, the economic impacts of COVID-19, and the economic impacts of the social distancing rule. The problem is one of optimally timing a switch, at time  $t_c$ , between a regime where mandated social distancing is not active ( $X = 1$ ) and a regime where mandated social distancing is active ( $X = 2$ ).<sup>9</sup> The planner must then determine when to move back to an open economy, at time  $t_r$ . Given the discount rate  $\rho$ , the optimal time to close and reopen the economy satisfies

$$V(S_0, I_0, 1) = \min_{t_c} E_0 \left[ \int_0^{t_c} \left[ v \frac{dD}{dt} + c_I I + c_S S \right] N e^{-\rho t} dt + [V(S(t_c), I(t_c), 2) + K] e^{-\rho t_c} \right] \quad (5)$$

and

$$V(S_0, I_0, 2) = \min_{t_r} E_0 \left[ \int_0^{t_r} \left[ v \frac{dD}{dt} + c_I I + k_S S \right] N e^{-\rho t} dt + [V(S(t_r), I(t_r), 1) - \mu K] e^{-\rho t_r} \right] \quad (6)$$

subject to the stochastic evolution of the outbreak in (1)-(3) and the initial condition taken to be cases estimated by the CDC on March 17, 2020:  $I(0) = 4,165/N$ , and  $S(0) = 1 - I(0)$ . The evaluation at each instant in time minimizes the expected impacts of the disease from that point forward by making a choice to continue to delay social distancing (whose payoff is defined as  $V(S, I, 1)$ ) or to take action and lower the spread of the disease through a costly social distancing mandate (whose payoff is defined as  $V(S, I, 2)$ ). When  $\mu < 1$ , closing the economy has a permanent effect and there is a conditional value of information (an option value) that creates an incentive to delay closing the economy longer than suggested by traditional benefit-cost analysis of deterministic discounted cash flows. The delay allows the decision maker to gain valuable information on how bad the outbreak will be. The benefit of this additional information must be balanced by the additional deaths that result from delay.

<sup>9</sup> Our methodology is generalizable to more than two regimes. A third regime could be designed to accommodate a partial reopening where the transmission rate increases but is lower than before the social distancing rule was implemented. For exposition, we assume the transmission rate returns to levels seen before social distancing.

According to Brekke and Øksendal (1994), the optimal switching problem can be rewritten as a set of variational inequalities with a straightforward economic interpretation. The outbreak is treated as a public health obligation whose magnitude  $V(S, I, 1)$  must be minimized.<sup>10</sup> The minimized value of the obligation prior to implementing social distancing satisfies

$$\rho \leq \frac{\left[ v \frac{\gamma m}{1-m} I(t) + c_I I(t) + c_S S(t) \right] N}{V(S, I, 1)} + \left[ \tilde{\beta} I(t) S(t) - \frac{\gamma}{1-m} I(t) \right] \frac{\partial V(S, I, 1) / \partial I}{V(S, I, 1)} - \tilde{\beta} I(t) S(t) \frac{\partial V(S, I, 1) / \partial S}{V(S, I, 1)} + \Omega(S, I, 1) \tag{7}$$

and value matching condition

$$V(S, I, 1) \leq V(S, I, 2) + K \tag{8}$$

with one of the conditions satisfied at each point in the state space of  $S(t)$  and  $I(t)$ . If the rate of return condition in (7) holds as an equality, it is optimal to delay social distancing and remain in the regime with the economy open. The discount rate is the opportunity cost of cash flows associated with delaying social distancing, or the rate of loss a public health official would tolerate to delay social distancing. The right-hand side is the expected (own) rate of loss from delaying social distancing over the interval  $dt$ , with a deterministic component (first three terms)

and a risk adjustment  $\Omega(S, I, 1) = \frac{(\sigma SI)^2}{2} \frac{\left[ \frac{\partial^2 V(S, I, 1)}{\partial I^2} - \frac{\partial^2 V(S, I, 1)}{\partial S^2} \right]}{V(S, I, 1)}$ . Current costs and impacts are given in the first term on the right hand side of (7), the second term is the expected rate of increase in the marginal cost of the outbreak due to additional infections that arise due to delay, and the third term is the expected rate of decrease in the marginal cost of the outbreak due to fewer susceptible individuals. The risk adjustment captures how the inability to predict the evolution of the outbreak alters the expected rate of loss due to delaying social distancing. When  $\Omega(S, I, 1) < 0$ , the risk adjustment lowers the expected rate of loss which requires an increase in the deterministic rate of loss (percentage increase in costs) tolerated by public health officials to delay social distancing and remain in the business as usual regime. The implication is that

<sup>10</sup> In traditional financial or capital accumulation applications, a decision maker chooses when to invest in a financial asset to maximize a measure of return such as profit. The current application is identical except the asset is COVID-19 impacts and the decision maker wishes to minimize this asset.

transmission risk results in the decision maker being willing to absorb higher growth in losses which creates an incentive to delay social distancing – a risk premium. Likewise, when  $\Omega(S, I, 1) > 0$ , the risk adjustment pushes the own rate of loss for delaying social distancing up – a risk reward. Here, the uncertainty in the evolution of the outbreak requires a lower deterministic rate of increase in costs for the decision maker to be willing to delay social distancing, creating an incentive to hasten social distancing. Equation (8) compares the total obligation when social distancing is inactive and active and acts as a boundary condition for the open economy regime. If (8) holds as an equality, it is optimal to immediately implement social distancing and switch to the closed economy regime.

With social distancing rules in place,  $k_R$  taken to be zero, and the economy closed, the minimized value of the outbreak obligation satisfies

$$\rho \leq \frac{\left[ v \frac{\gamma m}{1-m} I(t) + c_I I(t) + k_S S(t) \right] N}{V(S, I, 2)} + \left[ \tilde{\beta}_{sd} I(t) S(t) - \frac{\gamma}{1-m} I(t) \right] \frac{\partial V(S, I, 2) / \partial I}{V(S, I, 2)} - \tilde{\beta}_{sd} I(t) S(t) \frac{\partial V(S, I, 2) / \partial S}{V(S, I, 2)} + \Omega(S, I, 2) \tag{9}$$

and

$$V(S, I, 2) \leq V(S, I, 1) - \mu K. \tag{10}$$

Equation (9) is a rate of return condition that compares the discount rate and the expected rate of loss from delaying suspension of the social distancing program and reopening the economy. If (9) holds as an equality, it is optimal to keep the social distancing program active and remain in the regime with a closed economy. Equation (9) reveals a similar risk adjustment associated

with the decision to reopen the economy  $\Omega(S, I, 2) = \frac{(\sigma_{sd} S I)^2}{2} \frac{\left[ \frac{\partial^2 V(S, I, 2)}{\partial I^2} - \frac{\partial^2 V(S, I, 2)}{\partial S^2} \right]}{V(S, I, 2)}$ . When

$\Omega(S, I, 2) < 0$ , the risk adjustment pushes the expected rate of loss down requiring an increase in the deterministic rate of loss tolerated to delay reopening the economy – a risk premium. This risk premium raises the acceptable increase in loss associated with staying closed, creating an incentive to delay reopening. Likewise, when  $\Omega_2 > 0$ , the risk adjustment pushes the deterministic rate of return for delaying reopening down (a risk reward) which creates an incentive to hasten reopening. Equation (10) acts a boundary condition for the closed regime. If

(10) holds as an equality, it is optimal to suspend the social distancing program and reopen the economy.

In general, the sign of  $\frac{\partial^2 V(S, I, \cdot)}{\partial I^2}$  and  $\frac{\partial^2 V(S, I, \cdot)}{\partial S^2}$  are ambiguous. Stochasticity in transmission impacts both the expected discounted costs of the outbreak and the magnitude of two option values associated with closing and reopening the economy. Whether risk creates a premium or a reward depends on whether the effect of the option value outweighs the possibly countervailing effect of risk on the expected net present value of outbreak impacts. Because the decision to close the economy creates potentially permanent economic impacts,  $V(S, I, 1)$  includes an option value that delays implementing a social distancing mandate. The decision to reopen the economy is also irreversible since a second round of social distancing mandates is necessary if the number of infections after reopening is worse than expected. Because the decision to reopen the economy is irreversible,  $V(S, I, 2)$  includes an additional option value that delays reopening the economy. However, uncertainty in transmission also influences the expected discounted impacts of the outbreak with and without social distancing. If uncertainty in transmission decreases the expected discounted impacts of the outbreak, this effect reinforces the effect of the option value leading to a definitive risk premium. However, if uncertainty in transmission increases the expected discounted impacts of the outbreak, the effect of the option value may be overwhelmed, leading to a risk reward.


The solution to the variational inequalities in equations (7)-(10) can be characterized by a closed economy threshold curve,  $I_c(S)$ , that partitions the state space (combinations of  $I(t)$  and  $S(t)$ ) where social distancing should occur from those where it should not occur, given the social distancing mandate is not currently in place. Likewise, a reopened economy threshold curve  $I_r(S)$  divides the state space into regions where social distancing should and should not occur given the social distancing mandate is currently in place. A hysteresis window emerges when the number of infected cases is in between these two curves:  $I_c(S) > I(t) > I_r(S)$ . The presence of a hysteresis window can be seen by comparing the boundary conditions in equations (8) and (10). These two equations will not generally hold with equality at the same values of  $S(t)$  and  $I(t)$ , suggesting that the boundary of the closed economy and open economy regime will differ ( $I_c(S)$  and  $I_r(S)$  will not overlap). The hysteresis window disappears when the boundary conditions in equations (8) and (10) are equivalent:  $I_c(S) = I_r(S)$ . This equivalence occurs for two cases. The first is when none of the consumption losses due to mandated social distancing

are deferrable ( $K = 0$ ). Here the boundary of the open economy regime and the closed economy regime are the same and satisfy  $V(S, I, 1) = V(S, I, 2)$ . The second is when mandated social distancing has no permanent effects on consumption ( $\mu = 1$ ). Here the boundary of the open economy regime and the closed economy regime are also the same and satisfy  $V(S, I, 1) - K = V(S, I, 2)$ .

The two stochastic variables  $I(t)$  and  $S(t)$  and the dual program regimes require numerical methods to approximate the unknown value functions needed to solve for the policy thresholds and risk adjustments (Judd 1998; Miranda and Fackler 2004). We approximate  $V(S, I, 1)$  and  $V(S, I, 2)$  over a subset of the state space using piecewise linear basis functions. The approximation procedure solves for the basis function coefficients which satisfy equations (7) – (10) at a set of  $n$  nodal points spread evenly over the two-dimensional state-space. Instead of an explicit solution, the threshold curve to close,  $I_c(S)$ , and reopen,  $I_r(S)$ , the economy are sets of  $n$  points where conditions (7) – (10) are met. Details of our computational approach are provided in the appendix.

### 3 Results

Figure 1 shows how a change in the susceptible population (movement along the x-axis) influences the critical levels of infected cases that trigger and cancel social distancing (movement along the y-axis). The boundary of the gray region and the white region in the northeastern portion of the figure is the implementation curve to close the economy,  $I_c(S)$ , while the boundary of the gray region and the white region in the southwestern portion of the figure is the curve to reopen the economy,  $I_r(S)$ , for our benchmark specification. Both curves are downward sloping in the susceptible population. A negative relationship indicates a greater propensity to intervene in the outbreak when the susceptible population and thus force of infection is high. Social distancing is optimal for combinations of  $S(t)$  and  $I(t)$  in the northeastern white region. Reopening the economy is optimal for combinations of  $S(t)$  and  $I(t)$  in the southwestern white region. The gray region in Figure 1 is the hysteresis region. When the outbreak is in this region, it is not optimal to implement a social distancing policy but it is also not yet optimal to suspend a social distancing program that is already in place.



**Figure 1. Optimal timing of social distancing and the effect of hysteresis on reopening.** The critical threshold to implement social distancing and close the economy,  $I_c(S)$ , is the northeastern border of the hysteresis region (gray region). The critical threshold to cancel social distancing and reopen the economy,  $I_r(S)$ , is the southwestern border of the hysteresis region. Dashed lines show the expected trajectory of the outbreak with red (green) lines indicating the economy is closed (open).

The expected trajectory of the outbreak (dashed lines) indicate where we are expected to cross the policy thresholds. Starting from early in the outbreak when the infected population is low and the susceptible population is large (the lower right corner of the figure), it is optimal to delay social distancing until the expected trajectory of the outbreak crosses into the white region (approximately 0.43 percent of the population infected). With social distancing in place, the expected percentage of the population that is infected eventually reaches its peak at 6.5 percent; substantially lower than the peak with no social distancing. It is optimal to terminate social



distancing and reopen the economy when the percent infected falls to 2.9 percent or 49 percent of the population is recovered. A decision maker may be tempted to reopen the economy when the hysteresis window opens just past the peak of the infection curve. This leads to a second wave of infected cases where over 10 percent of the population is infected. Waiting until the hysteresis window closes and the recovered population rises to 49 percent will result in a much smaller second wave of infection and will lower the probability that a second wave of social distancing will be needed.

Table 2 summarizes the results under our benchmark scenario and Figure 2 highlights key dynamic processes. Figure 2A shows the stochastic dynamics of COVID-19 transmission. Before and after social distancing when the economy is open (solid line), the effective reproduction number varies daily between 6 (highly transmissible) and less than 1 (not transmissible) with a mean of 2.4. The spikes in transmission reflect temporary changes in contact rates that may arise due to sporting events, music festivals, or other mass gatherings. For example, a Champions League football match on February 19<sup>th</sup> between Italy's Atalanta and Spain's Valencia has been cited as a major factor in the virus spread to Spain (Rudan 2020). The optimally timed social distancing program lasts 86 days, achieves moderation in 13 days, and control in 84 days following the benchmarks proposed in Tellis, et al. (2020). With social distancing in place (dashed line), the expected reproduction number drops to 1.49 and spikes in the reproduction number are much less severe.

Table 2. Summary of results from benchmark specification

Expected duration of social distancing	86 days
Expected percent of population infected to close economy	0.43%
Expected percent of population recovered to open economy	49%
Hysteresis window	27 days
Economic costs of social distancing	\$3.0 trillion
Value of lives saved	\$15.8 trillion
Risk premium when closing economy	0.28%
Risk premium when reopening economy	0.77%

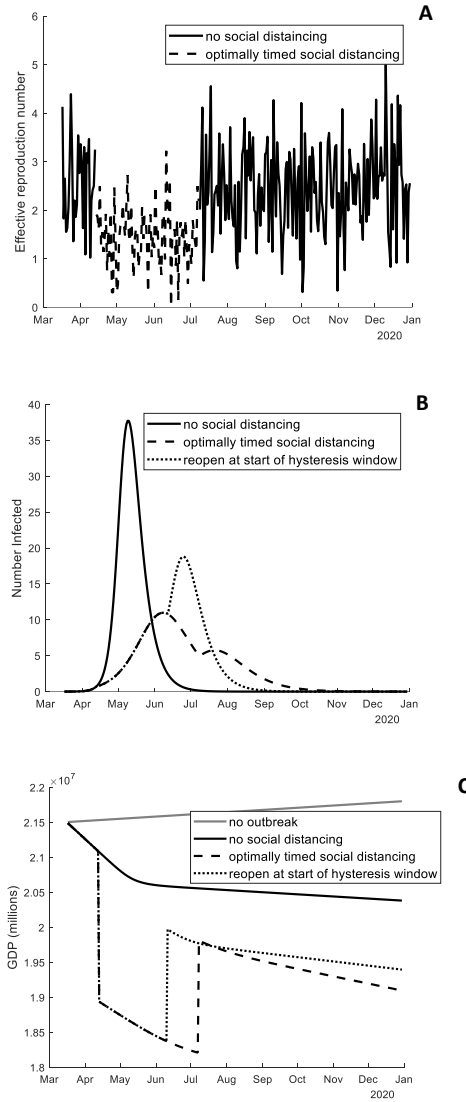


Figure 2. Dynamics of (A) COVID-19 transmission, (B) expected cases of COVID-19, and (C) economic cost of social distancing

Figure 2B, shows the expected number of infected cases without social distancing, with the optimally timed social distancing program, and when the economy is prematurely reopened at the start of the hysteresis window. The optimal “curve flattening” strategy available to decision makers due to the permanent economic impacts finds social distancing rules implemented before the expected sharp increase in infected cases that arises from leaving the outbreak unchecked. This is in sharp contrast to the “curve shaving” policies that delay severe social distancing policies until disease prevalence is closer to the capacity of the medical system and then gradually relax those policies. Suspending social distancing and reopening the economy at the optimal time results in a small secondary peak in infected cases. In contrast, prematurely reopening the economy at the beginning of the hysteresis window can be expected to lead to a larger secondary peak in infected cases. The optimally timed social distancing program lowers the number of infected cases and deaths. The value of the lives saved from this optimally timed social distancing program is \$15.8 trillion.


The optimally timed social distancing program costs \$3 trillion; approximately 1/5<sup>th</sup> the value of lives saved. The time path of the expected total cost of social distancing (shown in Figure 2C) shows a sharp downward drop in GDP due to consumption deferred when social distancing mandates are implemented. Economic impacts accumulate while social distancing is in effect due to the non-deferrable consumption losses. The sharp increase in GDP marks the recovery of a portion of the deferred consumption losses triggered by reopening the economy. Hysteresis in the economy is present since GDP does not return to the trend line before the outbreak.

The uncertainty-induced caution created by unpredictable transmission and the permanent economic impacts incurred to fight the outbreak creates an incentive to delay reopening the economy 27 days. In other words, there will no longer be an economic justification for implementing the social distancing program 59 days after the economy is closed. However, finding that a social distancing program should not be implemented is not a viable justification for canceling a social distancing program that is already in place. Reopening after 59 days ignores the possibility that reopening the economy may lead to a bigger increase in infected cases than expected which would necessitate a second round of social distancing and incur another round of permanent economic impacts.

Consistent with our analysis of the variational inequalities in equation (7) - (10), the hysteresis window vanishes and benefit-cost analysis can provide reasonable guidance on the timing of reopening when mandated social distancing has no permanent effects on the economy. This occurs under two conditions. The first is when none of the consumption losses attributable to social distancing are deferred until the economy reopens. As shown in Figure 3 panel A, the hysteresis window shrinks when deferred consumption makes up a smaller portion of the consumption losses associated with social distancing. The second is when all of the deferred consumption from social distancing is recouped when the economy is reopened (see Figure 3 panel B). In either case, there is no value in delaying reopening because there are no sunk costs associated with a second round of social distancing if the number of cases is higher than expected after reopening. If social distancing is characterized by permanent economic impacts (i.e., all of the deferred consumption does not occur upon reopening the economy), the hysteresis window will likely push the optimal time to reopen the economy far past the point when the number of infected cases peaks.

Also consistent with our analysis of the variational inequalities in equations (7)-(10), the uncertainty in transmission can create incentives to both delay and hasten closing and reopening the economy (see Appendix for details). When focusing on optimal social distancing policies, transmission risk creates an incentive to delay both implementing social distancing and reopening the economy. Due to the stochastic dynamics of the disease, the point at which the disease first crosses  $I_c(S)$  is uncertain. However, at nearly all points along  $I_c(S)$ , transmission risk creates an incentive to delay social distancing. Furthermore, at the values of  $I(t)$  and  $S(t)$  where the expected trajectory of the outbreak crosses this boundary, uncertainty in transmission lowers the own rate of loss from delaying social distancing by 0.28% (a 9% decrease in the expected rate of loss tolerated to delay social distancing). At the values of  $I(t)$  and  $S(t)$  where the expected trajectory of the outbreak crosses  $I_r(S)$ , uncertainty in transmission lowers the own rate of loss from delaying reopening by 0.77% (a 25% decrease in the expected loss tolerated to delay reopening). Depending on the parameters in the model, transmission risk can shift the timing of closing and reopening the economy by 1 to 8 days. However, due to the fact transmission risk creates an incentive to delay both closing and reopening the economy, the uncertainty in

transmission only increases the expected duration of social distancing by 2 days and shortens the expected duration of the hysteresis window by 1 day.<sup>11</sup>



**Figure 3. Relationship between economic irreversibility and the duration of hysteresis window.** Dots represent our benchmark specification.

<sup>11</sup> We characterize the influence on our results of changes in all parameters in the Appendix.

Since mandated social distancing in the United States was shorter than the 12-week program of social distancing suggested in our benchmark model, we identify cases when more rapid reopening could be justified (Table 3). Ineffective policies, significant deferred consumption, low value of mortality risk reduction, and higher expected transmission rates can lower the duration of social distancing to 7-10 weeks. Social distancing programs less than 7 weeks are only consistent with preemptively reopening at the start of the hysteresis window (when it is optimal to not reopen), shown in the second column of Table 3.

Table 3. The case for short duration social distancing

	<b>Optimally reopen at end of hysteresis window</b>	<b>Pre-emptively reopen at beginning of hysteresis window</b>
<b>Ineffective policy:</b> social distancing only reduces $\tilde{\beta}$ by 20 percent	7 weeks	5 weeks
<b>Large deferred consumption:</b> $K$ accounts for over 68 percent of total costs of social distancing	9 weeks	5 weeks
<b>Low value of mortality risk reduction:</b> $v = \$5.6$ million	10 weeks	6 weeks
<b>High expected transmission:</b> $\tilde{\beta} = 0.40$	9.5 weeks	8 weeks

#### 4 Conclusion

Considerable effort has been devoted to estimating the portion of COVID-related economic impacts that are attributable to policy responses (Andersen, et al. 2020;Aum, et al. 2020;Coibion, et al. 2020b;Goolsbee and Syverson 2020;Gupta, et al. 2020;Kahn, et al. 2020;Rojas, et al. 2020). These studies illuminate the costs of social distancing mandates and also indicate the extent to which reopening the economy versus lowering COVID-related deaths alleviates economic impacts.

Our work highlights an additional policy-relevant metric –the degree to which these consumption losses are permanent. The extended period of social distancing mandates undertaken in many parts of the world will likely have permanent economic impacts due to

confidence effects, lost corporate income in the service sectors, bankruptcies, and credit constraints (Bénassy-Quéré, et al. 2020). While these permanent effects have not yet been measured, they are a defining characteristic of other economic disasters (Barro and Ursúa 2008) and have been documented for previous pandemics (Barro, et al. 2020). There is also considerable uncertainty about the severity of a second wave of infection upon reopening the economy and how this uncertainty should influence decisions about reopening.

This paper illustrates how this hysteresis in the economy and uncertainty about the evolution of the outbreak will cascade to cause hysteresis in the optimal policy response to COVID-19. Our paper illustrates one mechanism through which consumption losses may become permanent – deferred consumption. Specifically, we highlight two key parameters that influence the magnitude of the policy hysteresis: 1) the amount of consumption that is deferred until the economy is reopened and 2) the proportion of that deferred consumption that is permanently lost. While benefit-cost analyses have been used as a tool to analyze social distancing rules (Scherbina 2020;Thunstrom, et al. 2020), our results show how basic applications of benefit cost analysis, that omit this hysteresis window, will suggest reopening too soon.

While our paper adds to the literature on the role of uncertainty and learning in public health responses to COVID-19 (Bandyopadhyay, et al. 2020;Davies and Grimes 2020;Gollier 2020;Manski 2020), it only considers the uncertainty in the evolution of the outbreak. Determining how reopening policies should be adjusted to account for uncertainty about the fraction of the population that have had the virus (i.e., state variable uncertainty) and the true infection fatality rate (i.e., parametric uncertainty) remains a critical research need. These sources of uncertainty cannot be characterized through a stochastic epidemiological model and will require approaches with a foundation in Bayesian updating. Instead of the risk adjustments presented in our paper, these other important sources of uncertainty will necessitate a precautionary approach to reopening the economy. Identifying policies that efficiently balance the benefits and costs of this precaution can help guide reopening.

## References

- Alvarez, F.E., D. Argente, and F. Lippi. 2020. "A simple planning problem for covid-19 lockdown." *Covid Economics* 14:1-32.
- Andersen, A.L., et al. 2020. "Consumer responses to the COVID-19 crisis: Evidence from bank account transaction data." *Covid Economics* 7:92-118.
- Aum, S., S.Y.T. Lee, and Y. Shin. "COVID-19 Doesn't Need Lockdowns to Destroy Jobs: The Effect of Local Outbreaks in Korea." National Bureau of Economic Research, No. w27264.
- Baker, S.R., et al. 2020. "How does household spending respond to an epidemic? consumption during the 2020 covid-19 pandemic." *Covid Economics* 18:73-108.
- Bandyopadhyay, S., et al. "Learning or habit formation? Optimal timing of lockdown for disease containment." Department of Economics, University of Birmingham.
- Barro, R.J., and T. Jin. 2020. "Rare events and long-run risks." *Review of Economic Dynamics*.
- Barro, R.J., and J.F. Ursúa. "Macroeconomic crises since 1870." National Bureau of Economic Research, No. w13940.
- Barro, R.J., J.F. Ursúa, and J. Weng. "The coronavirus and the great influenza pandemic: Lessons from the "spanish flu" for the coronavirus's potential effects on mortality and economic activity." National Bureau of Economic Research.
- Bénassy-Quéré, A., et al. (2020) "13 COVID-19: Europe needs a catastrophe relief plan." In R. Baldwin, and B. Weder di Mauro eds. *Mitigating the COVID Economic Crisis: Act Fast Do Whatever It Takes*. London, CEPR Press, pp. 121-128.
- Bolzoni, L., et al. 2019. "Optimal control of epidemic size and duration with limited resources." *Mathematical Biosciences* 315:108232.
- Brekke, K., and B. Øksendal. 1994. "Optimal Switching in an Economic Activity under Uncertainty." *SIAM Journal on Control and Optimization* 32:1021-1036.
- Caley, P., D.J. Philp, and K. McCracken. 2008. "Quantifying social distancing arising from pandemic influenza." *Journal of the Royal Society Interface* 5:631-639.
- Coibion, O., Y. Gorodnichenko, and M. Weber. 2020b. "The cost of the covid-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending." *Covid Economics* 20.
- . 2020a. "Labor markets during the covid-19 crisis: A preliminary view." *Covid Economics* 21:40-58.



- Davies, B., and A. Grimes. 2020. "COVID-19, lockdown and two-sided uncertainty." *New Zealand Economic Papers*:1-6.
- Dixit, A. 1992. "Investment and hysteresis." *Journal of Economic Perspectives* 6:107-132.
- Eichenbaum, M.S., S. Rebelo, and M. Trabandt. "The macroeconomics of epidemics." National Bureau of Economic Research, No. w26882.
- Fenichel, E.P. 2013. "Economic considerations for social distancing and behavioral based policies during an epidemic." *Journal of Health Economics* 32:440-451.
- Fenichel, E.P., et al. 2011. "Adaptive human behavior in epidemiological models." *Proceedings of the National Academy of Sciences* 108:6306-6311.
- Ferguson, N., et al. "Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand." Imperial College London (16-03-2020),doi:<https://doi.org/10.25561/77482>.
- Fornaro, L., and M. Wolf. "Covid-19 coronavirus and macroeconomic policy." CEPR Discussion Paper No. DP14529.
- Gollier, C. 2020. "Pandemic economics: optimal dynamic confinement under uncertainty and learning." *The Geneva Risk and Insurance Review* 45:80-93.
- Gonzalez-Eiras, M., and D. Niepelt. 2020. "On the Optimal" Lockdown" During an Epidemic." *Covid Economics* 7:72-91.
- Goolsbee, A., and C. Syverson. "Fear, lockdown, and diversion: Comparing drivers of pandemic economic decline 2020." National Bureau of Economic Research, No. w27432.
- Gray, A., et al. 2011. "A stochastic differential equation SIS epidemic model." *SIAM Journal on Applied Mathematics* 71:876-902.
- Greenstone, M., and V. Nigam. 2020. "Does social distancing matter?" *Covid Economics* 7:1-22.
- Gupta, S., et al. "Tracking public and private response to the covid-19 epidemic: Evidence from state and local government actions." National Bureau of Economic Research, No. w27027.
- Hansen, E., and T. Day. 2011. "Optimal control of epidemics with limited resources." *Journal of Mathematical Biology* 62:423-451.
- Hethcote, H.W. 2000. "The mathematics of infectious diseases." *SIAM review* 42:599-653.
- Judd, K.L. 1998. *Numerical methods in economics*: MIT press.
- Kahn, L.B., F. Lange, and D.G. Wiczer. "Labor Demand in the time of COVID-19: Evidence from vacancy postings and UI claims." National Bureau of Economic Research, No. w27061.

- Kermack, W.O., and A.G. McKendrick. 1927. "A contribution to the mathematical theory of epidemics." *Proceedings of the royal society of london. Series A, Containing papers of a mathematical physical character* 115:700-721.
- Kniesner, T.J., and W.K. Viscusi. 2019. "The Value of a Statistical Life." *Oxford Research Encyclopedia of Economics and Finance*:19-15.
- Kruse, T., and P. Strack. 2020. "Optimal Control of an Epidemic through Social Distancing." *Covid Economics* 21:168-193.
- Kucharski, A.J., et al. 2020. "Early dynamics of transmission and control of COVID-19: a mathematical modelling study." *The Lancet Infectious Diseases* 20:553-558.
- Lauer, S.A., et al. 2020. "The incubation period of coronavirus disease 2019 (COVID-19) from publicly reported confirmed cases: estimation and application." *Annals of Internal Medicine* 172:577-582.
- Lee, S., G. Chowell, and C. Castillo-Chávez. 2010. "Optimal control for pandemic influenza: the role of limited antiviral treatment and isolation." *Journal of Theoretical Biology* 265:136-150.
- Liu, Y., et al. 2020. "The reproductive number of COVID-19 is higher compared to SARS coronavirus." *Journal of Travel Medicine* 27:taaa021.
- Manski, C.F. 2020. "Forming COVID-19 Policy Under Uncertainty." *Journal of Benefit-Cost Analysis*:1-16.
- Martin, I.W.R., and R.S. Pindyck. "Welfare costs of catastrophes: Lost consumption and lost lives." National Bureau of Economic Research, No. w26068.
- Miranda, M.J., and P.L. Fackler. 2004. *Applied computational economics and finance*: MIT press.
- Nakamura, E., et al. 2013. "Crises and recoveries in an empirical model of consumption disasters." *American Economic Journal: Macroeconomics* 5:35-74.
- Newbold, S.C., et al. 2020. "Effects of Physical Distancing to Control COVID-19 on Public Health, the Economy, and the Environment." *Environmental and Resource Economics* 76:705-729.
- Piguillem, F., and L. Shi. "The optimal covid-19 quarantine and testing policies." Einaudi Institute for Economics and Finance (EIEF).
- Pindyck, R.S. "COVID-19 and the welfare effects of reducing contagion." National Bureau of Economic Research, No. w27121.
- Qi, H.-D., and L.-Z. Liao. 1999. "A smoothing Newton method for extended vertical linear complementarity problems." *SIAM Journal on Matrix Analysis Applications* 21:45-66.

- Rojas, F.L., et al. "Is the cure worse than the problem itself? immediate labor market effects of covid-19 case rates and school closures in the us." National Bureau of Economic Research, No. w27127.
- Rowthorn, R.E., R. Laxminarayan, and C.A. Gilligan. 2009. "Optimal control of epidemics in metapopulations." *Journal of the Royal Society Interface* 6:1135-1144.
- Rudan, I. 2020. "A cascade of causes that led to the COVID-19 tragedy in Italy and in other European Union countries." *Journal of global health* 10:010335-010335.
- Scherbina, A. 2020. "Determining the optimal duration of the COVID-19 suppression policy: A cost-benefit analysis." Available at SSRN 3562053.
- Sims, C., D. Finnoff, and S.M. O'Regan. 2016. "Public control of rational and unpredictable epidemics." *Journal of Economic Behavior & Organization* 132:161-176.
- Tellis, G.J., A. Sood, and N. Sood. 2020. "How Long Should Social Distancing Last? Predicting Time to Moderation, Control, and Containment of COVID-19." *USC Marshall School of Business Research Paper*, Available at SSRN: <https://ssrn.com/abstract=3562996> or <http://dx.doi.org/10.2139/ssrn.3562996>.
- Thunstrom, L., et al. 2020. "The benefits and costs of using social distancing to flatten the curve for COVID-19." *Journal of Benefit-Cost Analysis* 11:179-195.
- Toxvaerd, F. 2020. "Equilibrium social distancing." *Covid Economics* 15:110-133.
- Tuite, A.R., D.N. Fisman, and A.L. Greer. 2020. "Mathematical modelling of COVID-19 transmission and mitigation strategies in the population of Ontario, Canada." *CMAJ* 192:E497-E505.

## Appendix

### *Numerical methodology*

$V(S, I, 1)$  and  $V(S, I, 2)$  are the solutions to the ordinary differential equations in condition (7) and (9). The multiple policy regimes require numerical methods to approximate these unknown value functions (Judd 1998; Miranda and Fackler 2004). Further,  $V(S, I, :)$  can be

approximated over a subset of the state space using piecewise linear basis functions. Specifically, the unknown value functions are approximated with a linear spline constructed using upwind finite difference approximations. The approximated state space extends from 0 to  $N$  in the  $I(t)$  and  $S(t)$  dimensions with 100 nodal points in the  $S$  dimension and 500 nodal points in the  $I$  dimension. Increasing the number of nodal points or extending the state space in either the  $I(t)$  or  $S(t)$  directions does not alter our general results. The resulting complementarity problem is solved in MATLAB using the smoothing-Newton root finding method (Qi and Liao 1999).

### *Sensitivity analysis*

The sensitivity of our results to assumptions about economic and epidemiological parameter values is presented in Tables A1 and A2. In general, changes in parameter values that delay closing the economy also hasten reopening the economy. The optimal expected duration of social distancing is reduced by increases in all costs associated with social distancing and increases in the expected benefits of social distancing (i.e., a higher value of mortality risk reduction and a more effective social distancing program). This intuitive economic finding leads to the counterintuitive conclusion for public policy - the more effective social distancing is, the longer the program should remain in effect. As shown in Figure A1, a more effective and thus more beneficial social distancing program hastens the expected implementation of the program and delays its expected suspension. A more effective policy is implemented sooner than a less effective policy because the value of lives saved and the economic benefits of the policy are higher and thus outweigh the costs of the policy sooner in the outbreak. A more effective policy remains in effect longer than a less effective policy because relaxing an effective social distancing policy implies a larger second wave of infections after the economy is reopened. Any social distancing policy that fails to reduce the expected rate of transmission by less than 12 percent should not be implemented due to the smaller expected benefits this program would provide.

Several key epidemiological parameters alter the efficient timing of a social distancing program. Higher and more uncertain transmission shortens the duration of a public social distancing program, as do lower values of the probability of mortality and a higher recovery rate. In fact, if the probability of mortality is comparable to the seasonal flu ( $m = 0.001$ ), it is not

optimal to ever implement a social distancing program. A higher health care stress threshold lowers the duration of social distancing.

A primary contribution of this paper is that the efficient reopening of the economy may be delayed longer than suggested by benefit-cost analysis due to a hysteresis window. In general, changes in parameter values that increase the optimal duration of social distancing also increase the duration of the hysteresis window. The hysteresis window is most sensitive to the expected transmission rate. The hysteresis window may be as long as 67 days when the expected transmission rate is lower than our benchmark level or as short as 11 days when the expected transmission rate is higher than our benchmark level.

### *Risk adjustments*

Figure A2 shows how the risk adjustment influences the decision to close (Panel A) and reopen (Panel B) the economy. The boundary of the white and green region in Panel A is the optimal policy threshold for closing the economy,  $I_c(S)$ . The boundary of the white and red region in Panel B is the optimal policy threshold for reopening the economy,  $I_r(S)$ . The black lines separate regions of the state space where the risk adjustment is negative (a risk premium) from those where it is positive (a risk reward). Transmission risk creates incentives to delay closing and reopening the economy when the prevalence of the disease is low. The delay is consistent with the option value that arises due to the partial irreversibility of the decisions to close and reopen the economy. When the proportion infected approaches the medical system stress threshold, risk begins to create incentives to hasten social distancing. The increasing probability of additional deaths due to limited medical resources causes the value function to become convex as  $I(t)$  approaches the stress threshold:  $\frac{\partial^2 V(S, I, 1)}{\partial I^2} > 0$ . Through Jensen's inequality, risk increases the expected impacts of the disease with no social distancing which creates an incentive to hasten the start of social distancing rules. Risk also increases the expected impacts of the disease when social distancing is in place through a similar application of Jensen's inequality which creates an incentive to hasten reopening.

With unlimited medical resources, the uncertainty in transmission always creates an incentive to delay implementing social distancing and reopening the economy. When  $m_l = m_h$ ,

$\frac{\partial^2 V(S, I, i)}{\partial I^2} < 0$  for all  $S(t)$  and  $I(t)$  and there is always a risk premium associated with closing the economy and reopening the economy.

Table A1. Percent change in results when economic parameters change by 20 percent

Benchmark specification		Expected duration of social distancing (days)	Threshold % of population infected to close economy	Threshold % of population recovered to reopen economy	Hysteresis window (days)	Risk premium for closing economy (%)	Risk premium for reopening economy (%)
<b>Benchmark specification</b>		<b>86</b>	<b>0.43</b>	<b>49</b>	<b>27</b>	<b>0.28</b>	<b>0.77</b>
Deferred consumption losses	High: $K = \$2.6T$	-20	91	-8	-11	-32	-65
	Low: $K = \$1.7T$	9	0	6	7	-18	43
Non-deferrable consumption losses	High: $k_S = \$43$	-5	26	0	-4	4	18
	Low: $k_S = \$29$	0	0	0	0	-4	-6
Recovered do not return to work	$k_R = \$35$	-5	0	-4	-11	0	-64
Consumption impact per infected	High: $c_I = \$66$	0	0	0	0	0	-4
	Low: $c_I = \$44$	0	0	0	0	0	-4
Consumption impact per susceptible	High: $c_S = \$60$	0	0	0	0	-7	25
	Low: $c_S = \$40$	-5	26	0	-4	4	-19
Deferred consumption recouped	More: $\mu = 0.9$	-13	26	-8	-37	4	-55
	Less: $\mu = 0.6$	9	0	6	33	-18	43
Value of mortality risk reduction	High: $v = \$8.4M$	7	0	4	7	-11	0
	Low: $v = \$5.6M$	-21	91	-10	-7	-36	-60
Effectiveness of social distancing	More: $\tilde{\beta}_{sd} = 0.23^*$ $\sigma_{sd} = 0.07$	56	26	-16	96	-43	-40
	Less: $\tilde{\beta}_{sd} = 0.26^{**}$ $\sigma_{sd} = 0.10$	-29	0	2	-59	-7	-87

\* The equivalent reproduction number is  $R_0^{sd} = \frac{.23(1-.005)}{.154} \approx 1.49$  in this scenario

\*\* The equivalent reproduction number is  $R_0^{sd} = \frac{.26(1-.005)}{.154} \approx 1.68$  in this scenario

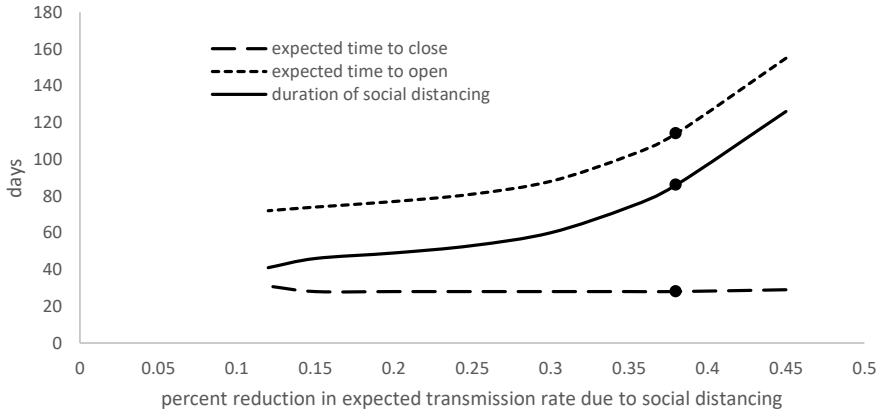
Table A2. Percent change in results when epidemiological parameters change by 20 percent

Benchmark specification		Expected duration of social distancing (days)	Threshold % of population infected to close economy	Threshold % of population recovered to reopen economy	Hysteresis window (days)	Risk premium for closing economy (%)	Risk premium for reopening economy (%)
<b>Benchmark specification</b>		<b>86</b>	<b>0.43</b>	<b>49</b>	<b>27</b>	<b>0.28</b>	<b>0.77</b>
Healthcare stress threshold	High: $\bar{I} = 0.130$	-3	26	0	11	4	-1
	Low: $\bar{I} = 0.086$	0	0	0	-11	-11	0
Probability of mortality	High: $m_l = 0.006$	6	2	4	4	-11	-1
	Low: $m_l = 0.004$	-21	91	-10	-7	-36	-66
Probability of mortality	High: $m_h = 0.018$	0	0	0	0	0	-1
	Low: $m_h = 0.012$	-3	26	0	0	4	0
Recovery rate	High: $\gamma = 0.18$	-7	-12	6	0	182	-126
	Low: $\gamma = 0.12$	0	77	-10	4	-21	-77
Expected transmission	High: $\tilde{\beta} = 0.45^*$	-21	-58	20	-59	193	-90
	Low: $\tilde{\beta} = 0.30^{**}$	5	288	-39	148	-29	43
Transmission volatility	High: $\sigma = 0.18$	-3	26	0	0	61	56
	Low: $\sigma = 0.12$	0	0	0	4	-43	-40

\* The basic reproduction number becomes  $R_0 = \frac{.45(1-.005)}{.154} \approx 2.91$  in this scenario

\*\* The basic reproduction number becomes  $R_0 = \frac{.30(1-.005)}{.154} \approx 1.94$  in this scenario





**Figure A1. Effect of social distancing effectiveness on optimal duration of social distancing program.** Dots represent our benchmark specification.

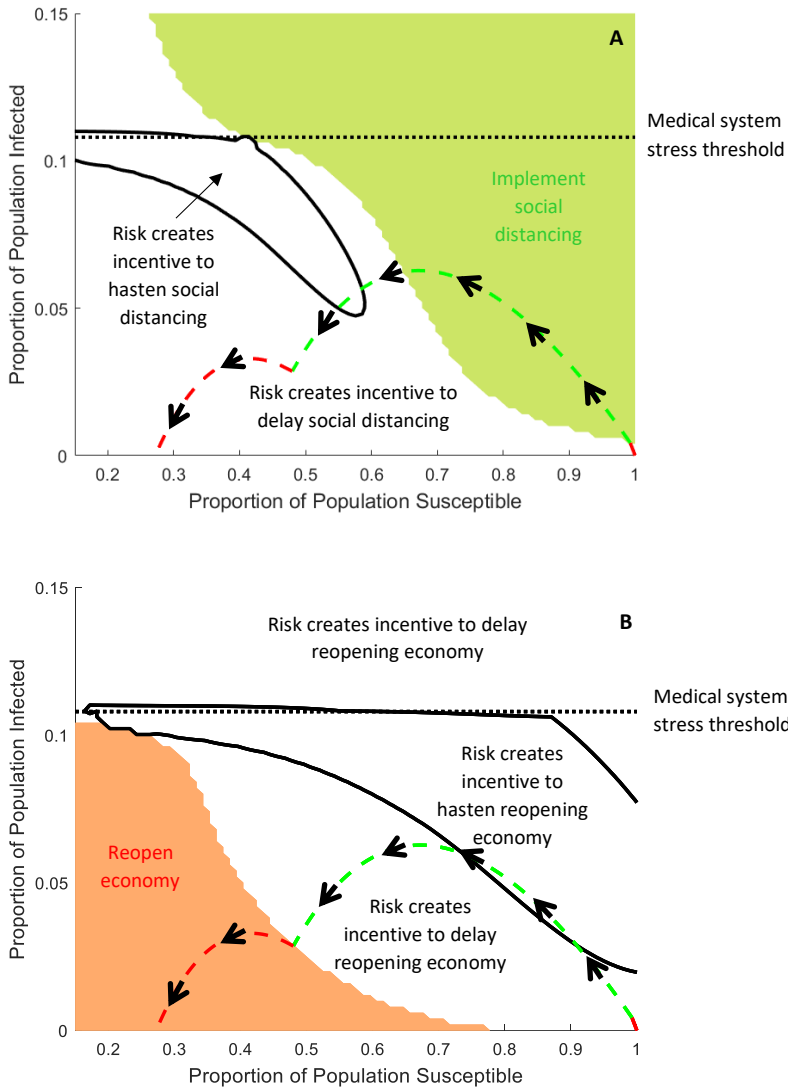


Figure A2. Risk adjustments for the decision to (A) close and (B) reopen the economy.

# Social learning along international migrant networks<sup>1</sup>

Yuan Tian,<sup>2</sup> Maria Esther Caballero<sup>3</sup> and Brian K. Kovak<sup>4</sup>

Date submitted: 25 October 2020; Date accepted: 25 October 2020

*We document the transmission of social distancing practices from the United States to Mexico along migrant networks during the early 2020 Covid-19 pandemic. Using data on pre-existing migrant connections between Mexican and U.S. locations and mobile-phone tracking data revealing social distancing behavior, we find larger declines in mobility in Mexican regions whose emigrants live in U.S. locations with stronger social distancing practices. We rule out confounding pre-trends and use a variety of controls and an instrumental variables strategy based on U.S. stay-at-home orders to rule out the potential influence of disease transmission and migrant sorting between similar locations. Given this evidence, we conclude that our findings represent the effect of information transmission between Mexican migrants living in the U.S. and residents of their home locations in Mexico. Our results demonstrate the importance of personal connections when policymakers seek to change fundamental social behaviors.*

- 1 We thank Edith Soto Ramirez at the Instituto de los Mexicanos en el Exterior for providing the MCAS data, UNACAST for sharing mobility data in the U.S., and Facebook's Data for Good initiative for sharing Facebook mobility data in the U.S. and in Mexico. Thanks to Brian Cadena, Michael Poyker, and participants in the 2020 NBER Summer Institute Development Economics meeting for helpful comments and to Treb Allen, Caué Dobbin, and Melanie Morten for helpful discussions regarding the MCAS data. All errors are our own.
- 2 Assistant Professor, School of Economics, University of Nottingham.
- 3 PhD Candidate, Heinz College, Carnegie Mellon University.
- 4 Associate Professor, Heinz College, Carnegie Mellon University.

Copyright: Yuan Tian, Maria Esther Caballero and Brian K. Kovak

## 1 Introduction

Social networks are a critical source of new information. By interacting with family, friends, and acquaintances, individuals learn new facts, observe the implications of others' decisions, and encounter new social norms. This type of social learning can be valuable when facing uncertainty about the nature of the choices one faces or the efficacy of one choice in comparison to others. Such uncertainty is particularly acute when one faces a novel set of choices and when the stakes are high. For example, in a time of pandemic when people must quickly learn about the nature of a disease and the appropriate actions to take in response, social learning can play an especially important role.

We document the transmission of social distancing practices from the United States to Mexico along migrant networks during the early 2020 Covid-19 pandemic. Social distancing is considered effective in reducing the spread of the novel coronavirus that causes Covid-19 and has been encouraged by public health organizations and most national and local governments.<sup>1</sup> The outbreak of Covid-19 in the United States emerged about two weeks earlier than in Mexico, and in the United States there was substantial spatial variation in timing of and compliance with social distancing policies. Using data on pre-existing migrant connections between Mexican and U.S. locations and mobile-phone tracking data revealing social distancing behavior, we find larger declines in mobility in Mexican regions whose emigrants live in U.S. locations with stronger social distancing practices. After ruling out confounding pre-trends and the potential influence of disease transmission and migrant sorting between similar locations (e.g. urban vs. rural areas), we conclude that our findings represent the effect of information transmission between Mexican migrants living in the U.S. and residents of their home locations in Mexico.

Key to our analysis is the ability to observe pre-existing migrant connections between Mexican source regions (*municipios*) and U.S. counties. We do so using administrative data from the *Matrícula Consular de Alta Seguridad* (MCAS) program, which provides identity cards to Mexicans living in the U.S. Prior work has confirmed the quality and representativeness of these data, which allow us to measure the extent to which each Mexican *municipio* was exposed to each U.S. county through the migrant network.<sup>2</sup> We combine these data with observed social distancing measures derived from smartphone geolocation data collected by Facebook and Unacast for the U.S. and Facebook for Mexico.<sup>3</sup> The Facebook data report the reduction in the number of 0.6 km-square

<sup>1</sup>Examples include the World Health Organization (WHO 2020), the U.S. Centers for Disease Control and Prevention (CDC 2020), and the Mexican Health Ministry (Secretaría de Salud 2020).

<sup>2</sup>Caballero et al. [2018] confirm the quality and representativeness of the MCAS data by comparing it against gold-standard household survey data. Other papers using data derived from the same underlying source include Albert and Monras [2019]; Allen et al. [2019]; and Caballero et al. [2020]. These papers and Caballero et al. [2018] each use slightly different extracts from the same underlying data source. As in Caballero [2020], we use the most detailed geographic information available (*municipio* by county) and use a version of the 2008 data that were cleaned and matched to valid *municipio* and county names by the *Instituto de los Mexicanos en el Exterior*.

<sup>3</sup>Similar data have been used to study migration (Blumenstock et al. 2019), segregation (Athey et al. 2019), commuting (Kreindler and Miyauchi 2019), friendship (Kreindler and Miyauchi 2019), and the spreading of disease (Kuchler et al. 2020), and are used in many of the papers cited below that focus on social distancing in response to Covid-19.

tiles visited each day, while Unacast reports the reduction in daily distance traveled. These data sources allow us to directly observe the behavior of interest (social distancing) and to do so with high frequency and fine geographic granularity. For each *municipio*, we calculate the migration-network-weighted average of social distancing across U.S. counties. Because social distancing varied substantially across U.S. counties, and migrants from different *municipios* go to very different sets of U.S. destinations, there is significant variation in exposure to U.S. social distancing across *municipios*.

Our empirical analysis examines how observed reductions in movement among people living in Mexico relate to this variation in migrants' exposure to U.S. social distancing. We find that *municipios* with a one-standard-deviation larger exposure to U.S. social distancing had a 0.47-standard-deviation larger decline in mobility. This finding is not driven by pre-existing trends and is robust to controlling flexibly for the number of local Covid-19 cases; the number of cases in migrant-connected U.S. counties; and baseline local characteristics including population density, urban status, age distribution, education, income, and the employment rate. The effect estimate also remains nearly unchanged when using U.S. state government stay-at-home orders as an instrument for U.S. social distancing behavior and when controlling for Mexican state government stay-at-home orders. When investigating heterogeneity in this effect, we find it is stronger in *municipios* with initially higher education levels, higher population density, and higher urbanization rate, but does not differ significantly with the characteristics of migrant-connected U.S. counties.

How should one interpret these results? There are several mechanisms that might generate an observed relationship between social distancing behavior in a Mexican *municipio* and in migrant-connected U.S. counties. First, migrants in the U.S. may observe the importance of social distancing during the U.S. outbreak, and may communicate that information back to people in their home region in Mexico, leading to more social distancing there as well. We refer to this as the "information" channel. Second, return migrants or others may have moved between the U.S. and Mexico, transmitting the disease and leading to correlated outbreaks in the two locations, which may in turn lead to correlated social distancing. We refer to this as the "disease transmission" channel. Third, migrants from locations with a higher likelihood of Covid-19 outbreak or with a higher likelihood of compliance with public health orders may choose similar locations in the U.S. If this is the case, then observed correlations between migrant-connected locations simply result from migrants' selection of destinations rather than reflecting a causal effect. We refer to this as the "migrant sorting" channel.<sup>4</sup>

Our empirical findings strongly reject the disease transmission and migrant sorting channels. The

<sup>4</sup>Another hypothetical channel would involve changes in remittances. If U.S. regions facing larger increases in social distancing also experience larger declines in economic activities, migrants living in those regions may reduce their remittance payments, leading to less economic activity and perhaps less mobility in their source regions in Mexico. While plausible in theory, this mechanism is unlikely to be relevant in our context. First, there was no substantial decline in remittances. In fact, according to remittance data collected by the Bank of Mexico, aggregate remittances in March 2020 surged, exceeding those of March 2019 by 35%, while remittances in April and May 2020 were within  $\pm 3\%$  of the values in the same months of 2019 (authors' calculations). Also, Mexican social distancing responds very quickly to declines in U.S. mobility, within one to two weeks. In contrast, the vast majority (68%) of Mexicans who send home remittances from the U.S. do so at monthly or longer frequencies, while only 15.3% send home remittances weekly (Serrano Herrera and Jiménez Uribe 2019).

observed relationship between U.S. and Mexican social distancing is barely affected when controlling flexibly for the number of cases in either location, implying that disease transmission is not driving our results. We address the possibility of migrant sorting first by controlling for pre-pandemic characteristics in the relevant *municipio*, including population density, urban status, age distribution, education, income, and the employment rate. As discussed below, these features are relevant for disease transmission and compliance with social distancing, but controlling flexibly for them has minimal effect on our results. We also use government stay-at-home orders as an instrument for U.S. social distancing behavior and again find nearly identical results. Together, these findings reject the disease transmission and migrant sorting channels, leaving the information channel as the most likely explanation for the observed relationship between U.S. and Mexican social distancing.

Our analysis relates to the large literature examining how social network connections reduce information frictions and facilitate learning. Papers in this literature cover a wide range of topics including technology adoption, labor markets, international trade, and many others.<sup>5</sup> A minority of these papers implements randomized controlled field trials, which include baseline network measures, randomized information interventions, and follow-up surveys measuring information transmission.<sup>6</sup> In contrast, the majority of this literature infers the presence of social learning based on equilibrium outcomes in the absence of a well-defined information shock. We contribute a clear example of social learning in an observational setting where we have a well-defined and credibly exogenous information shock, a high-quality measure of spatial network connections, and observed changes in behavior that are closely linked to the new information.

As in our setting, a number of papers in this broader literature focus on situations where immigrants transmit information across international borders. Examples include studies finding that immigrants increase trade with their source countries (Gould 1994, Head and Ries 1998, Rauch and Trinidad 2002), transfer knowledge through co-ethnic patent citations (Kerr 2008), influence source country political preferences (Barsbai et al. 2017, Karadja and Prawitz 2019) or fertility norms (Beine et al. 2013), and facilitate FDI and venture capital funding relationships with the source country (Dimmock et al. 2019, Kugler and Rapoport 2005, Li 2020, Pandya and Leblang 2017). We introduce a new example of cross-country information transmission through migrant networks, documenting migrants' role in spreading public-health information with potential life-and-death consequences. Moreover, we show that these responses can arise very quickly, with migrant source regions benefiting from destination-country information nearly in real time.

Our paper also contributes to the emerging literature examining the determinants of compliance with public health recommendations in the midst of Covid-19 outbreaks. Contemporaneous work shows that social distancing compliance varies with civic capital (Barrios et al. 2020), trust in

<sup>5</sup>BenYishay and Mobarak [2019] and Miller and Mobarak [2015] study information transmission in agricultural technology adoption; Barwick et al. [2019], Beaman [2012], Dustmann et al. [2016], Edin et al. [2003], and Munshi [2003] study the role of social networks and immigrant enclaves in job referrals and labor market outcomes; Büchel et al. [2019] examine how networks affect spatial mobility; Burchardi and Hassan [2013] show how social ties affect entrepreneurial activity and firm investment.

<sup>6</sup>Prominent examples include Beaman et al. [2018]; Banerjee et al. [2019]; BenYishay and Mobarak [2019]. See Breza et al. [2019] for a survey of the literature on networks in economic development.

science (Brzezinski et al. 2020), education and income (Brzezinski et al. 2020, Wright et al. 2020), partisanship (Allcott et al. 2020, Fan et al. 2020), media consumption (Ananyev et al. 2020, Simonov et al. 2020), political leaders' speech (Ajzenman et al. 2020), and whether workers can telework (Mongey et al. 2020). Additional work finds that many of these factors can impact the realized number of Covid-19 cases and resulting deaths (Bursztyjn et al. 2020, Desmet and Wacziarg 2020). Our work shows how migrants' experiences with U.S. Covid-19 outbreaks affect the social distancing behavior of those remaining in Mexico. This cross-country context is (to our knowledge) novel in this literature, and it helps avoid a number of potential pitfalls present in single-country designs.

For example, in a closely related paper Holtz et al. [2020] examine spillover effects of social distancing policies across U.S. counties, based on pre-existing mobility patterns and social-network friendship connections. Although we address similar questions, Holtz et al. [2020] face a much more challenging causal identification problem, because they examine spillovers between U.S. counties. It is quite likely that a U.S. county's choice of social-distancing policy is affected by those of neighboring counties, both for public health and political reasons, so reverse causality is a substantial concern. In our context, it is far less likely that U.S. social distancing practices or policies were influenced by Mexican practices or policies, mitigating concerns about reverse causality. The primary remaining threat to causal inference is the possibility of migrant sorting. As discussed above, we are able to allay these concerns using flexible controls for regional characteristics that may be relevant for sorting and an instrumental variables strategy. Thus, our setting provides a relatively clean test of the importance of social connections in driving compliance with public health recommendations during the pandemic.

The rest of the paper is organized as follows. Section 2 presents the institutional background on the Covid-19 epidemic and the U.S.-Mexico ties. Section 3 discusses the data on mobility and migrant networks. Section 4 shows the main empirical results on the effect of exposure to U.S. social distancing, and Section 5 investigates heterogeneous effects by origin and destination characteristics. The last section concludes.

## 2 Institutional background

### 2.1 The Covid-19 epidemic and the situations in the United States and in Mexico

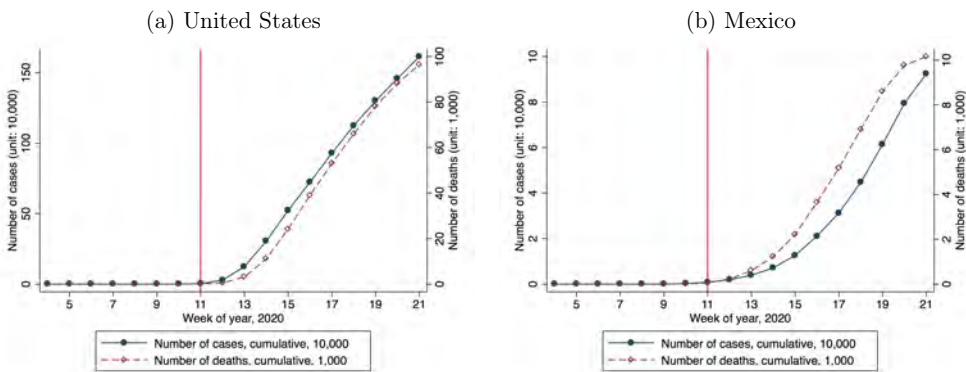
Covid-19 is a respiratory disease caused by a novel coronavirus (SARS-CoV-2). After the first case was reported in Wuhan, China on December 31, 2019, it spread across the world rapidly, despite containment efforts by various governments and organizations.<sup>7</sup> The World Health Organization (WHO) characterized Covid-19 as a pandemic on March 11, 2020, and by June 12, 2020, there were 7,533,182 cases, 423,349 confirmed deaths, and 216 countries, areas, or territories with cases worldwide.<sup>8</sup>

<sup>7</sup>See the detailed WHO timeline at <https://www.who.int/news-room/detail/27-04-2020-who-timeline---covid-19>.

<sup>8</sup>The following declaration was accessed on June 13, 2020: <https://www.who.int/emergencies/diseases/novel-coronavirus-2019/events-as-they-happen>.

The epicenter of the outbreak has been shifting over time. After China's initial outbreak and lockdown measures in January and February, the epicenter shifted to Europe in mid March, followed by the United States starting from late March, and by June, further shifted to Latin American countries. In the United States, the first case was reported on January 22, 2020, and President Trump declared a national emergency on March 13 (in the 11th week of 2020, shown in Figure 1 with a red vertical line).<sup>9</sup> As of June 12, 2020, the total number of U.S. cases was 2,016,027 and the number of deaths was 113,914.<sup>10</sup> Figure 1 Panel (a) shows the number of cases (solid circles) and number of deaths (hollow diamonds) from Week 4 of 2020 (Jan 20–26) to Week 21 (May 18–24). The numbers of U.S. cases and deaths began increasing rapidly after Week 13. The outbreak in Mexico emerged slightly later. The first case was confirmed on February 28, 2020, and the increase in the number of cases and number of deaths accelerated after Week 15 (Figure 1 Panel b). By the end of Week 22, there were 47.7 cases and 2.9 deaths per 10,000 U.S. population, and there were 8.2 cases and 0.9 deaths per 10,000 Mexican population.<sup>11</sup>

Figure 1: U.S. outbreak began earlier and was more severe than that of Mexico by Week 21.



Note: The number of cases and deaths in the United States are from Johns Hopkins University: <https://coronavirus.jhu.edu/>. The corresponding information in Mexico is from the Mexican Ministry of Health: <https://coronavirus.gob.mx/>. The horizontal axis represents the week of the year in 2020. For example, Week 4 is the Week of Jan 20 to Jan 26, and Week 21 is the week of May 18 to May 24. The vertical line at Week 11 denotes the week when a national emergency was declared in the United States.

The Covid-19 outbreak was unexpected, and in many ways unprecedented, meaning that governments and public health organizations had much to learn regarding how to appropriately respond.<sup>12</sup>

<sup>9</sup>Declaration of a national emergency: <https://www.whitehouse.gov/presidential-actions/proclamation-declaring-national-emergency-concerning-novel-coronavirus-disease-covid-19-outbreak/>.

<sup>10</sup>The following source was accessed on June 13, 2020: <https://www.cdc.gov/coronavirus/2019-ncov/cases-updates/cases-in-us.html>.

<sup>11</sup>Note that observed cases and deaths are subject to testing capacity and reporting errors. In the case of Mexico, for example, there are concerns about the hidden death toll: <https://www.nytimes.com/2020/05/08/world/americas/mexico-coronavirus-count.html>.

<sup>12</sup>In WHO's announcement of the pandemic, the WHO Director-General said that "we have never before seen a pandemic sparked by a coronavirus. This is the first pandemic caused by a coronavirus. And we have never before seen a pandemic that can be controlled, at the same time." (<https://www.who.int/dg/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19---11-march-2020>)



As an example, Italy declared a state of emergency on Jan 31, 2020 and subsequently halted air traffic to and from China.<sup>13</sup> However, the disease continued to spread, and a national lockdown was imposed on March 9, 2020, when Italy became the epicenter of the pandemic. Strict travel restrictions were in place, only essential businesses were allowed to open, and people were required to maintain at least one meter of distance from one another in public spaces.<sup>14</sup> In the case of the U.S., although international travel restrictions with China were in place relatively early, the effectiveness of this and other policies has been debated. After one week of the outbreak in the State of Washington, the White House issued social-distancing guidelines on March 16; recommendations regarding the use of cloth face coverings were issued by the Centers for Disease Control and Prevention (CDC) on April 3.<sup>15</sup>

Individuals also report learning from the experiences of others in their social networks. In Prato, Italy, where a quarter of the population is ethnic Chinese, residents voluntarily quarantined and practiced social distancing much earlier than those in the rest of the country, after learning about the success of similar measures in China, leading to very low rates of infection and transmission.<sup>16</sup> Similarly, restaurants owned by Chinese immigrants in the U.S. began scaling up takeout and delivery operations prior to the U.S. outbreak, based on information from similar businesses in China.<sup>17</sup> Holtz et al. [2020] find that social distancing in U.S. regions significantly influenced policies and behaviors in other parts of the country.

## 2.2 Mexico-U.S. migration

The U.S. and Mexico have long been closely linked in terms of trade and migration. The U.S. is Mexico's most important trading partner, accounting for 76% of Mexican exports in 2018, and 96% of those who reported living abroad five years prior to the 2010 Mexican Census.<sup>18</sup> Mexican migrants in the U.S. maintain close ties with their friends and family in Mexico. According to data from the Mexican Migration Project (MMP), an average Mexican migrant sends 27% of income earned in the U.S. back to Mexico, a much higher share than saving (20%), food budget (19%), or rent (18%).<sup>19</sup> During their first trip to the U.S., 61% of migrants received financial help from people in their home community. Such close ties do not deteriorate much along repeated migration trips; even in their

<sup>13</sup><https://www.reuters.com/article/china-health-italy/italy-govt-agrees-state-of-emergency-after-confirmed-coronavirus-cases-govt-source-idUSR1N282044>

<sup>14</sup>See details of the timeline and measures at: [https://en.wikipedia.org/wiki/COVID-19\\_pandemic\\_lockdown\\_in\\_Italy](https://en.wikipedia.org/wiki/COVID-19_pandemic_lockdown_in_Italy).

<sup>15</sup>The details and timeline of the Washington outbreak: [https://en.wikipedia.org/wiki/COVID-19\\_pandemic\\_in\\_Washington\\_\(state\)#March\\_1%E2%80%931935](https://en.wikipedia.org/wiki/COVID-19_pandemic_in_Washington_(state)#March_1%E2%80%931935). Social distancing guidelines: <https://www.whitehouse.gov/articles/president-trump-actions-to-confront-pandemic/>. Face covering recommendations: <https://www.cdc.gov/coronavirus/2019-ncov/prevent-getting-sick/cloth-face-cover.html>.

<sup>16</sup>Source: <https://www.reuters.com/article/us-health-coronavirus-italy-chinese/from-zero-to-hero-italys-chinese-help-beat-coronavirus-idUSKBN21I3I8>

<sup>17</sup><https://www.npr.org/2020/05/13/855791740/episode-999-the-restaurant-from-the-future>


<sup>18</sup>Sources: trade data: <https://wits.worldbank.org/countrysnapshot/en/MEX>; 2010 Mexican Census: IPUMS International (Minnesota Population Center 2020).

<sup>19</sup>The Mexican Migration Project is a collaborative research project based at Princeton University and the University of Guadalajara. The data are publicly available at: <https://mmp.opr.princeton.edu/>. The figures here were calculated using a sample of 8,823 individuals who had a previous trip to the U.S., using survey years from 1982 to 2018.

last trip to the U.S., 51% received financial help. It is therefore entirely plausible that information regarding pandemic response would be transmitted from U.S. migrants to contacts in their home communities.

During a pandemic, the intensive flow of goods and people between the U.S. and Mexico can transmit both disease and information.<sup>20</sup> However, due to the travel restrictions imposed early in the pandemic, the number of trips across the U.S.-Mexico border fell substantially, as shown in Figure 2, which reports the number of trips between the two countries as recorded among Facebook mobile app users. Initially there were more than 134,000 trips per day from Mexico to the U.S., and more than 137,000 trips from the U.S. to Mexico, but the numbers declined sharply after Week 11 when the U.S. declared a national emergency and imposed more strict travel restrictions. By Week 15, the number of trips declined to 40,000 per day on both sides, with a slight increase afterwards. Although cross-border flows have fallen by about two-thirds since early March, many people still cross the border each day. Our empirical analysis will therefore address the possibility of physical disease transmission along with potential information flows through migrant networks.

Figure 2: Number of trips between the U.S. and Mexico declined after Week 11



Note: The trip counts are calculated using Facebook mobile app users who opt into location services. The horizontal axis is the week of the year, and the vertical axis is the average number of trips per day during the week. The vertical line at Week 11 denotes the week when the national emergency was declared in the United States.

<sup>20</sup>For example, Kuchler et al. [2020] use Facebook friendship data between U.S. counties to show that the outbreak followed these connections.

### 3 Data and measurement

#### 3.1 Migrant network between Mexico and the U.S.

We use administrative information from the *Matrícula Consular de Alta Seguridad* (MCAS) program to measure migration networks between the U.S. and Mexico at the sub-national level. The MCAS card, which acts as an official form of identification for banking purposes and other transactions, is issued by Mexican consulates to Mexican citizens living in the U.S.<sup>21</sup> The MCAS administrative dataset contains annual counts of newly issued MCAS cards by place of birth in Mexico and place of residence in the U.S.

Caballero et al. [2018] validate the migration network measures obtained from the MCAS data by showing that although they likely over-represent unauthorized Mexican-born migrants, who have the strongest incentive to obtain a *matrícula*, they have strong agreement with the source and destination distributions of Mexican-born migrants obtained from high quality household surveys both from Mexico and the U.S. In this paper, we construct the migration network measure as the share of *matrículas* issued in 2008 to Mexican-born migrants from each Mexican *municipio* living in each U.S. county. Summary statistics appear in Table 1. There are 174,281 *municipio*-county pairs, with 2,412 origin *municipios* and 2,468 destination U.S. counties in the 2008 MCAS dataset. The average number of migrants per link is 5.5, but it varies substantially, ranging from 1 to 5,253.

Table 1: Summary of statistics for Mexico-U.S. migration networks using the 2008 MCAS data

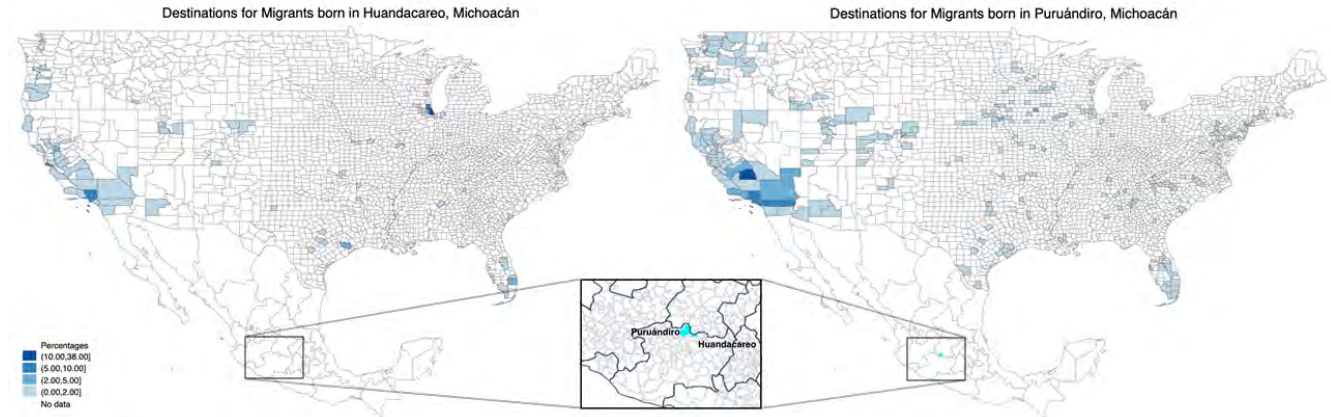
Variable	Value
Number of Mexican <i>municipios</i>	2,412
Number of U.S. counties	2,468
Number of county- <i>municipio</i> pairs (links)	174,281
Mean (s.d.) # of migrants per link	5.5 (28)
Min (max) # of migrants per link	1 (5,253)

Note: 2008 MCAS data. A link is a *municipio*-county pair, and the number of migrants per link is the number of Mexicans from the origin *municipio* who reside in the corresponding destination county in the U.S.

Our empirical analysis relies on the fact that migrants from different *municipios* choose quite different destinations in the U.S. and therefore are exposed to different social distancing practices in different parts of the country. Figure 3 shows the destination distribution for two different *municipios* in the state of Michoacán: Huandacareo and Puruándiro. Despite these two sources being located very close to each other (less than an hour apart by car) and thus roughly equal distances from particular U.S. labor markets, there are large differences in the U.S. destinations selected by migrants from these two *municipios*. The vast majority of migrants from Huandacareo live in Chicago (Cook County), while the most common destination for migrants from Puruándiro is Tulare county in California's Central valley. Because social distancing behavior differed across these U.S. destinations (shown in Figure 4 below), migrants from Huandacareo and Puruándiro will be exposed to different degrees of social distancing in the U.S. This example is representative in

<sup>21</sup>See Caballero et al. [2018] for more detail on the MCAS program and data.

Figure 3: Differences in migrant destination distributions



Note: 2008 MCAS data. The left panel shows the distribution across U.S. counties for migrants from Huandacareo, and the right panel shows the same for migrants from Puruándiro.

the sense that migrants from otherwise similar *municipios* often exhibit quite different destination distributions in the U.S. (Caballero et al. 2018), leading to variation in exposure to U.S. social distancing across *municipios*.

### 3.2 Unacast and Facebook data on local mobility

We use two data sources to measure changes in mobility. Due to the nature of Covid-19 transmission, scientists have identified social distancing as one of the key measures to combat the pandemic (Hsiang et al. 2020 and Anderson et al. 2020).<sup>22</sup> One way to measure the extent of social distancing behaviors is to use the reduction in geographic mobility. Our first mobility measure is from Unacast, a New York based technology company (Unacast 2020). The dataset uses location information from 15-17 million smartphones to calculate the average distance travelled each day. We measure the county-level mobility reduction as the percentage reduction in the average distance traveled compared to the same day of the week during the four weeks before March 8, 2020 (prior to the outbreak). As shown in Table 2, the measure covers 3,054 counties in the U.S., with an average decline in mobility of 19% during the period of Week 9 to Week 21.<sup>23</sup>

Table 2: Mobility data summary statistics

Source	Country	Moment	Value
Unacast	U.S.	# of counties	3,054
		Mean (s.d.) of decline in mobility, Week 9–21	-0.19 (0.17)
Facebook	U.S.	# of counties	2,691
		Mean (s.d.) of decline in mobility, Week 9–21	-0.13 (0.14)
	Mexico	# of <i>municipios</i>	1,084
		with exposure to US measure	1,014
	Mean (s.d.) of decline in mobility, Week 9–21	-0.21 (0.15)	

Sources: Unacast and Facebook Data for Good. Unacast data covers 3,054 U.S. counties, while the coverage of the Facebook data varies by week (see Appendix Table 10 for details).

Our second mobility measure is from Facebook’s Data for Good program.<sup>24</sup> The dataset uses the location information of users who enable location services on their mobile Facebook app. The mobility metric is the proportional change in the average number of 0.6 km by 0.6 km tiles visited during a 24 hour period compared to same day of the week in February 2020 (excluding President’s day).<sup>25</sup> The data cover 2,691 counties in the U.S. and 1,084 *municipios* in Mexico, since only regions

<sup>22</sup>See CDC recommendation at: <https://www.cdc.gov/coronavirus/2019-ncov/prevent-getting-sick/social-distancing.html>

<sup>23</sup>There are two filters applied to the sample to ensure data reliability. Unacast define a “dwell” as a set of location records observed within 80 meters of each other within an 8-minute to 4-hour time period. Only devices with at least two dwells per day or one dwell longer than three hours in duration are included in the analysis. The data also exclude counties with population less than 1,000 or that did not have at least 100 devices on at least 70% of the days during the pre-outbreak period.

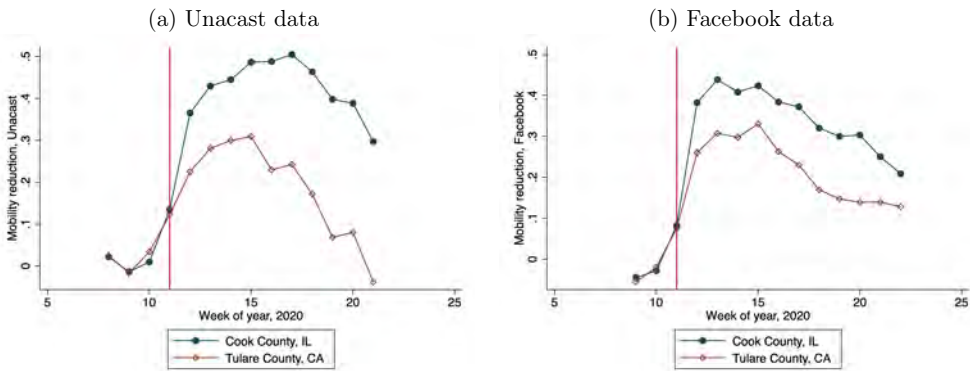
<sup>24</sup>Source: <https://dataforgood.fb.com/docs/covid19/>.

<sup>25</sup>Details of the tile system are available at: <https://docs.microsoft.com/en-us/bingmaps/articles/bing-maps-tile-system>.

with more than 300 unique users are included. During the period of Week 9 to Week 21, the average decline in mobility in the U.S. is 13%, and the decline in Mexico is 21%. (Table 2)

Places in the U.S. vary in the extent of social distancing. We measure social distancing based on the observed mobility reduction, with more positive values corresponding to larger declines in mobility. Figure 4 uses Cook County in Illinois (solid circles) and Tulare County in California (hollow diamonds) as an example. The reduction in mobility is more pronounced and persistent in Cook County than in Tulare County. In the Unacast data (Panel a) both counties started around zero in Week 10, and by Week 12, the decline in mobility was 37% in Cook County and 23% in Tulare County. In Week 21, Cook County’s mobility reduction declined to 30%, while in Tulare County it fell just below zero, indicating no reduction in mobility compared to the pre-pandemic period. Although the differences are less extreme in the Facebook data (Panel b), mobility in Cook County clearly decreased far more than in Tulare County in each week. Appendix Figure 10 maps the increase in social distancing from Week 9 to Week 21 for all U.S. counties in the Unacast and Facebook datasets, documenting substantial variation in social distancing behavior across counties.

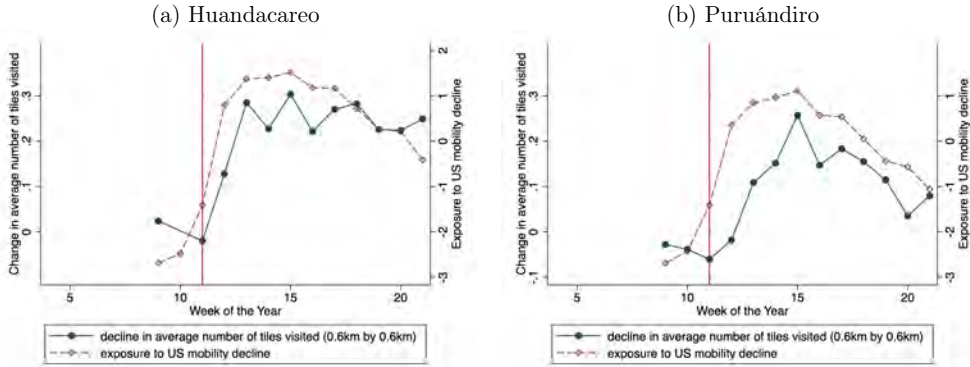
Figure 4: Larger and more persistent mobility reduction in Cook County than in Tulare County, Week 9–21.



Note: The solid vertical line is at Week 11, when a national emergency was announced in the U.S.

Our Facebook data also cover mobility in Mexican *municipios*. Figure 5 shows the trends in social distancing (solid circles, left y-axis) in Huandacareo (Panel a) and Puruándiro (Panel b), the two *municipios* in Michoacán considered in Figure 3. Social distancing in Mexico lagged that of the U.S. by a few weeks, consistent with the somewhat later emergence of the pandemic in Mexico. Both Huandacareo and Puruándiro exhibit substantial reductions in mobility by Weeks 13 and 14, although Huandacareo exhibits more social distancing than Puruándiro. As we will discuss below, this is consistent with the fact that migrants from Huandacareo, who tend to migrate to Chicago, were exposed to more social distancing in the U.S. than migrants from Puruándiro, who tend to migrate to California’s Central Valley.

Figure 5: Trends in social distancing in Huandacareo and Puruándiro, Mexico and their exposure to U.S. mobility declines



Note: The solid vertical line is at Week 11, when a national emergency was announced in the U.S.

### 3.3 Mexican exposure to U.S. social distancing


Intuitively, if a *municipio* happened to have more migrants residing in a U.S. county where more social distancing measures were taken, the migrants’ relatives and friends remaining in that *municipio* may have received more information about the severity of Covid-19 and the importance of social distancing, and may have further transmitted this information to other residents of the *municipio*. Thus, we measure a Mexican *municipio*’s exposure to U.S. social distancing practices as follows:

$$exposure_{it}^s = \sum_j \frac{m_{ij}}{\sum_{j'} m_{ij'}} socdist_{jt}^s, \tag{1}$$

where  $m_{ij}$  is number of MCAS cards issued to migrants from *municipio*  $i$  living in county  $j$  in 2008,  $socdist_{jt}^s$  is the social distancing measure in county  $j$  week  $t$  using data sources  $s \in \{\text{Facebook, Unacast}\}$ . In our main analysis, we reduce noise by using the principal component of the two social distancing measures, denoted as  $exposure_{it}^{pc}$ , but our results are robust to using either data source individually.

Figure 6 maps the change in the exposure measure in (1) for each Mexican *municipio* from Week 9 to Week 21, using the Unacast data. This exposure measure ranges from  $-0.02$  to  $0.47$ , and the mean is  $0.2$ , indicating that migrants lived in U.S. counties with a 20 percentage-point average decline in mobility from Week 9 to Week 21. Variation in exposure derives from a combination of the variation in social distancing across U.S. counties, shown in Figure 4 and Appendix Figure 10, and the variation in the migrant destination distribution across *municipios*, shown in Figure 3. These two sources of variation lead to significant differences in exposure to U.S. social distancing across *municipios*, and our empirical analysis will examine how this exposure influenced social distancing behavior in Mexico.

Figure 6: Change in exposure to U.S. social distancing, Week 9 to Week 21



Note: The change in exposure to U.S. social distancing is calculated as  $exposure_{21}^{Unacast} - exposure_9^{Unacast}$ , using the Unacast data. See Appendix Figure 9 for versions using Facebook or the principal component of the Unacast and Facebook measures together.

Returning to Figure 5, we plot the exposure measure (hollow diamonds, right y-axis) for Huandacareo (Panel a) and Puruándiro (Panel b). Because Huandacareo’s migrants concentrate in Cook County (Chicago), which had a large increase in social distancing, and Puruándiro’s migrants concentrate in Tulare County (CA Central Valley), which had much less social distancing, Huandacareo was exposed to a larger U.S. mobility decline throughout the pandemic period. A *municipio*’s mix of migrant destinations combines with variation in social distancing behavior across U.S. counties to create variation in exposure to U.S. social distancing, as measured in (1). After being more exposed to more U.S. social distancing through its migrant network, Huandacareo exhibited larger declines in mobility than Puruándiro. As we will document below, this relationship between U.S. and Mexican social distancing holds on average across *municipios*.

### 3.4 Other datasets

We use various data sources to measure characteristics of U.S. counties and Mexican *municipios* that may have affected disease or information transmission. The number of weekly Covid-19 cases and deaths by U.S. county come from John Hopkins University, and the corresponding information for Mexico come from the Mexican Ministry of Health.<sup>26</sup> Paralleling our measure of exposure to U.S. social distancing, we also construct *municipio i*’s exposure to U.S. Covid-19 cases as follows.

<sup>26</sup>Sources: <https://coronavirus.jhu.edu/> and <https://coronavirus.gob.mx/>, respectively.



$$exposure_{it}^{case} = \sum_j \frac{m_{ij}}{\sum_{j'} m_{ij'}} \sinh^{-1}(\text{cumulative cases}_{jt}), \quad (2)$$

where  $\sinh^{-1}(\text{cumulative cases}_{jt})$  is the inverse-hyperbolic-sine transformation of the cumulative number of cases in county  $j$  and week  $t$ . We use the inverse hyperbolic sine transformation to include the counties with zero cases, and in the remaining text, we use “log cumulative cases” or “log new cases” as a shorthand, given the close correspondence between the natural log and inverse hyperbolic sine, particularly for large numbers.

U.S. county characteristics are from the 2010 Census and 2005–2009 American Community Surveys (ACS). Specifically, we use the 2010 Census to calculate the county-level Hispanic or Latino share of the population, the Mexican population share, total population, and information on the land area (used to calculate population density). We use the 2005–2009 ACS to calculate the county-level (1) number of Hispanic and Latino individuals aged 25 and over by educational attainment, and similar numbers for the overall population; (2) number of Hispanic and Latino families by income group; (3) mean and median household income in the entire population; (4) number of employed persons by industry (NAICS); and (5) mode of transportation to work.

Mexican *municipio* characteristics are from the 2015 Intercensal Count (*Coneto*), including the share of working age population (aged 16 to 65), schooling attainment, share employed, and income earned in the working age population.<sup>27</sup> We obtain population density and percent of urban population from Mexican Statistical Office (INEGI) tabulations.

The timing of issuing and lifting stay-at-home orders by U.S. states are obtained from Raifman et al. [2020], and similar data for Mexican states were collected from Mexican states’ official decrees (see Appendix B).

## 4 Social learning across borders: main empirical results

### 4.1 Empirical specification

Our empirical analysis examines the impact of exposure to U.S. social distancing practices on social distancing in Mexican *municipios*, and seeks to isolate the portion of that impact driven by social learning. Our baseline estimation equation is as follows:

$$socdist_{it} = \alpha + \beta exposure_{it}^{pc} + \Gamma z_{it} + I_i + I_t + \epsilon_{it} \quad (3)$$

where  $socdist_{it}$  is the mobility reduction in *municipio*  $i$  week  $t$  using the Facebook data, and  $exposure_{it}$  measures exposure to U.S. social distancing in the same week. We include a variety of *municipio*-week-specific controls  $z_{it}$  to take into account time-varying region-specific factors that could affect people’s social distancing behavior, such as the severity of the local disease outbreak. *Municipio* fixed effects  $I_i$  are included to control for *municipio*-specific factors such as population

<sup>27</sup>Source: IPUMS International [Minnesota Population Center, 2020].

density, income level, education level, and means of transportation to work. Week fixed effects  $I_t$  are used to account for national policies that affect social distancing behaviors across all regions in a week. We present robust standard errors, but clustering at the *municipio* level gives very similar results.

The parameter  $\beta$  captures the relationship between U.S. social distancing behaviors and network-connected Mexican *municipios*' social distancing practices. A positive value of  $\beta$  indicates that *municipios* connected to U.S. counties practicing more social distancing experienced on average larger reductions in mobility. In order to interpret  $\beta$  as the causal effect of U.S. social distancing on Mexican social distancing, the key identification assumption is that changes in social distancing behaviors across Mexican *municipios* with similar observable characteristics would not have differed systematically in the absence of differential exposure to U.S. social distancing practices.

This identification assumption may be violated if, for example, Mexican regions with higher population density tend to send more migrants to U.S. counties with higher population density. Since the probability of infection is higher in denser areas, people in both regions may practice more social distancing even in the absence of information transmission. A similar issue may arise if migrant origins and destinations are selected along other dimensions that affect the severity of the local Covid-19 outbreak. We refer to this as the “migrant sorting” channel and rule out its effects on our estimates by including extensive controls flexibly capturing the effects of relevant regional characteristics over time. Another threat to causal interpretation would arise if more exposed *municipios* had different trends in mobility even before the outbreak. However, as seen in Figures 4 and 5, both U.S. and Mexican mobility reductions were very close to zero in Week 9 compared to the pre-Covid period. This is not particular to the regions examined in those figures; the mean mobility reduction across *municipios* was 0.01 in Week 9 (standard deviation 0.05). Thus, pre-Covid social distancing trends were nearly identical and approximately equal to zero across *municipios*.

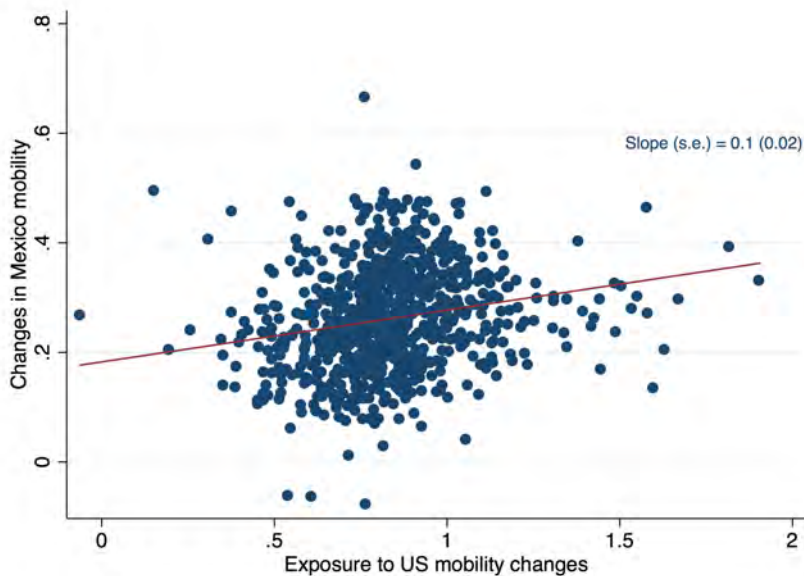
A final concern is that a positive estimate of  $\beta$  may reflect the effects of disease transmission rather than information transfer along migrant networks. If disease transmission operates along the migration network, then migrant-connected locations in the U.S. and Mexico will have similar severity and timing of outbreaks and may have similar degrees of social distancing as a result. Although a result driven by this “disease transmission” channel could potentially be interpreted as a causal effect of exposure to the U.S. outbreak, it would not reflect the information channel of interest. In order to rule out this disease transmission channel, we include flexible controls for the severity of the local outbreak in the *municipio* and in network-connected U.S. counties (following (2)). In addition, we address a variety of these causal inference concerns using government stay-at-home order as an instrument for U.S. social distancing behavior.<sup>28</sup>

<sup>28</sup>See footnote 4 for a discussion ruling out the role of remittances in driving correlated social distancing between migrant-connected regions in Mexico and the U.S.

## 4.2 Main results

Before reviewing the main estimation results from Equation (3), we present visual evidence on the relationship between Mexican social distancing and declines in mobility in migrant-connected U.S. counties. For each *municipio*  $i$ , we calculate the long-difference change in local social distancing ( $socdist_{i21} - socdist_{i11}$ ) and the change in the *municipio*'s exposure to U.S. social distancing ( $exposure_{i21}^{pc} - exposure_{i11}^{pc}$ ) from Week 11 to Week 21. Figure 7 shows a scatter plot relating these two measures, where each point represents a *municipio*. The fitted line has a slope of 0.1 (significant at the 1% level), indicating that a one-standard-deviation larger exposure to U.S. social distancing is associated with a 0.2-standard-deviation larger decrease in Mexican mobility.

Figure 7: Positive correlation between changes in social distancing in Mexico and the U.S. (Week 11 to Week 21)



Note: This figure includes all Mexican *municipios* with at least one Covid-19 case in Week 21, and each point represents a *municipio*. The horizontal axis is the exposure to U.S. social distancing in Week 21 minus that in Week 11 ( $exposure_{i21}^{pc} - exposure_{i11}^{pc}$ ), and the vertical axis is the change in social distancing in a Mexican *municipio* between Week 21 and Week 11 ( $socdist_{i21} - socdist_{i11}$ ). The mean (s.d.) of the x-axis is 0.8 (0.2), and the mean (s.d.) of the y-axis is 0.3 (0.1).

We now turn to the main specification in equation (3) to investigate cross-border learning about social distancing. Table 3 shows the estimation results. Columns (1)–(4) only include *municipios* with at least one case by Week 21 and Columns (5)–(8) include all *municipios*.<sup>29</sup> In Column (1), we regress social distancing in Mexican *municipios* on the exposure to U.S. social distancing ( $exposure_{it}^{pc}$ ), controlling for *municipio* fixed effects and week fixed effects. The coefficient is 0.05

<sup>29</sup>For details on the sample restrictions, see Appendix 11.

(statistically significant at the 1% level), indicating that a one-standard-deviation larger exposure to U.S. social distancing (1.4) led to a 0.47-standard deviation larger increase in social distancing in Mexico. Column (5) uses the same specification for all *municipios*. The coefficient is 0.03, smaller than the Column (1) estimate, suggesting that the learning effect might be weaker in areas with no active Covid-19 cases.<sup>30</sup>

Table 3: Larger exposure to U.S. social distancing led to more social distancing in Mexico, Week 9 to Week 21

Outcome:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mexico social dist.	<i>Municipios</i> with cases>0				All <i>municipios</i>			
Exposure to U.S. social dist.	0.05*** (0.005)	0.05*** (0.005)	0.04*** (0.005)	0.05*** (0.005)	0.03*** (0.004)	0.03*** (0.004)	0.03*** (0.004)	0.03*** (0.004)
Exposure to log U.S. cum. cases		-0.01*** (0.003)		-0.01*** (0.003)		-0.01*** (0.002)		-0.01*** (0.002)
Log cum. cases in Mexico muni.			0.02*** (0.001)	0.02*** (0.001)			0.01*** (0.001)	0.02*** (0.001)
Constant	0.22*** (0.000)	0.28*** (0.015)	0.19*** (0.001)	0.28*** (0.015)	0.21*** (0.000)	0.26*** (0.013)	0.19*** (0.001)	0.26*** (0.013)
Observations	10,051	10,051	10,051	10,051	13,036	13,036	13,036	13,036
R-squared	0.91	0.91	0.92	0.92	0.91	0.91	0.91	0.91

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Week fixed effects and *municipio* fixed effects are controlled in all columns. Columns (1)–(4) include the *municipios* with at least one Covid-19 case at the end of Week 21, and Columns (5)–(8) include all *municipios*. The mean (s.d.) of Mexican social distancing in the first four columns is 0.21 (0.15), and the mean (s.d.) of the exposure to U.S. social distancing is -0.02 (1.4). The mean (s.d.) of the log cumulative cases in Mexico is 1.4 (1.8), and the mean (s.d.) of the exposure to U.S. cases is 5.1 (2.7). The corresponding numbers for the last four columns are: 0.21 (0.15), -0.1 (1.4), 1.1 (1.7), and 5.1 (2.7).

In Columns (2)–(5) and (6)–(8), we introduce controls for the cumulative numbers of cases in the relevant *municipio* or in the U.S. destinations to which it is connected via the migrant network, using the measure in equation (2). As expected, when the local outbreak is more severe, people in Mexico practice more social distancing.<sup>31</sup> In contrast, the U.S. case-exposure variable consistently has a negative coefficient, suggesting that observing more cases in U.S. destination regions actually decreased people’s incentive to practice social distancing in Mexico. One potential explanation is that, conditional on the realized level of U.S. social distancing, an increase in the number of U.S. cases sends the signal that social distancing is not very effective in stopping the spread of the disease.<sup>32</sup> That said, the most important conclusion for this portion of the analysis is that the

<sup>30</sup>These numbers are smaller than the slope coefficient in Figure 7, since the previous estimate only uses data from Week 11 and Week 21. It is evident from Figure 4 that the U.S. social distancing was the strongest around Week 15 and declined afterwards due to reopening. In the meantime, Mexican social distancing had not declined by the end of the study period (Week 21). Thus, when we use the full panel of Week 9 to Week 21 instead of restricting to the starting and the ending weeks alone, the changes in U.S. social distancing are larger, and as a result, the coefficient estimate on exposure is smaller.

<sup>31</sup>Similarly, Brzezinski et al. [2020] find that in the United States, people engage in social distancing even in the absence of lockdown policies, once the virus occurs in their area.

<sup>32</sup>Briscese et al. [2020] present a related finding, showing that Italian residents were less likely to follow self-isolation

estimated effect of exposure to U.S. social distancing is essentially unchanged when including these controls for the number of cases. This finding rules out the disease transmission channel discussed in the prior subsection. If the observed correlation between U.S. and Mexican social distancing were the result of disease transmission along the migrant network, the inclusion of these controls would absorb the variation driving the observed correlation, and our results would disappear. In Appendix Table 17, we further reinforce this conclusion by controlling for flexible functional forms of the number of cases, with nearly identical results.

Although disease transmission is not an important mechanism driving the relationship between U.S. and Mexican social distancing behavior, it remains possible that correlated social distancing behavior results from underlying similarities in migrants' source and destination regions – what we have called the “migrant sorting” channel. For example, concurrent research finds that individuals and regions with higher education levels and higher incomes are more likely to practice social distancing (Brzezinski et al. 2020, Wright et al. 2020, Mongey et al. 2020, and Fan et al. 2020, among others). If migrants from higher income areas of Mexico are more likely to choose higher income destinations in the U.S., then one might observe correlated social distancing even without information transmission. Although our inclusion of *municipio* fixed effects addresses level differences in social distancing, it does not capture the likelihood that higher income locations (for example) *increasingly* practicing social distancing as the pandemic evolves.

We address this migrant sorting concern in Table 4. First, we measure *municipio* features that the literature has shown are correlated with baseline social distancing behavior, including population density, urban share, working-age share, average years of education, mean log income, and the employment to population rate. Then, we control for each feature interacted with separate indicators for each week of our sample, allowing for the effect of the relevant feature to vary arbitrarily over time. As an example, Column (1) interacts the initial population density with week indicators, controlling for the possibility that migrants from more densely populated *municipios* choose to live in more densely populated counties. Across the columns of Table 4, it is apparent that the six regional features generally drive larger gaps in Mexican social distancing between Week 9 and Week 12, after which the effects are largely stable. Most importantly for our purposes, the effect of exposure to U.S. social distancing is very stable when comparing the estimates in Table 4 to those in Columns (1)–(4) of Table 3.<sup>33</sup> This rules out the effects of migrant sorting based on the characteristics investigated in Table 4.

Another potentially important regional characteristic is the share of jobs that facilitate working from home. In Appendix Section C.6, we use the industry-level measure of the ability to work from home constructed by Dingel and Neiman [2020], and the industry mix of local employment in the 2015 Intercensal Count to construct the share of jobs in each Mexican *municipio* that facilitate working from home. We repeat the analysis in Table 4 by using the interaction of this share with policies when the policies are kept in place longer than expected. In our context, Mexican residents may also be discouraged by the observation that U.S. cases kept increasing despite the social distancing policies, and were less likely to follow suit.

<sup>33</sup>Table 4 uses only *municipios* with a positive case count, but we find similar agreement when using all *municipios* - see Appendix Table 22.

week fixed effects, finding that the effect of exposure to U.S. social distancing remain unchanged. Thus, as with the characteristics examined in 4, this regional characteristic is not likely to drive the relationship between U.S. and Mexican social distancing through migrant sorting.

Table 4: The main results are robust to controlling for differential effects of socio-economic conditions across weeks

Variable for interaction	(1) pop. density	(2) % urban	(3) % age 16–65	(4) years edu.	(5) log income	(6) % employed
Exposure to U.S. social distancing	0.04*** (0.005)	0.04*** (0.005)	0.04*** (0.005)	0.04*** (0.005)	0.05*** (0.005)	0.04*** (0.005)
Week 10 Interaction	0.002 (0.001)	0.01 (0.01)	0.13 (0.10)	0.002 (0.003)	0.01 (0.01)	0.05 (0.05)
Week 11 Interaction	0.003** (0.001)	0.02** (0.01)	0.22** (0.09)	0.006** (0.003)	0.03*** (0.01)	0.10** (0.05)
Week 12 Interaction	0.01*** (0.001)	0.06*** (0.01)	0.53*** (0.09)	0.01*** (0.002)	0.06*** (0.01)	0.22*** (0.04)
Week 13 Interaction	0.01*** (0.001)	0.04*** (0.01)	0.64*** (0.08)	0.01*** (0.002)	0.05*** (0.01)	0.21*** (0.04)
Week 14 Interaction	0.01*** (0.001)	0.04*** (0.01)	0.65*** (0.08)	0.01*** (0.002)	0.04*** (0.01)	0.23*** (0.04)
Week 15 Interaction	0.01*** (0.001)	0.05*** (0.01)	0.69*** (0.08)	0.02*** (0.002)	0.06*** (0.01)	0.24*** (0.04)
Week 16 Interaction	0.01*** (0.001)	0.05*** (0.01)	0.70*** (0.08)	0.02*** (0.002)	0.05*** (0.01)	0.24*** (0.04)
Week 17 Interaction	0.01*** (0.001)	0.05*** (0.01)	0.73*** (0.08)	0.02*** (0.002)	0.05*** (0.01)	0.23*** (0.04)
Week 18 Interaction	0.01*** (0.001)	0.04*** (0.01)	0.73*** (0.08)	0.02*** (0.002)	0.05*** (0.01)	0.22*** (0.04)
Week 19 Interaction	0.01*** (0.001)	0.04*** (0.01)	0.89*** (0.09)	0.02*** (0.002)	0.05*** (0.01)	0.26*** (0.04)
Week 20 Interaction	0.01*** (0.001)	0.05*** (0.01)	0.90*** (0.08)	0.02*** (0.002)	0.05*** (0.01)	0.25*** (0.04)
Week 21 Interaction	0.01*** (0.001)	0.04*** (0.01)	0.89*** (0.09)	0.02*** (0.002)	0.04*** (0.01)	0.23*** (0.04)
Observations	10,051	9,882	10,051	10,025	10,025	10,051
R-squared	0.92	0.91	0.92	0.92	0.91	0.92

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Municipio fixed effects are controlled in all columns. Each column replicates the regression in Table 3 and adds the interaction of a city characteristic with week fixed effects. Week 9 is the baseline week. The sample is the Week-9-to-21 panel of municipios with at least one Covid-19 case by the end of Week 21.

In the Appendix, we present a wide variety of robustness tests, including analyses using the Unacast and Facebook measures of U.S. social distancing separately, introducing flexible controls for realized cases in the U.S. and Mexico, introducing leads and lags of exposure to U.S. cases, and controlling for Mexican state-level stay-at-home orders. In all cases, the results presented here

are confirmed, and all specification checks yield favorable results.<sup>34</sup> Using estimates with Facebook exposures directly (Table 14), a 14-percentage-point larger decline in average mobility faced by migrants to the U.S. leads to a 4-percentage-point larger decline in mobility in the *municipio*.

To further reinforce our interpretation that U.S. social distancing causes changes in Mexican social distancing, we implement an instrumental variables analysis using U.S. stay-at-home orders as an instrument for observed U.S. social distancing. State-level stay-at-home orders were first imposed in the third week of March and started to be phased out in the last week of April (see Appendix Figure 14). In order for stay-at-home orders to serve as a valid instrument for U.S. social distancing, the orders must drive substantial increases in social distancing in the relevant counties (confirmed shortly in the first-stage analysis), must not be subject to confounding from reverse causality or omitted variables, and must affect Mexican social distancing only through U.S. social distancing. The latter two conditions are likely satisfied in our context, since U.S. policies are unlikely to be influenced by Mexican social distancing behavior and are likely to affect Mexican behavior only through information transmission. Migrant sorting could still pose a concern if migrants from *municipios* that are more likely to comply with social distancing recommendations are more likely to choose destinations that impose stay-at-home orders. While possible, we do not find this concern compelling, as migrants' destinations are primarily driven by enclave locations and economic considerations, and few would have anticipated the emergence of the pandemic or how different states would respond to it.

We begin by constructing each *municipio*'s exposure to U.S. stay-at-home orders as the share of its migrant network in U.S. states with a stay-at-home order in week  $t$ .

$$stayhome\_IV_{it} = \sum_j \frac{m_{ij}}{\sum_{j'} m_{ij'}} \mathbf{1}(stayhome_{jt}) \quad (4)$$

Appendix Figure 15 shows that there is substantial variation in exposure to U.S. stay-at-home orders across *municipios*, even in mid April, when a majority of states had active stay-at-home orders in place. Table 5 shows the first-stage regression relating exposure to U.S. social distancing (Equation 1) to the stay-at-home exposure instrument (Equation 4) and controls for cumulative U.S. and Mexican case counts. In all cases, the coefficient on the instrument is positive and highly statistically significant, yielding first-stage F-statistics of at least 596 and ruling our weak-instrument concerns. The magnitude of the coefficient on the stay-at-home instrument is 0.271, implying that even after controlling for the actual numbers of cases, a one-standard-deviation larger increase in exposure to U.S. stay-at-home order led to a 0.07-standard-deviation larger increase in exposure to U.S. social distancing.<sup>35</sup>

<sup>34</sup>In the main analysis, since some counties are not covered in the Facebook data, the migrant shares do not sum to 1 in Equation 1 when using the Facebook mobility measure in the U.S. In Appendix Section C.4, we show that the results are very similar when we rescale the shares to sum to 1. Out of the 959,089 migrants in the MCAS data, only 1,536 are not in counties covered by the Facebook data (less than 0.2%).

<sup>35</sup>Appendix Figure 16 shows a first-stage residual plot corresponding to Column (1) of Table 5, and Appendix Table 25 performs a similar analysis at the U.S. county level showing the determinants of social distancing behavior in the U.S.

Table 5: First stage: *municipios* with larger exposure to U.S. stay-at-home policies were also more exposed to U.S. social distancing

Outcome: Exposure to U.S. social distancing	(1)	(2)	(3)	(4)
Exposure to U.S. stay-at-home orders	0.271*** (0.017)	0.271*** (0.016)	0.271*** (0.017)	0.271*** (0.016)
Exposure to log U.S. cumulative cases		0.045*** (0.014)		0.045*** (0.014)
Log cum. cases Mexican muni.			0.005*** (0.002)	0.005*** (0.002)
Constant	-0.163*** (0.009)	-0.425*** (0.078)	-0.170*** (0.009)	-0.427*** (0.078)
Observations	10,051	10,051	10,051	10,051
R-squared	0.993	0.993	0.993	0.993
First-stage F-statistic	597	602	596	600

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. All columns include controls for week fixed effects and *municipio* fixed effects. The mean (s.d.) of exposure to U.S. social distancing is -0.02 (1.4), and the mean (s.d.) of the exposure to U.S. stay-at-home orders is 0.54 (0.36). The sample is the Week-9-to-21 panel of *municipios* with at least one Covid-19 case by the end of Week 21.

The instrumental variable results appear in Table 6. We restrict attention to *municipios* with a positive number of cases by Week 21, corresponding to the OLS regressions in Columns (1)–(4) of Table 3.<sup>36</sup> The estimates are quite similar to those in Table 3, confirming our main findings and further ruling out concerns regarding potential migrant sorting in driving the observed relationship between U.S. and Mexican social distancing behaviors.

Table 6: IV results confirm main findings in Table 3.

Outcome: Mexican social distancing	(1)	(2)	(3)	(4)
Exposure to U.S. social distancing	0.046** (0.019)	0.046** (0.019)	0.042** (0.019)	0.042** (0.019)
Exposure to log U.S. cum. cases		-0.012*** (0.003)		-0.014*** (0.003)
Log cum. cases Mexican muni.			0.015*** (0.001)	0.016*** (0.001)
Observations	10,051	10,051	10,051	10,051

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Week fixed effects and *municipio* fixed effects are included in all columns. The sample is the Week-9-to-21 panel of *municipios* with at least one Covid-19 case by the end of Week 21. The exposure to U.S. social distancing is instrumented with the exposure to U.S. stay-at-home orders in all columns.

Together, the various results and robustness tests in this section document a strong and robust relationship between social distancing behavior in the U.S. and reductions in mobility in migrant-

<sup>36</sup>See Appendix Table 26 for the corresponding reduced-form regressions.



connected regions in Mexico. This appears to be a causal relationship that was not driven by disease transmission or migrant sorting between similar regions in the U.S. and Mexico. Instead, the results support the conclusion that receiving information about social distancing from acquaintances, friends, and family living in the U.S. led to increased social distancing in Mexico.

## 5 Heterogeneous effects by origin and destination characteristics

Information transmission and social learning depend not only on the information content itself, but also crucially on how the information is spread and who communicates with whom. For example, BenYishay and Mobarak [2019] show that the social standing of the communicators matters in the process of promoting agricultural technology adoption, and that people who share the same group identity and face comparable agricultural conditions are especially influential. Büchel et al. [2019] show that in migrant networks, local contacts who migrated recently or are more central in the social network have larger impacts on reducing information frictions. In the context of the Covid-19 pandemic, Fan et al. [2020] find that there are substantial gaps in behaviors and beliefs across gender, income, and partisanship lines. These gaps may also influence the effects of information transmission. In this section, we therefore investigate heterogeneity in the effect of exposure to U.S. social distancing based on the characteristics of Mexican *municipios* and of connected U.S. counties in the following sections.

### 5.1 Origin characteristics

We first focus on origin characteristics. As an example, even when facing the same exposure to U.S. social distancing, people living in a *municipio* with higher average educational attainment may react differently than those in a less educated area. For example, people with more education may have more trust in science, which facilitates the adoption of social distancing (Brzezinski et al. 2020). We test for heterogeneity along this and other dimensions by interacting various *municipio* characteristics with the exposure measure in equation (1). The regression is as follows.

$$sodist_{it} = \alpha + \beta \text{exposure}_{it}^{pc} + \gamma \text{exposure}_{it}^{pc} \times C_i + I_i + I_t + \epsilon_{it}, \quad (5)$$

where  $C_i$  is a time-invariant baseline characteristic of *municipio*  $i$ , including population density, urban share of population, share of working age population, average years of education, log earnings per person, and share employed. Note that the *municipio* fixed effects,  $I_i$ , capture the level effect of the characteristic  $C_i$ . To interpret the size of the heterogeneous effects, we first evaluate the impact of the exposure to U.S. social distancing ( $\text{exposure}_{it}^{pc}$ ) at the mean value of  $C_i$  and call it  $\hat{\beta}_1$ . Then we evaluate the effect at the mean plus one-standard deviation of  $C_i$  and call it  $\hat{\beta}_2$ . Finally, we compare the two by calculating  $\hat{\delta} = \hat{\beta}_2/\hat{\beta}_1 - 1$ . A more positive value of  $\hat{\delta}$  indicates that more positive values of  $C_i$  drive more positive effects of U.S. social distancing on Mexican social distancing.<sup>37</sup>

<sup>37</sup>The expressions for  $\hat{\beta}_1$  and  $\hat{\beta}_2$  are as follows:  $\hat{\beta}_1 = \hat{\beta} + \hat{\gamma}\bar{C}$ , and  $\hat{\beta}_2 = \hat{\beta} + \hat{\gamma}(\bar{C} + sd(C))$ , where  $\bar{C}$  is the mean of  $C_i$ , and  $sd(C)$  is the standard deviation of  $C_i$ .

Generally, we find that *municipios* with more favorable socio-economic conditions responded more strongly to U.S. social distancing. Table 7 Column (1) evaluates the heterogeneous effect by the population density. Compared to the effect on a *municipio* with average population density, the effect is 8% larger when the population density is one-standard deviation larger. We find similar heterogeneity when considering the urban population share (7%), working age population share (12%), average years of schooling (10%), log average earnings (7%), and employment share (9%). In Appendix C.6, we also show the heterogeneous effect for the ability to work from home. The results are similar to those in Table 7, and this is consistent with Dingel and Neiman [2020]’s finding that regions with higher incomes also have higher shares of jobs in which working home is feasible.

Table 7: *Municipios* with more favorable socio-economic conditions responded more strongly to U.S. social distancing

Outcome: Mexico social dist.	(1)	(2)	(3)	(4)	(5)	(6)
Exposure to US social dist.	0.04*** (0.005)	0.03*** (0.005)	-0.04*** (0.008)	0.02*** (0.006)	-0.04*** (0.013)	0.02*** (0.006)
Interact: population density	0.002*** (0.000)					
Interact: share urban		0.01*** (0.002)				
Interact: aged 16-65 share			0.15*** (0.012)			
Interact: yrs of schooling				0.003*** (0.000)		
Interact: log income					0.01*** (0.001)	
Interact: % employed						0.05*** (0.006)
Constant	0.22*** (0.000)	0.21*** (0.000)	0.22*** (0.000)	0.22*** (0.000)	0.22*** (0.000)	0.22*** (0.000)
Mean (s.d.) of the interaction $\hat{\delta}$	0.56 (1.8) 8%	0.59 (0.27) 7%	0.62 (0.04) 12%	8.6 (1.4) 10%	8.4 (0.34) 7%	0.51 (0.08) 9%
Observations	10,051	9,882	10,051	10,025	10,025	10,051
R-squared	0.92	0.91	0.92	0.92	0.91	0.92

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Week fixed effects and *municipio* fixed effects are included in all columns. Each column replicates the regressions in Column (1) of Table 3 and adds the interaction of a *municipio* characteristic with the exposure to U.S. social distancing. The sample is the Week-9-to-21 panel of *municipios* with at least one Covid-19 case by the end of Week 21.

## 5.2 Destination characteristics

Migrant destination characteristics may also influence the information transmission process. In addition to the examples mentioned above, Kerr [2008] shows that ethnic ties to home countries among scientific and entrepreneurial communities in the U.S. facilitate international knowledge transfer.

Mexican migrants in the United States may be more likely to learn from people with a similar background. For example, if a destination region has a larger Hispanic community or has a higher share of residents of Mexican descent, the connected *municipios* may learn from them more easily. Learning about social distancing may also be more effective if the destination regions' Hispanic population has higher socio-economic status. If Mexican migrants learn from the general population, then the average education and income level of U.S. counties may also be important.

Table 8: The effect of exposure to U.S. social distancing is not significantly different across *municipios* that are connected to different types of U.S. counties

Outcome: Mexico soc dist.	(1)	(2)	(3)	(4)	(5)	(6)
Exposure to U.S. soc dist.	0.05*** (0.005)	0.05*** (0.005)	-0.03 (0.026)	-0.01 (0.019)	-0.004 (0.077)	0.005 (0.056)
Interact: % Hispanic	-0.001 (0.005)					
Interact: % Mexican		0.007 (0.004)				
Interact: Hispanic education			0.006*** (0.002)			
Interact: education				0.004*** (0.001)		
Interact: log Hispanic income					0.005 (0.007)	
Interact: log income						0.004 (0.005)
Constant	0.22*** (0.000)	0.22*** (0.000)	0.22*** (0.000)	0.22*** (0.000)	0.22*** (0.000)	0.22*** (0.000)
Mean (s.d.) of the interaction $\hat{\delta}$	0.29 (0.08)	0.22 (0.08)	10.3 (0.21) 3.5%	13.1 (0.28) 3%	10.9 (0.06)	11.2 (0.08)
Observations	10,051	10,051	10,051	10,051	10,051	10,051
R-squared	0.91	0.91	0.91	0.91	0.91	0.91

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Week fixed effects and *municipio* fixed effects are included in all columns. Each column replicates the regression in Table 3 and adds the interaction of the average U.S. destination characteristic faced by migrants from each *municipio* characteristic with the *municipio*'s exposure to U.S. social distancing. The sample is the Week-9-to-21 panel of *municipios* with at least one Covid-19 case by the end of Week 21.

In Table 8, we evaluate how the effect of exposure to U.S. social distancing differs by the average characteristics of migrant-connected destination regions. For a destination county characteristic  $x_j$ , we calculate the average value faced by migrants from *municipio*  $i$  as follows.

$$x_i = \sum_j \frac{m_{ij}}{\sum_{j'} m_{ij'}} x_j$$

We then estimate specifications paralleling equation (5), using  $x_i$  as the interaction variable. We find that the effect of exposure is not significantly influenced by the share of Hispanics, the share

of population of Mexican descent, the log Hispanic household income, or the log average income (Columns (1), (2), (5), and (6)). In contrast, we do find significant heterogeneity based on the overall education level in the destination regions and by the education level of Hispanic individuals in the destinations (Columns (3) and (4)). However, the extent of heterogeneity is quite small; using the same  $\hat{\delta}$  measure described in the previous subsection, compared to the effect of exposure of a *municipio* with average education level at the destination, the effect is only 3-4% larger when the destination's education level is one-standard deviation larger.

In sum, we primarily find heterogeneity in learning based on origin characteristics. More affluent Mexican *municipios* responded more strongly to exposure to U.S. social distancing, while the effects do not differ much by observable destination characteristics. One potential explanation is that Mexican residents may not distinguish much between different types of U.S. counties, which is consistent with the fact that U.S. counties are much more homogeneous than Mexican *municipios*.<sup>38</sup>

## 6 Conclusion

People are social entities who learn about information and form beliefs through their social connections. Among various sources of information, friends and family can be especially important when forming beliefs, particularly when there is considerable uncertainty and the stakes are high. In the context of the early-2020 Covid-19 pandemic, we study the effects of migrants' exposure to U.S. social distancing practices on social distancing behavior in Mexico.

Using detailed *municipio*-to-county migrant network data and observed social distancing behavior in U.S. counties based on smartphone tracking data, we construct the exposure to U.S. social distancing for the residents of each Mexican *municipio*. We find that this exposure had a positive impact on the Mexican residents' social distancing behavior, and that this effect was likely driven by learning, rather than assortative matching between origin places and destination places, or the possibility of disease transmission along the network. Mexican regions with more favorable socio-economic conditions responded more strongly to U.S. social distancing exposure, but the effect did not differ significantly based on the characteristics of migrants' locations in the U.S.

Together, these findings highlight the importance of social networks in influencing individuals' compliance with or rejection of public health recommendations in the context of an emerging pandemic. We chose to examine this kind of social learning in the international context because it resolves difficult identification issues that arise in other contexts, since events in Mexico were unlikely to have a significant influence on U.S. social distancing behaviors or policies. However, our conclusions are nonetheless informative regarding the broader importance of personal connections when policy makers seek to change fundamental social behaviors, such as social distancing or wearing masks during a disease outbreak.

<sup>38</sup>For example, as shown in Table 7 and Table 8, the standard deviation of years of schooling across Mexican *municipios* is 1.4, and the standard deviation of years of schooling across connected U.S. counties is 0.28. The standard deviation of years of schooling across 3,195 U.S. counties is 0.73.

## References

- Nicolás Ajzenman, Tiago Cavalcanti, and Daniel Da Mata. More than words: Leaders' speech and risky behavior during a pandemic. *Working paper*, 2020.
- Christoph Albert and Joan Monras. Immigration and spatial equilibrium: the role of expenditures in the country of origin. *working paper*, 2019.
- Hunt Allcott, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Y Yang. Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic. Technical Report w26946, 2020.
- Treb Allen, Cauê de Castro Dobbin, and Melanie Morten. Border walls. *NBER Working Paper*, (25267), 2019.
- Maxim Ananyev, Mikhail Poyker, and Tian Yuan. The safest time to fly: Pandemic response in the era of fox news. Technical report, May 2020.
- Roy M Anderson, Hans Heesterbeek, Don Klinkenberg, and T Déirdre Hollingsworth. How will country-based mitigation measures influence the course of the covid-19 epidemic? *The Lancet*, 395(10228):931–934, 2020.
- Susan Athey, Billy Ferguson, Matthew Gentzkow, and Tobias Schmidt. Experienced segregation. Technical report, Technical Report, Stanford University Working Paper, 2019.
- Abhijit Banerjee, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. Using gossips to spread information: Theory and evidence from two randomized controlled trials. *Review of Economic Studies*, 86: 2453–2490, 2019.
- John M Barrios, Efraim Benmelech, Yael V Hochberg, Paola Sapienza, and Luigi Zingales. Civic capital and social distancing during the covid-19 pandemic. Working Paper 27320, National Bureau of Economic Research, June 2020. URL <http://www.nber.org/papers/w27320>.
- Toman Barsbai, Hillel Rapoport, Andreas Steinmayr, and Christoph Trebesch. The effect of labor migration on the diffusion of democracy: evidence from a former soviet republic. *American Economic Journal: Applied Economics*, 9(3):36–69, 2017.
- Panle Jia Barwick, Yanyan Liu, Eleonora Patacchini, and Qi Wu. Information, mobile communication, and referral effects. Working Paper 25873, National Bureau of Economic Research, May 2019. URL <http://www.nber.org/papers/w25873>.
- Lori Beaman, Ariel BenYishay, Jeremy Magruder, and Ahmed Mushfiq Mobarak. Can network theory-based targeting increase technology adoption? *NBER Working Paper*, (24912), 2018.
- Lori A. Beaman. Social networks and the dynamics of labour market outcomes: Evidence from refugees resettled in the u.s. *Review of Economic Studies*, 79:128–161, 2012.
- Michel Beine, Frédéric Docquier, and Maurice Schiff. International migration, transfer of norms and home country fertility. *Canadian Journal of Economics/Revue canadienne d'économique*, 46(4):1406–1430, 2013.
- Ariel BenYishay and A Mushfiq Mobarak. Social learning and incentives for experimentation and communication. *The Review of Economic Studies*, 86(3):976–1009, 2019.
- Joshua Evan Blumenstock, Guanghua Chi, and Xu Tan. Migration and the value of social networks. 2019.
- Emily Breza, Arun Chandrasekhar, Benjamin Golub, and Aneesha Parvathaneni. Networks in economic development. *Oxford Review of Economic Policy*, 35(4):678–721, 2019.
- Guglielmo Briscese, Nicola Lacetera, Mario Macis, and Mirco Tonin. Compliance with covid-19 social-distancing measures in italy: The role of expectations and duration. Working Paper 26916, National Bureau of Economic Research, March 2020. URL <http://www.nber.org/papers/w26916>.
- Adam Brzezinski, Guido Deiana, Valentin Kecht, and David Van Dijke. The covid-19 pandemic: Government vs. community action across the united states. *Covid Economics* 7, pages 115–156, 2020.

- Konstantin Büchel, Diego Puga, Elisabet Viladecans-Marsal, and Maximilian von Ehrlich. Calling from the outside: The role of networks in residential mobility. 2019.
- Konrad B Burchardi and Tarek A Hassan. The economic impact of social ties: Evidence from german reunification. *The Quarterly Journal of Economics*, 128(3):1219–1271, 2013.
- Leonardo Bursztyjn, Aakaash Rao, Christopher Roth, and David Yanagizawa-Drott. Misinformation during a pandemic. Technical Report 2020-44, 2020.
- Maria Esther Caballero. Origin-country education choices and destination immigration policies. *Working paper*, 2020.
- Maria Esther Caballero, Brian C Cadena, and Brian K Kovak. Measuring geographic migration patterns using matrículas consulares. *Demography*, 55(3):1119–1145, 2018.
- Maria Esther Caballero, Brian C. Cadena, and Brian K. Kovak. The international transmission of local shocks through migration networks. *working paper*, 2020.
- CDC. Social distancing, 2020. URL [cdc.gov/coronavirus/2019-ncov/prevent-getting-sick/social-distancing.html](https://www.cdc.gov/coronavirus/2019-ncov/prevent-getting-sick/social-distancing.html).
- Klaus Desmet and Romain Wacziarg. Understanding spatial variation in covid-19 across the united states. Working Paper 27329, National Bureau of Economic Research, June 2020. URL <http://www.nber.org/papers/w27329>.
- Stephen G. Dimmock, Jiekun Huang, and Scott J. Weisbenner. Give me your tired, your poor, your high-skilled labor: H-1b lottery outcomes and entrepreneurial success. *NBER Working Paper*, (26392), 2019.
- Jonathan I Dingel and Brent Neiman. How many jobs can be done at home? Working paper, 2020.
- Christian Dustmann, Albrecht Glitz, Uta Schönberg, and Herbert Brücker. Referral-based job search networks. *Review of Economic Studies*, 83:514–546, 2016.
- Per-Anders Edin, Peter Frederiksson, and Olof Aslund. Ethnic enclaves and the economic success of immigrants - evidence from a natural experiment. *Quarterly Journal of Economics*, 118(1):329–357, 2003.
- Ying Fan, A. Yesim Orhun, and Dana Turjeman. Heterogeneous actions, beliefs, constraints and risk tolerance during the covid-19 pandemic. Working Paper 27211, National Bureau of Economic Research, May 2020. URL <http://www.nber.org/papers/w27211>.
- David M Gould. Immigrant links to the home country: empirical implications for us bilateral trade flows. *The Review of Economics and Statistics*, pages 302–316, 1994.
- Keith Head and John Ries. Immigration and trade creation: Econometric evidence from canada. *Canadian Journal of Economics*, 31(1):47–62, 1998.
- David Holtz, Michael Zhao, Seth G Benzell, Cathy Y Cao, M Amin Rahimian, Jeremy Yang, Jennifer Nancy Lee Allen, Avinash Collis, Alex Vernon Moehring, Tara Sowrirajan, et al. Interdependence and the cost of uncoordinated responses to covid-19. 2020.
- Solomon Hsiang, Daniel Allen, Sebastien Annan-Phan, Kendon Bell, Ian Bolliger, Trinetta Chong, Hannah Druckenmiller, Andrew Hultgren, Luna Yue Huang, Emma Krasovich, et al. The effect of large-scale anti-contagion policies on the coronavirus (covid-19) pandemic. *medRxiv*, 2020.
- Mounir Karadja and Erik Prawitz. Exit, voice, and political change: Evidence from swedish mass migration to the united states. *Journal of Political Economy*, 127(4):1864–1925, 2019.
- William R Kerr. Ethnic scientific communities and international technology diffusion. *The Review of Economics and Statistics*, 90(3):518–537, 2008.
- Gabriel E Kreindler and Yuhei Miyauchi. Measuring commuting and economic activity inside cities with cell phone records. 2019.
- Theresa Kuchler, Dominic Russel, and Johannes Stroebel. The geographic spread of covid-19 correlates with structure of social networks as measured by facebook. Working Paper 26990, National Bureau of Economic Research, April 2020.

- Maurice Kugler and Hillel Rapoport. Skilled emigration, business networks, and foreign direct investment. *CESifo Working Paper*, (1455), 2005.
- Shan Li. High-skilled immigrant workers and u.s.firms' access to foreign venture capital. *working paper*, 2020.
- Grant Miller and A Mushfiq Mobarak. Learning about new technologies through social networks: experimental evidence on nontraditional stoves in bangladesh. *Marketing Science*, 34(4):480–499, 2015.
- Minnesota Population Center. Integrated public use microdata series, international: Version 7.2. <https://doi.org/10.18128/D020.V7.2>, 2020.
- Simon Mongey, Laura Pilossoph, and Alex Weinberg. Which workers bear the burden of social distancing policies? Working Paper 27085, National Bureau of Economic Research, May 2020. URL <http://www.nber.org/papers/w27085>.
- Kaivan Munshi. Networks in the modern economy: Mexican migrants in the us labor market. *The Quarterly Journal of Economics*, 118(2):549–599, 2003.
- Sonal Pandya and David Leblang. Risky business: Institutions vs. social networks in fdi. *Economics and Politics*, 29(2):91–117, 2017.
- J Raifman, K Nocka, D Jones, J Bor, S Lipson, J Jay, and P Chan. Covid-19 us state policy database. Technical Report Available at: [www.tinyurl.com/statepolicies](http://www.tinyurl.com/statepolicies), 2020.
- James E. Rauch and Vitor Trindade. Ethnic chinese networks in international trade. *Review of Economics and Statistics*, 84(1):116–130, 2002.
- Gobierno de México Secretaría de Salud. Social distancing, 2020. URL <https://www.gob.mx/salud/documentos/sana-distancia>.
- Carlos Serrano Herrera and Rodrigo Jiménez Uribe. Yearbook of migration and remittances: Mexico 2019. Mexico, July 2019.
- Andrey Simonov, Szymon K Sacher, Jean-Pierre H Dubé, and Shirsho Biswas. The persuasive effect of fox news: Non-compliance with social distancing during the covid-19 pandemic. Working Paper 27237, National Bureau of Economic Research, May 2020. URL <http://www.nber.org/papers/w27237>.
- Unacast. Unacast social distancing dataset. Technical Report <https://www.unacast.com/data-for-good>. Version from 18 April 2020., 2020.
- WHO. Coronavirus disease (covid-19) advice for the public, 2020. URL <https://www.who.int/emergencies/diseases/novel-coronavirus-2019/advice-for-public>.
- Austin L Wright, Konstantin Sonin, Jesse Driscoll, and Jarnicka Wilson. Poverty and economic dislocation reduce compliance with covid-19 shelter-in-place protocols. Technical Report 2020-40, 2020.

## Appendix

### A Data

This section presents additional summary statistics on the geographic variation in exposure to social distancing behavior across U.S. counties for each Mexican source region (*municipio*).

#### A.1 Weeks of the year in 2020

Table 9 shows the dates for each week of the year covered in both Facebook and Unacast datasets used to measure local mobility, as explained in section 3.

Table 9: Week of the year table, 2020

Week Number	From Date	To date
Week 9	February 24	March 1
Week 10	March 2	March 8
Week 11	March 9	March 15
Week 12	March 16	March 22
Week 13	March 23	March 29
Week 14	March 30	April 5
Week 15	April 6	April 12
Week 16	April 13	April 19
Week 17	April 20	April 26
Week 18	April 27	May 3
Week 19	May 4	May 10
Week 20	May 11	May 17
Week 21	May 18	May 24



## A.2 Additional mobility data summary statistics

Table 10 shows the Facebook data coverage by week in Mexico and in the United States. The coverage varies by week since the number of unique active users may change from week to week.

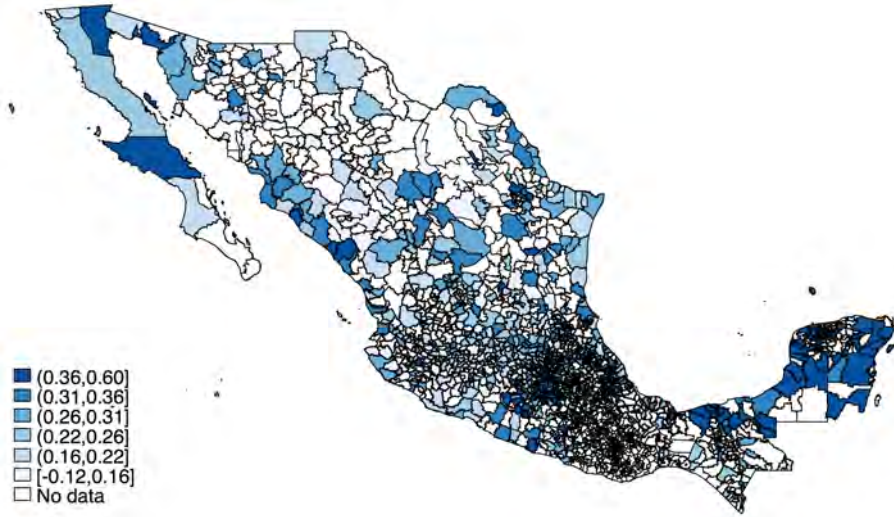
Table 10: Geographic coverage of Facebook mobility data in the U.S. and in Mexico

Week	Num. US counties	Num. MX <i>municipios</i>
9	2,656	1,050
10	2,662	1,060
11	2,662	1,066
12	2,655	1,068
13	2,656	1,074
14	2,658	1,078
15	2,658	1,083
16	2,653	1,081
17	2,650	1,078
18	2,644	1,081
19	2,641	1,081
20	2,637	1,079
21	2,645	1,076
Any week	2,691	1,084

Note: This table presents the number of U.S. counties and Mexican *municipios* covered by the Facebook mobility data. The number of regions covered vary by week due to the constraint that only regions with more than 300 unique users are included. In the Unacast data, 3,054 US counties are covered for all weeks (9–21).

As discussed in section 3, Figure 8 maps the change in the social distancing measure from the Facebook dataset across Mexican *municipios*, used as our main dependent variable in equation 3. There was substantial geographic variation in the increase in social distancing across Mexican *municipios* from Week 9 to Week 21, with Mexican regions in dark blue representing places with larger declines in mobility.

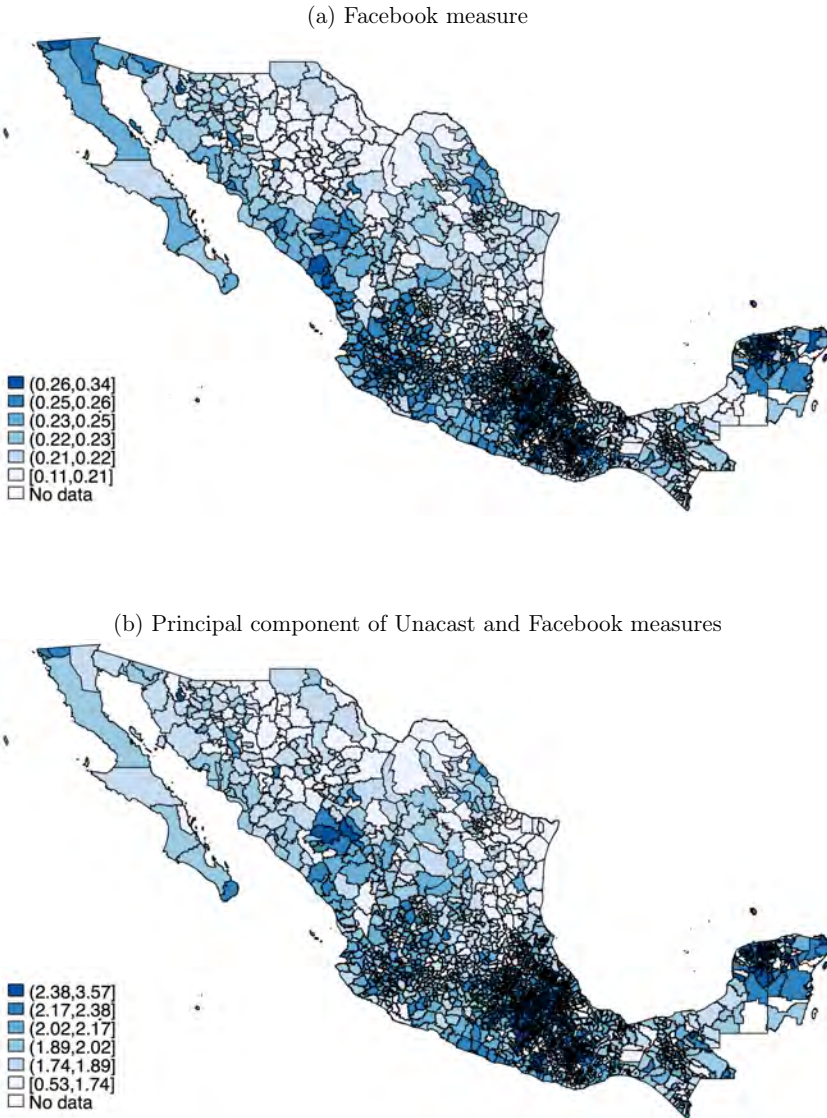
Figure 8: Distribution of changes in social distancing in Mexico, Week 9 to Week 21



Note: The changes in social distancing in Mexico are calculated as  $socdist_{21} - socdist_9$ , using the Facebook data. There are 989 municipios with non-missing values of social distancing in Week 9, and 1,010 municipios in Week 21.

Panel (a) of Figure 9 maps the change in exposure faced by each Mexican *municipio* to U.S. social distancing from the Facebook dataset, while Panel (b) maps the principal component of the Unacast and Facebook social distancing measures as defined in 1. These measures combined geographic variation in U.S. social distancing behavior with geographic variation in the destination distribution of Mexican source regions. This creates the geographic differences in exposure for each Mexican region to different social distancing practices in the U.S. observed in Figure 9.

Figure 9: Distribution of changes in exposure to social distancing in the United States, Week 9 to Week 21

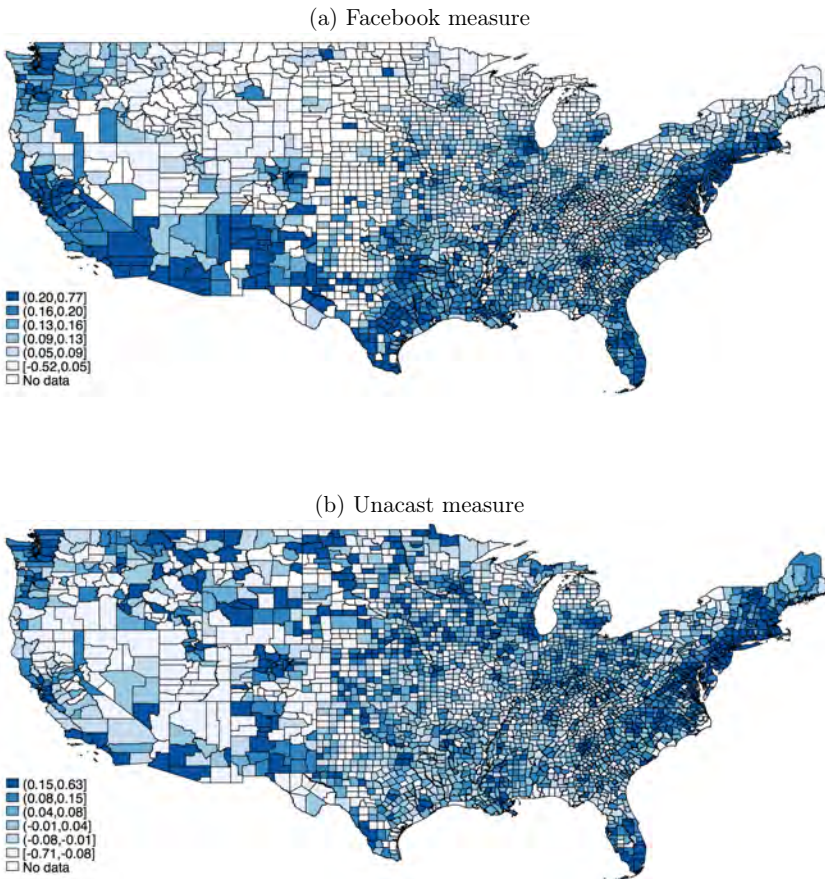


Note: The changes in exposure to social distancing in the United States are calculated as  $exposure_{21}^s - exposure_9^s$ , where  $s =$  Facebook in Panel (a) and  $s = pc$  in Panel (b). There are 46 municipios with no data, and 2,415 municipios with data.

Covid Economics 54, 29 October 2020: 62-120

Panel (a) of Figure 10 maps the change in the social distancing measure from the Facebook dataset across U.S. counties, while panel (b) of Figure 10 shows the same measure using data from Unacast. There was a great deal of geographic variation in the increase in social distancing across U.S. counties from Week 9 to Week 21, with counties in dark blue representing places with larger mobility declines. These maps show part of the geographic variation in social distancing behavior that we use to construct our exposure measure as defined in equation 1.

Figure 10: Distribution of changes in social distancing in the United States, Week 9 to Week 21



Note: The changes in social distancing in the United States are calculated as  $socdist_{21} - socdist_9$ . Panel (a) uses the Facebook data and includes 2,531 counties, and Panel (b) uses the Unacast data and includes 3,033 counties. Hawaii and Alaska are not included.

Covid Economics 54, 29 October 2020: 62-120

### A.3 Sample restrictions

Table 11 shows the sample size restrictions yielding the 13,036 observations in Table 3 Columns (5)–(8). There are 2,411 *municipios* from the MCAS dataset, after excluding Yaxkukul in the State of Yucatan with only one migrant in one U.S. county. There are 1,083 *municipios* with Facebook mobility measures. There are 1,013 *municipios* and 13,037 *municipio*-week observations satisfying both conditions. One *municipio*, San Miguel De Horcasitas in the State of Sonora, only has the mobility measure in Week 21 and is excluded in the panel regression as a singleton.

Table 11: How we arrive at the final sample size


Week	Num. MX <i>municipios</i>		
	(1)	(2)	(3)
9	2,411	1,049	984
10	2,411	1,059	991
11	2,411	1,065	997
12	2,411	1,067	998
13	2,411	1,073	1,002
14	2,411	1,077	1,007
15	2,411	1,082	1,012
16	2,411	1,080	1,009
17	2,411	1,077	1,006
18	2,411	1,080	1,009
19	2,411	1,080	1,009
20	2,411	1,078	1,008
21	2,411	1,075	1,005
Any week	2,411	1,083	1,013
With mobility exposure measure	Yes		Yes
With mobility measure		Yes	Yes
Total number of obs.	31,343	13,942	13,037

Note: This table presents sample size for Mexican *municipios* covered in the analysis. Yaxkukul in the State of Yucatan is dropped since the population size is very small (2,868 in 2010) and it is only has one destination county with one migrant count in the MCAS dataset, Horry in South Carolina. MCAS data includes 2412 *municipios*. Thus, 2,411 *municipios* have the measure of exposure to U.S. social distancing after dropping Yaxkukul (Column 1). In Column (2), there are 1,083 *municipios* with Facebook mobility measure. When we restrict to the *municipio*-weeks with both the Facebook mobility measure and the exposure to U.S. social distancing, we have 1,013 *municipios*. The panel regression in the main analysis with all *municipios* includes 13,036 observations instead of 13,037 in Column (3) since San Miguel De Horcasitas in the State of Sonora only has mobility measure in Week 21 and is excluded in the panel regression as a singleton.

## B Mexican state-level stay-at-home orders

Figure 11 and Table 12 describe state-level stay-at-home orders across Mexican states, based on Mexican States' official decrees. Table 12 provides details on the specific measures imposed by each state, along with the date of the relevant decree, and Figure 11 depicts the decrees graphically, with blue bars showing weeks in which relevant decrees were in place. States without specific stay-at-home orders are omitted from Figure 11 and Table 12 (see the note to Table 12 for a list). These states declared states of emergency and closure of nonessential businesses in the first week of April following the federal government order.

Figure 11: Mexico State-level Stay-at-home Orders, by week



Note: This figure shows the Mexican states imposing mandatory stay-at-home orders or mobility restrictions in the weeks under study (see Table 12 for details) based on Mexican States' Official decrees. The blue bars represent the week in which a state had an active stay-at-home order.

Table 12: Mexico State-level Stay-at-home Orders

State	Measures	Date
Baja California Sur	Imposed measures to restrict mobility within the state. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Friday, April 24, 2020
Chihuahua	Installed check points in main highways and roads.	Sunday, April 19, 2020
Coahuila	Imposed measures to restrict mobility within the state. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Wednesday, April 22, 2020
Colima	Imposed measures to restrict mobility within the state. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Thursday, April 9, 2020
Distrito Federal	Imposed measures to restrict mobility within the state. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Wednesday, April 22, 2020
Durango	Imposed measures to restrict mobility within the state. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Sunday, April 26, 2020
Jalisco	Mandatory stay-at-home measures were imposed. Penalties included fines.	Monday, April 20, 2020
México	Imposed measures to restrict mobility within the state. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Wednesday, April 22, 2020
Michoacán	Mandatory stay-at-home measures were imposed. Penalties included fines and jail time.	Monday, April 20, 2020
Nayarit	Imposed measures to restrict mobility within the state. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Saturday, April 18, 2020
Sinaloa	Following the federal government announcement, the state of emergency was extended and the closure of nonessential businesses continued. In addition, measures to restrict mobility within the state were imposed. Lowered public transportation capacity, and limit the number of persons who could travel in personal vehicles.	Wednesday, April 22, 2020
Sonora	State of emergency was declared and nonessential businesses were ordered to close, before the announcement from the federal government was made.	Wednesday, March 25, 2020
	Mandatory stay-at-home measures were imposed. Penalties included fines and jail time.	Monday, April 13, 2020
Tabasco	Following the federal government announcement, the state of emergency was extended and the closure of nonessential businesses continued. In addition, measures to restrict mobility within the state were imposed. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Tuesday, April 21, 2020
Tamaulipas	Following the federal government announcement, the state of emergency was extended and the closure of nonessential businesses continued. In addition, measures to restrict mobility within the state were imposed. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Thursday, April 23, 2020
Yucatán	Following the federal government announcement, the state of emergency was extended and the closure of nonessential businesses continued. In addition, measures to restrict mobility within the state were imposed. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Thursday, April 23, 2020
Zacatecas	Following the federal government announcement, the state of emergency was extended and the closure of nonessential businesses continued. In addition, measures to restrict mobility within the state were imposed. Lowered public transportation capacity and limit the number of persons who could travel in personal vehicles.	Wednesday, April 8, 2020

Note: This table presents a description of the mandatory stay-at-home orders or mobility restrictions imposed by each Mexican state government as well as the dates for each mandate, based on Mexican States' Official decrees. The following states declared states of emergency and closure of nonessential businesses on the first week of April along with the federal government order: Aguascalientes, Baja California, Hidalgo, Morelos, Nuevo León, Oaxaca, and Tlaxcala. Between the third and fourth week of April the following states extended the state of emergency and maintained closure of nonessential businesses: Campeche, Chiapas, Guanajuato, Guerrero, Puebla, Querétaro, Quintana Roo, San Luis Potosí, and Veracruz.

## C Additional empirical results

This section outlines several robustness checks to support the validity of our main results presented in section 4. Our main results are robust to 1) including controls for Mexican state-level stay-at-home orders, 2) dropping outlier regions in Mexico, 3) introducing lagged exposure measures, 4) using the exposure measure constructed from Facebook and Unacast data separately instead the principal component exposure measure, 5) flexibly controlling for the local cases and the exposure to U.S. cases, and 6) including Mexican *municipios* with no cases.

### C.1 Robustness of of the main results after controlling for Mexican state-level stay-at-home orders

Table 13 replicates Table 3 in our main analysis with an additional control for stay-at-home orders imposed in Mexican states that differ from those imposed by the federal government (as described in Appendix B). The results are nearly identical to those of Table 3.

Table 13: Larger exposure to U.S. social distancing led to more social distancing in Mexico, Week 9 to Week 21

Outcome:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mexico social dist.	<i>Municipios with cases&gt;0</i>				<i>All municipios</i>			
Exposure to U.S. social dist.	0.05*** (0.005)	0.05*** (0.005)	0.04*** (0.005)	0.05*** (0.005)	0.03*** (0.004)	0.03*** (0.004)	0.03*** (0.004)	0.03*** (0.004)
Exposure to log U.S. cum. cases		-0.01*** (0.003)		-0.01*** (0.003)		-0.01*** (0.002)		-0.01*** (0.002)
Log cum. cases in Mexico muni.			0.02*** (0.001)	0.02*** (0.001)			0.01*** (0.001)	0.02*** (0.001)
Mexico state-level stay-at-home orders	0.003 (0.002)	0.004** (0.002)	-0.000 (0.002)	0.001 (0.002)	-0.006*** (0.002)	-0.005*** (0.002)	-0.009*** (0.001)	-0.008*** (0.001)
Constant	0.21*** (0.0001)	0.29*** (0.015)	0.19*** (0.001)	0.28*** (0.015)	0.21*** (0.001)	0.26*** (0.013)	0.20*** (0.001)	0.26*** (0.013)
Observations	10,051	10,051	10,051	10,051	13,036	13,036	13,036	13,036
R-squared	0.91	0.91	0.92	0.92	0.91	0.91	0.91	0.91

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. This table replicates Table 4 by including controls for Mexican state-level stay-at-home orders. Week fixed effects and *municipio* fixed effects are included in all columns. Columns (1)–(4) include the *municipios* with at least one Covid-19 case at the end of Week 21, and Columns (5)–(8) include all *municipios*. The mean (s.d.) of Mexican social distancing in the first four columns is 0.21 (0.15), and the mean (s.d.) of the exposure to U.S. social distancing is -0.02 (1.4). The mean (s.d.) of the log cumulative cases in Mexico is 1.4 (1.8), and the mean (s.d.) of the exposure to U.S. cases is 5.1 (2.7). The corresponding numbers for the last four columns are: 0.21 (0.15), -0.1 (1.4), 1.1 (1.7), and 5.1 (2.7).

### C.2 Robustness of the correlations

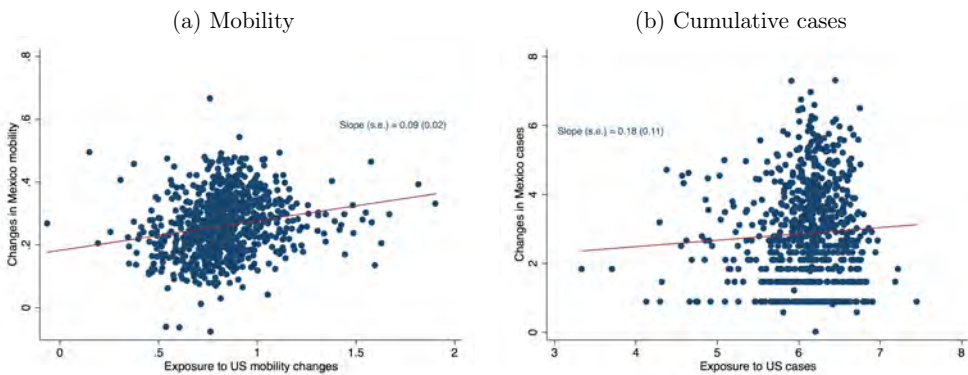
Figure 12 Panel (a) replicates Figure 7 by dropping an outlier *municipio*, San José Miahuatlán in Puebla State. This figure relates the long-difference change in local social distancing to the change



in the *municipio's* exposure to U.S. social distancing. It shows that the strong positive correlation between changes in social distancing in Mexico and the U.S. remains after dropping out the outlier *municipio*, suggesting that the results are not driven by outliers.

Panel (b) shows the corresponding relationship between changes in the log cumulative cases in Mexican *municipios* and changes in the exposure to cumulative U.S. cases. The horizontal axis is the change in exposure to U.S. cumulative cases ( $exposure_{i21}^{cases} - exposure_{i11}^{cases}$ ), and the vertical axis is the change in the log cumulative cases in Mexico ( $\ln(\text{cum cases})_{i21} - \ln(\text{cum cases})_{i11}$ ). The fitted line has slope of 0.18, statistically insignificant.

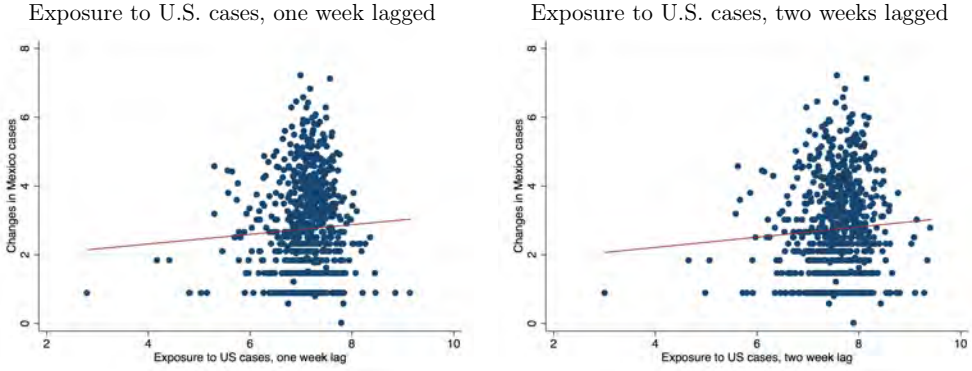
Figure 12: A strong positive correlation between changes in social distancing in Mexico and the U.S., replicating Figure 7 by dropping an outlier, San José Miahuatlán in Puebla State.



Note: This figure includes 769 Mexican *municipios* with at least one Covid-19 case in Week 21, and each dot is a *municipio*. It replicates Figure 7 by dropping an outlier, San José Miahuatlán in Puebla State. Panel (a) shows the mobility result, where the horizontal axis is the exposure to U.S. social distancing in Week 21 minus that in Week 11 ( $exposure_{i21}^{pc} - exposure_{i11}^{pc}$ ), and the vertical axis is the change in social distancing in a Mexican *municipio* between Week 21 and Week 11 ( $socdist_{i21} - socdist_{i11}$ ). The mean (s.d.) of the x-axis is 0.8 (0.2), and the mean (s.d.) of the y-axis is 0.3 (0.1). Panel (b) shows the cumulative case result, where the horizontal axis is the change in exposure to log U.S. cumulative cases from Week 11 to Week 21 ( $exposure_{i21}^{cases} - exposure_{i11}^{cases}$ ), and the vertical axis is the change in the log of cumulative cases in a Mexican *municipio* between Week 21 and Week 11 ( $\ln(\text{cum cases})_{i21} - \ln(\text{cum cases})_{i11}$ ). The mean (s.d.) of the x-axis is 6.1 (0.5), and the mean (s.d.) of the y-axis is 2.9 (1.4).

Figure 13 replicates Figure 7 Panel (b), but using one-week and two-week lagged values of the change in exposure to U.S. social distancing to allow for potential delays in information transmission. These figures show that the relationship between the number of cases in Mexican *municipios* and the exposure to U.S. cases remains unchanged.

Figure 13: The relationship between the number of cases in Mexican *municipios* and the exposure to U.S. cases does not change if we use lagged exposure.



Note: This figure includes 770 Mexican *municipios* with at least one Covid-19 case in Week 21, and each dot is a *municipio*. It replicates Figure 7 Panel (b) by using the one-week and two-week lagged exposure to U.S. cases as the horizontal axis. The horizontal axis in Panel (a) is the change in exposure to log U.S. cumulative cases from Week 10 to Week 20 ( $exposure_{i20}^{cases} - exposure_{i10}^{cases}$ ), and the vertical axis is the change in the log of cumulative cases in a Mexican *municipio* between Week 21 and Week 11 ( $\ln(\text{cum cases})_{i21} - \ln(\text{cum cases})_{i11}$ ). The horizontal axis in Panel (b) is the change in exposure to log U.S. cumulative cases from Week 9 to Week 19 ( $exposure_{i19}^{cases} - exposure_{i9}^{cases}$ ).

Covid Economics 54, 29 October 2020: 62-120

### C.3 Robustness of the main results using the Facebook and Unacast measures separately

Tables 14 and 15 replicate Table 3 in our main analysis, but separately use the exposure to U.S. social distancing constructed with the Facebook and Unacast data, respectively. The tables show that the results are robust to constructing the exposure to social distancing practices in the U.S. separately for each dataset as opposed to constructing it as the principal component of the two social distancing measures together.

Table 14: Larger exposure to U.S. social distancing led to more social distancing in Mexico, Week 9 to Week 21, Facebook measure as the outcome

Outcome:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mexico social dist.	<i>Municipios</i> with cases>0				All <i>municipios</i>			
Exposure to U.S. social dist. (Facebook measure)	0.27*** (0.04)	0.29*** (0.04)	0.25*** (0.04)	0.26*** (0.04)	0.12*** (0.04)	0.12*** (0.04)	0.10*** (0.04)	0.11*** (0.04)
Exposure to log U.S. cum. cases		-0.01*** (0.00)		-0.01*** (0.00)		-0.01*** (0.00)		-0.01*** (0.00)
Log cum. cases in Mexico muni.			0.02*** (0.00)	0.02*** (0.00)			0.01*** (0.00)	0.01*** (0.00)
Constant	0.15*** (0.01)	0.21*** (0.02)	0.13*** (0.01)	0.21*** (0.02)	0.18*** (0.01)	0.22*** (0.02)	0.17*** (0.01)	0.23*** (0.02)
Observations	10,051	10,051	10,051	10,051	13,036	13,036	13,036	13,036
R-squared	0.91	0.91	0.92	0.92	0.91	0.91	0.91	0.91

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. This table replicates Table 4 by using the exposure to U.S. social distancing that is measured using the Facebook data. Week fixed effects and *Municipio* fixed effects are controlled in all columns. Columns (1)–(4) include the *municipios* with at least one Covid-19 case at the end of Week 21, and Columns (5)–(8) include all *municipios*. The mean (s.d.) of Mexican social distancing in the first four columns is 0.21 (0.15), and the mean (s.d.) of the exposure to U.S. social distancing is 0.24 (0.14). The mean (s.d.) of the log cumulative cases in Mexico is 1.4 (1.8), and the mean (s.d.) of the exposure to U.S. cases is 5.1 (2.7). The corresponding numbers for the last four columns are: 0.21 (0.15), 0.24 (0.14), 1.1 (1.7), and 5.1 (2.7).

Table 15: Larger exposure to U.S. social distancing led to more social distancing in Mexico, Week 9 to Week 21, Unacast measure as the outcome

Outcome:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mexico social dist.	Municipios with cases>0				All municipios			
Exposure to U.S. social dist. (Unacast measure)	0.41*** (0.04)	0.43*** (0.04)	0.38*** (0.04)	0.40*** (0.04)	0.32*** (0.03)	0.34*** (0.03)	0.29*** (0.03)	0.31*** (0.03)
Exposure to log U.S. cum. cases		-0.01*** (0.00)		-0.02*** (0.00)		-0.01*** (0.00)		-0.01*** (0.00)
Log cum. cases in Mexico muni.			0.02*** (0.00)	0.02*** (0.00)			0.01*** (0.00)	0.01*** (0.00)
Constant	0.09*** (0.01)	0.16*** (0.02)	0.08*** (0.01)	0.16*** (0.02)	0.11*** (0.01)	0.17*** (0.02)	0.11*** (0.01)	0.17*** (0.02)
Observations	10,051	10,051	10,051	10,051	13,036	13,036	13,036	13,036
R-squared	0.91	0.91	0.92	0.92	0.91	0.91	0.91	0.91

Note: Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . This table replicates Table 3 by using the exposure to U.S. social distancing that is measured using the Unacast data. Week fixed effects and Municipio fixed effects are controlled in all columns. Columns (1)–(4) include the municipios with at least one Covid-19 case at the end of Week 21, and Columns (5)–(8) include all municipios. The mean (s.d.) of Mexican social distancing in the first four columns is 0.21 (0.15), and the mean (s.d.) of the exposure to U.S. social distancing is 0.29 (0.16). The mean (s.d.) of the log cumulative cases in Mexico is 1.4 (1.8), and the mean (s.d.) of the exposure to U.S. cases is 5.1 (2.7). The corresponding numbers for the last four columns are: 0.21 (0.15), 0.29 (0.16), 1.1 (1.7), and 5.1 (2.7).

**C.4 Robustness of the main results using rescaled Facebook exposure measure**

Since the Facebook data for the United States do not cover all U.S. counties, it is possible that counties covered in the MCAS data are not included in the Facebook data. When this is the case, the shares in equation (1) do not sum to 1.

Out of the 959,089 migrants in the MCAS data, 1,536 are in counties not covered by the Facebook data (less than 0.2%). We construct the share of migrants in MCAS data covered in Facebook counties for each *municipio*, and rescale the exposure to U.S. social distancing using Facebook data to make the shares to sum to 1. Then we construct the exposure to U.S. social distancing using the principal component of the rescaled Facebook exposure and the Unacast measure.

Table 16 presents the results using the rescaled measures, where Columns (1)–(4) replicate Table 3 Columns (1)–(4) with the principal component exposure measure, and Columns (5)–(8) replicate Table 14 Columns (1)–(4) with the Facebook exposure measure. The results are very similar. This is not surprising since in the sample used in Table 16, the mean (s.d.) of the share of migrants in counties covered by the Facebook data is 0.998 (0.007), with a minimum of 0.86 and the maximum of 1.

Table 16: Results robust to using rescaled Facebook exposure measures

Outcome: Mexico social dist.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Principal component measure				Facebook measure			
Exposure to U.S. social dist.	0.046*** (0.005)	0.048*** (0.005)	0.042*** (0.005)	0.045*** (0.005)	0.272*** (0.045)	0.281*** (0.045)	0.243*** (0.044)	0.253*** (0.044)
Log cum. cases in Mexico muni.		-0.012*** (0.003)		-0.014*** (0.003)		-0.010*** (0.003)		-0.013*** (0.003)
Exposure to log U.S. cum. cases			0.015*** (0.001)	0.016*** (0.001)			0.015*** (0.001)	0.016*** (0.001)
Constant	0.215*** (0.000)	0.283*** (0.015)	0.194*** (0.001)	0.276*** (0.015)	0.149*** (0.011)	0.207*** (0.019)	0.135*** (0.011)	0.207*** (0.019)
Observations	10,051	10,051	10,051	10,051	10,051	10,051	10,051	10,051
R-squared	0.913	0.913	0.919	0.919	0.912	0.913	0.918	0.918

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. This table replicates Table 4 Columns (1)–(4) and Table 14 (1)–(4) by using the exposure to U.S. social distancing using rescale Facebook exposure measure. Since the Facebook data in the U.S. does not cover all counties, the migrant shares in the migration network data does not sum up to 1. In the sample used in this table, the mean (s.d.) of the share of migrants in counties covered by the Facebook data is 0.998 (0.007), with a minimum of 0.86 and the maximum of 1. Here we rescale the exposure to Facebook U.S. social distancing such that the migrant shares sum up to 1. Columns (1)–(4) use the principal component of the rescaled Facebook measure and the Unacast measure, and Columns (5)–(8) use the rescale Facebook measure. Week fixed effects and *municipio* fixed effects are controlled in all columns. The sample is the Week-9-to-21 panel of *municipios* with at least one Covid-19 case by the end of Week 21.

**C.5 Robustness of the main results when flexibly controlling for local and U.S. cases**

Table 17 replicates Table 3 Columns (4) and (8) in our main analysis, but controls for flexible functional forms of the number of cases. The results are nearly identical after including these flexible controls, ruling out concerns about the disease transmission along the migrant network as the underlying channel of our main results.

Table 17: Larger exposure to U.S. social distancing led to more social distancing in Mexico, Week 9 to Week 21, Unacast measure as the outcome

Outcome:	(1)	(2)	(3)	(4)	(5)	(6)
Mexico social dist.	<i>Municipios with cases&gt;0</i>			<i>All municipios</i>		
Exposure to U.S. social dist.	0.045*** (0.005)	0.048*** (0.005)	0.048*** (0.005)	0.029*** (0.004)	0.032*** (0.004)	0.031*** (0.004)
Exposure to log U.S. cum. cases	-0.015*** (0.003)	-0.012*** (0.003)	-0.013*** (0.003)	-0.012*** (0.002)	-0.009*** (0.002)	-0.010*** (0.002)
Log cum. cases in Mexico muni.	0.005*** (0.001)			0.005*** (0.001)		
Log cum. cases in Mexico muni. squared	0.002*** (0.000)			0.002*** (0.000)		
I (number of cum. cases > 0)		-0.001 (0.002)			0.009*** (0.001)	
I (number of cases in (0, 100] )			0.007*** (0.002)			0.013*** (0.001)
I (number of cases in (100, 1000] )			0.055*** (0.004)			0.066*** (0.004)
I (number of cases > 1000)			0.108*** (0.007)			0.124*** (0.007)
Constant	0.286*** (0.015)	0.285*** (0.015)	0.286*** (0.015)	0.264*** (0.013)	0.256*** (0.013)	0.259*** (0.013)
Observations	10,051	10,051	10,051	13,036	13,036	13,036
R-squared	0.92	0.91	0.92	0.91	0.91	0.91

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. This table replicates Table 3 Columns (4) and (8) by using the different types of measures of severity of the local outbreak. Week fixed effects and *municipio* fixed effects are controlled in all columns. Columns (1)–(3) include the *municipios* with at least one Covid-19 case at the end of Week 21, and Columns (4)–(6) include all *municipios*. The mean (s.d.) of Mexican social distancing in the first four columns is 0.21 (0.15), and the mean (s.d.) of the exposure to U.S. social distancing is -0.02 (1.4). The mean (s.d.) of the log cumulative cases in Mexico is 1.4 (1.8), and the mean (s.d.) of the exposure to U.S. cases is 5.1 (2.7). The corresponding numbers for the last four columns are: 0.21 (0.15), -0.01 (1.4), 1.1 (1.7), and 5.1 (2.7). Columns (1) and (4) include the log cumulative cases in Mexico and the squared term. Columns (2) and (5) include a dummy variable indicating in this *municipio* and week, if there is a positive number of cumulative cases. Columns (3) and (6) include a dummy if the number of cumulative cases is in between of 0 and 100, 100 to 1000, and larger than 1000.

Covid Economics 54, 29 October 2020: 62-120

Table 18 replicates Table 3 in our main analysis, but includes leads and lags of the exposure to U.S. social distancing to rule out reverse causality. In this case, we see that the coefficients leading up to the beginning of our period of study are not statistically different from zero, suggesting that that U.S. social distancing are in fact transmitted to Mexico through the migrant network and not the other way around.

Table 18: Results in Table 3 are robust to controlling for leads and lags of exposure to U.S. cases.

Outcome: social distancing in Mexico	(1) Muni with cases > 0	(2)	(3) All muni.	(4)
Exposure to U.S. social dist.	0.05*** (0.005)	0.04*** (0.005)	0.03*** (0.004)	0.03*** (0.004)
Log cum. cases in Mexico muni.	0.02*** (0.001)	0.02*** (0.001)	0.02*** (0.001)	0.02*** (0.001)
Exposure to log U.S. cum. cases	-0.03*** (0.009)	-0.02*** (0.005)	-0.02** (0.008)	-0.01*** (0.004)
Exposure to log U.S. cum. cases, lagged one period	0.007 (0.005)		-0.001 (0.004)	
Exposure to log U.S. cum. cases, lead one period	0.006 (0.007)		0.008 (0.006)	
Exposure to log U.S. cum. cases, lagged one week		0.008*** (0.003)		-0.001 (0.002)
Exposure to log U.S. cum. cases, lead two weeks		0.004 (0.005)		0.006 (0.005)
Constant	0.266*** (0.020)	0.239*** (0.027)	0.240*** (0.018)	0.214*** (0.024)
Observations	10,051	9,276	13,036	12,032
R-squared	0.919	0.923	0.913	0.918

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. This table replicates Table 3 by controlling for leads and lags of the exposure to U.S. cases. Week fixed effects and *municipio* fixed effects are controlled in all columns. Columns (1)–(4) include the *municipios* with at least one Covid-19 case at the end of Week 21, and Columns (5)–(8) include all *municipios*. The mean (s.d.) of Mexican social distancing in the first four columns is 0.21 (0.15), and the mean (s.d.) of the exposure to U.S. social distancing is -0.02 (1.4). The mean (s.d.) of the log cumulative cases in Mexico is 1.4 (1.8), and the mean (s.d.) of the exposure to U.S. cases is 5.1 (2.7). The corresponding numbers for the last four columns are: 0.21 (0.15), -0.01 (1.4), 1.1 (1.7), and 5.1 (2.7).

## C.6 Robustness of the main results by taking into account the share of jobs facilitating work from home

As shown in Dingel and Neiman [2020], different industries and occupations have different shares of jobs that can be performed at home. In Table 19, we present the crosswalk of industries in Mexico and in the United States. The share of jobs facilitating work from home at the 2-digit NAICS sector level is from Table 3 in Dingel and Neiman [2020], and out of the 20 industries, 14 industries have direct matches with the IPUMS general industry code used in the 2015 Intercensal Count (Panel A), and 6 industries do not have exact matches (Panel B). In later analysis, we construct the *municipio*-level shares allowing work from home using the individual level industry code in the 2015 Intercensal Count and the Mexican industry level workable-at-home job shares. Given the imperfect matching, we use two matching methods. In the first one, “Other services” and “Private household services” in Mexico are assigned the value of 0.31 and 0.43 (unweighted and weighted by wage) to match “Other services (except for public administration)” in the U.S. The second method match these two Mexican industries to the average of unmatched U.S. industries, including “Professional, scientific, and technical services”, “Management of companies and enterprises”, “Information”, “Other services (except public administration)”, “Administrative and support and waste management and remediation services”, and “Arts, entertainment, and recreation”.



Table 19: Share of jobs facilitating work at home, by industry in Mexico

Panel A. Matched		Share of jobs doable at home	
Mexican industry (IPUMS general industry)	US industry (2-digit NAICS sector)	Unweighted	Weighted by wage
Agriculture, fishing, and forestry	Agriculture, forestry, fishing and hunting	0.08	0.13
Mining and extraction	Mining, quarrying, and oil and gas extraction	0.25	0.37
Manufacturing	Manufacturing	0.22	0.36
Electricity, gas, water and waste management	Utilities	0.37	0.41
Construction	Construction	0.19	0.22
Wholesale and retail trade	Wholesale trade	0.52	0.67
	Retail trade	0.14	0.22
Hotels and restaurants	Accommodation and food services	0.04	0.07
Transportation, storage, and communication	Transportation and warehousing	0.19	0.25
Financial services and insurance	Finance and insurance	0.76	0.85
Public administration and defense	Federal, state, and local government	0.41	0.47
Business services and real estate	Real estate and rental and leasing	0.42	0.54
Education	Educational services	0.83	0.71
Health and social work	Health care and social assistance	0.25	0.24
Panel B. Unmatched			
Other services	Professional, scientific, and technical services	0.80	0.86
Private household services	Management of companies and enterprises	0.79	0.86
	Information	0.72	0.80
	Other services (except public administration)	0.31	0.43
	Administrative and support and waste management and remediation services	0.31	0.43
	Arts, entertainment, and recreation	0.30	0.36

Notes: This table reports the crosswalk between Mexican industries and U.S. industries, where U.S. industries have the share of jobs doable at home from Dingel and Neiman [2020]. The U.S. industries are at the 2-digit NAICS sector level, and the Mexican industries are from the IPUMS International general industry code, where the grouping “roughly conform to the International Standard Industrial Classification (ISIC)” (IPUMS International). Panel A shows the list of matched industries, and Panel B shows the list of unmatched industries.

Table 20 controls flexibly for the share of jobs facilitating work from home by using the interaction of the share with week fixed effects. The coefficient estimates for the exposure to U.S. social distancing are very similar to those in Table 3, indicating that migrants are either not sorting into U.S. regions with similar ability to work from home, or that sorting is not influencing the effects of U.S. social distancing on Mexican social distancing.

Table 20: Results robust to flexibly controlling for the share of people whose job is workable at home

Outcome: Mexico social distancing	(1)	(2)	(3)	(4)
Variable for interaction:	Matching method 1		Matching method 2	
Workable at home share	Unweighted	Weighted	Unweighted	Weighted
Exposure to U.S. social distancing	0.043*** (0.005)	0.043*** (0.005)	0.043*** (0.005)	0.043*** (0.005)
Week 10 Interaction	0.003 (0.086)	0.006 (0.079)	-0.002 (0.079)	0.002 (0.074)
Week 11 Interaction	0.159* (0.084)	0.169** (0.076)	0.137* (0.076)	0.150** (0.071)
Week 12 Interaction	0.299*** (0.077)	0.295*** (0.070)	0.273*** (0.070)	0.273*** (0.065)
Week 13 Interaction	0.307*** (0.071)	0.279*** (0.065)	0.287*** (0.065)	0.265*** (0.061)
Week 14 Interaction	0.305*** (0.071)	0.276*** (0.065)	0.286*** (0.065)	0.263*** (0.060)
Week 15 Interaction	0.358*** (0.071)	0.325*** (0.065)	0.328*** (0.065)	0.304*** (0.060)
Week 16 Interaction	0.348*** (0.070)	0.302*** (0.064)	0.318*** (0.064)	0.283*** (0.060)
Week 17 Interaction	0.420*** (0.072)	0.349*** (0.065)	0.381*** (0.065)	0.326*** (0.061)
Week 18 Interaction	0.438*** (0.072)	0.362*** (0.065)	0.400*** (0.065)	0.340*** (0.061)
Week 19 Interaction	0.493*** (0.076)	0.404*** (0.069)	0.460*** (0.069)	0.387*** (0.064)
Week 20 Interaction	0.494*** (0.075)	0.414*** (0.069)	0.460*** (0.069)	0.395*** (0.064)
Week 21 Interaction	0.479*** (0.077)	0.391*** (0.070)	0.442*** (0.070)	0.371*** (0.065)
Observations	10,025	10,025	10,025	10,025
R-squared	0.915	0.915	0.915	0.915

Note: Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . *Municipio* fixed effects are controlled in all Columns. The workable-at-home measure is the mean of workable at home job shares using the industry code of working age population (aged 16–65) in a *municipio* and the industry-level workable at home job shares. The workable-at-home share for “Wholesale and retail trade” in Mexico is calculated as the mean of the workable-at-home shares for “Wholesale trade” and “Retail trade” in the U.S. (Table 19). In Columns (1)–(2), “Other services” and “Private household services” in Mexico are matched to “Other services (except public administration)” in the U.S., while in Columns (3)–(4), “Other services” and “Private household services” in Mexico are matched to “Other services (except public administration)”, “Professional, scientific, and technical services”, “Management of companies and enterprises”, “Information”, “Administrative and support and waste management and remediation services”, “Arts, entertainment, and recreation” in the U.S. (all unmatched service items in Table 19). Columns (1) and (3) use the unweighted shares, and Columns (2) and (4) use the weighted by wage shares.

Table 21 shows the heterogeneous effect of the exposure to U.S. social distancing with respect to the workable-at-home shares. We find that Mexican regions with higher workable-at-home job shares responding more strongly to U.S. social distancing. This is consistent with the heterogeneous effects found in Table 23, since as shown in Dingel and Neiman [2020], higher income is associated with higher shares of workable-at-home jobs (at the country level).

Table 21: *Municipios* with higher workable-at-home shares respond more strongly to U.S. social distancing

Outcome: Mexico social dist.	(1)	(2)	(3)	(4)
	Matching method 1		Matching method 2	
	Unweighted	Weighted	Unweighted	Weighted
Exposure to U.S. social distancing	0.025*** (0.006)	0.023*** (0.006)	0.025*** (0.006)	0.022*** (0.006)
Interaction with workable at home shares	0.080*** (0.010)	0.070*** (0.009)	0.074*** (0.009)	0.066*** (0.009)
Constant	0.215*** (0.000)	0.215*** (0.000)	0.215*** (0.000)	0.215*** (0.000)
Mean (s.d.) of workable at home shares	0.25 (0.04)	0.33 (0.05)	0.28 (0.05)	0.34 (0.05)
$\hat{\delta}$	7%	8%	8%	7%
Observations	10,025	10,025	10,025	10,025
R-squared	0.914	0.914	0.914	0.914

Note: Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . Week fixed effects and *municipio* fixed effects are included in all columns. Each column replicates the regressions in Column (1) of Table 3 and adds the interaction of a *municipio*'s workable-at-home job share. Similar as in Table 20, Columns (1) and (2) match "Other services" and "Private household services" in Mexico to "Other services (except public administration)" in the U.S., while Columns (3) and (4) match these two Mexican industries to the average of the unmatched service industries in Table 19 Panel B. Columns (1) and (3) use the unweighted shares, and Columns (2) and (4) use the weighted by wage shares.

C.7 Robustness of the main results when including all municipios

Table 22 replicates Table 4, but includes *municipios* with no cases. The results are similar in magnitude and significance to those in our main analysis, indicating that dropping *municipios* with no cases does not substantially affect the results.

Table 22: The results in Table 4 hold when all municipios are included

Variable for interaction	(1) pop. density	(2) % urban	(3) % age 16-65	(4) years edu.	(5) log income	(6) % employed
Exposure to U.S. social distancing	0.026*** (0.004)	0.025*** (0.004)	0.031*** (0.004)	0.028*** (0.004)	0.034*** (0.004)	0.028*** (0.004)
Week 10 Interaction	0.002 (0.002)	0.010 (0.011)	0.119 (0.083)	0.002 (0.002)	0.007 (0.008)	0.034 (0.038)
Week 11 Interaction	0.004*** (0.001)	0.031*** (0.011)	0.266*** (0.083)	0.006*** (0.002)	0.024*** (0.008)	0.088** (0.037)
Week 12 Interaction	0.007*** (0.001)	0.059*** (0.010)	0.517*** (0.079)	0.013*** (0.002)	0.048*** (0.008)	0.205*** (0.035)
Week 13 Interaction	0.010*** (0.001)	0.041*** (0.009)	0.547*** (0.071)	0.012*** (0.002)	0.033*** (0.007)	0.179*** (0.032)
Week 14 Interaction	0.011*** (0.001)	0.041*** (0.009)	0.551*** (0.070)	0.011*** (0.002)	0.025*** (0.007)	0.185*** (0.032)
Week 15 Interaction	0.012*** (0.001)	0.052*** (0.009)	0.617*** (0.070)	0.014*** (0.002)	0.041*** (0.007)	0.207*** (0.031)
Week 16 Interaction	0.011*** (0.001)	0.045*** (0.009)	0.570*** (0.070)	0.013*** (0.002)	0.029*** (0.007)	0.182*** (0.032)
Week 17 Interaction	0.013*** (0.001)	0.045*** (0.009)	0.630*** (0.070)	0.014*** (0.002)	0.033*** (0.007)	0.190*** (0.031)
Week 18 Interaction	0.014*** (0.001)	0.039*** (0.009)	0.632*** (0.071)	0.014*** (0.002)	0.034*** (0.007)	0.181*** (0.032)
Week 19 Interaction	0.016*** (0.001)	0.041*** (0.010)	0.822*** (0.075)	0.017*** (0.002)	0.036*** (0.007)	0.231*** (0.033)
Week 20 Interaction	0.016*** (0.001)	0.046*** (0.010)	0.812*** (0.074)	0.017*** (0.002)	0.033*** (0.007)	0.226*** (0.033)
Week 21 Interaction	0.016*** (0.001)	0.045*** (0.010)	0.815*** (0.077)	0.016*** (0.002)	0.032*** (0.008)	0.211*** (0.034)
Observations	13,036	12,841	13,036	13,010	13,010	13,036
R-squared	0.910	0.907	0.911	0.910	0.908	0.909

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Municipio fixed effects are controlled in all columns. Each column replicates the regression in Table 3 and adds the interaction of a city characteristic with week fixed effects. Week 9 is the baseline week. The sample is the Week-9-to-21 panel of all municipios.

Covid Economics 54, 29 October 2020: 62-120

Table 23 replicates Table 7, evaluating heterogeneity in the effects of exposure to U.S. social distancing based on the characteristics of Mexican *municipios*, but includes *municipios* with no cases. The results are similar in magnitude and significance to those in our main analysis, indicating that dropping *municipios* with no cases does not substantially affect the results.

Table 23: The results in Table 7 hold when all *municipios* are included

Outcome: Mexico social dist.	(1)	(2)	(3)	(4)	(5)	(6)
Exposure to US social dist.	0.028*** (0.004)	0.021*** (0.004)	-0.043*** (0.007)	0.007 (0.005)	-0.028*** (0.009)	0.009* (0.005)
Interact: population density	0.002*** (0.000)					
Interact: share urban		0.010*** (0.001)				
Interact: aged 16-65 share			0.124*** (0.010)			
Interact: yrs of schooling				0.003*** (0.000)		
Interact: log income					0.007*** (0.001)	
Interact: % employed						0.044*** (0.004)
Constant	0.209*** (0.000)	0.207*** (0.000)	0.209*** (0.000)	0.209*** (0.000)	0.209*** (0.000)	0.209*** (0.000)
Mean (s.d.) of the interaction $\hat{\delta}$	0.45 (1.6) 11%	0.56 (0.27) 10%	0.61 (0.04) 15%	8.3 (1.4) 13%	8.4 (0.40) 9%	0.50 (0.09) 13%
Observations	13,036	12,841	13,036	13,010	13,010	13,036
R-squared	0.908	0.907	0.909	0.908	0.908	0.908

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Week fixed effects and *municipio* fixed effects are included in all columns. Each column replicates the regression in Table 3 and adds the interaction of a *municipio* characteristic with the exposure to U.S. social distancing. The sample is the Week-9-to-21 panel of all *municipios*.

Table 24 replicates Table 8, evaluating heterogeneity in the effects of exposure to U.S. social distancing based on the characteristics of U.S. counties, but includes *municipios* with no cases. The results are similar in magnitude and significance to those in our main analysis, indicating that dropping *municipios* with no cases does not substantially affect the results.

Table 24: Results in Table 8 hold when all *municipios* are included

Outcome: Mexico soc dist.	(1)	(2)	(3)	(4)	(5)	(6)
Exposure to U.S. soc dist.	0.032*** (0.005)	0.030*** (0.004)	-0.049** (0.022)	-0.027* (0.016)	0.022 (0.067)	0.034 (0.050)
Interact: % Hispanic	-0.002 (0.004)					
Interact: % Mexican		0.007* (0.004)				
Interact: Hispanic education			0.007*** (0.002)			
Interact: education				0.004*** (0.001)		
Interact: log Hispanic income					0.001 (0.006)	
Interact: log income						-0.000 (0.004)
Constant	0.209*** (0.000)	0.209*** (0.000)	0.209*** (0.000)	0.209*** (0.000)	0.209*** (0.000)	0.209*** (0.000)
Mean (s.d.) of the interaction $\delta$	0.29 (0.08)	0.22 (0.08) 2%	10.3 (0.21) 6%	13.1 (0.29) 5%	10.9 (0.06)	11.2 (0.08)
Observations	13,036	13,036	13,036	13,036	13,036	13,036
R-squared	0.91	0.91	0.91	0.91	0.91	0.91

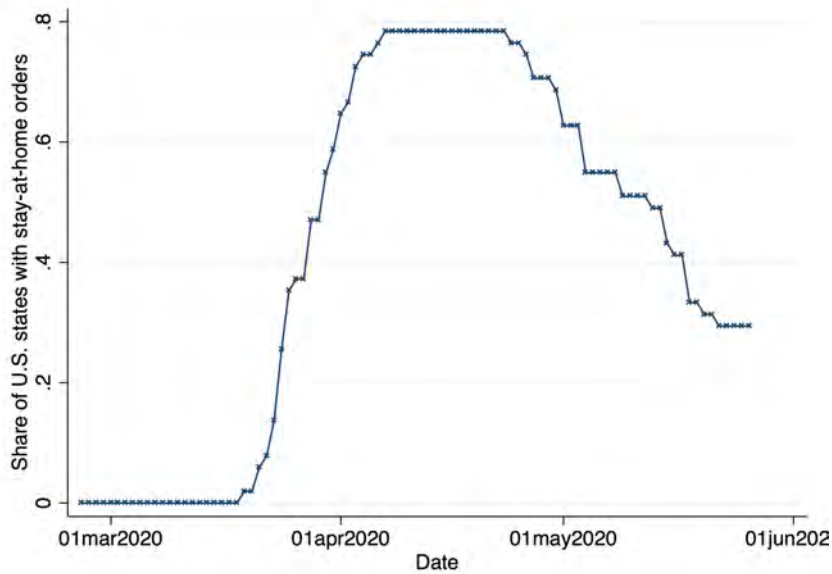
Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. Week fixed effects and *municipio* fixed effects are included in all columns. Each column replicates the regression in Table 3 and adds the interaction of a *municipio* characteristic with the exposure to U.S. social distancing. The *municipio* characteristics are characterized by the type of U.S. counties they are connected to. The sample is the Week-9-to-21 panel of all *municipios*.

### D Additional results for instrumental-variables analysis using stay-at-home orders

This section presents additional supporting evidence on the validity of the instrumental variables results from Section 4. It shows the first-stage residual plot and the reduced form results of the instrumental variables analysis using U.S. stay-at-home orders as an instrument for observed U.S. social distancing.

Figure 14 shows the proportion of U.S. states imposing stay-at-home orders since the beginning of the pandemic. We use indicators for U.S. state-level stay-at-home orders as an instrument for U.S. social distancing, as defined in equation 4.

Figure 14: The dynamic of stay-at-home orders in the U.S.



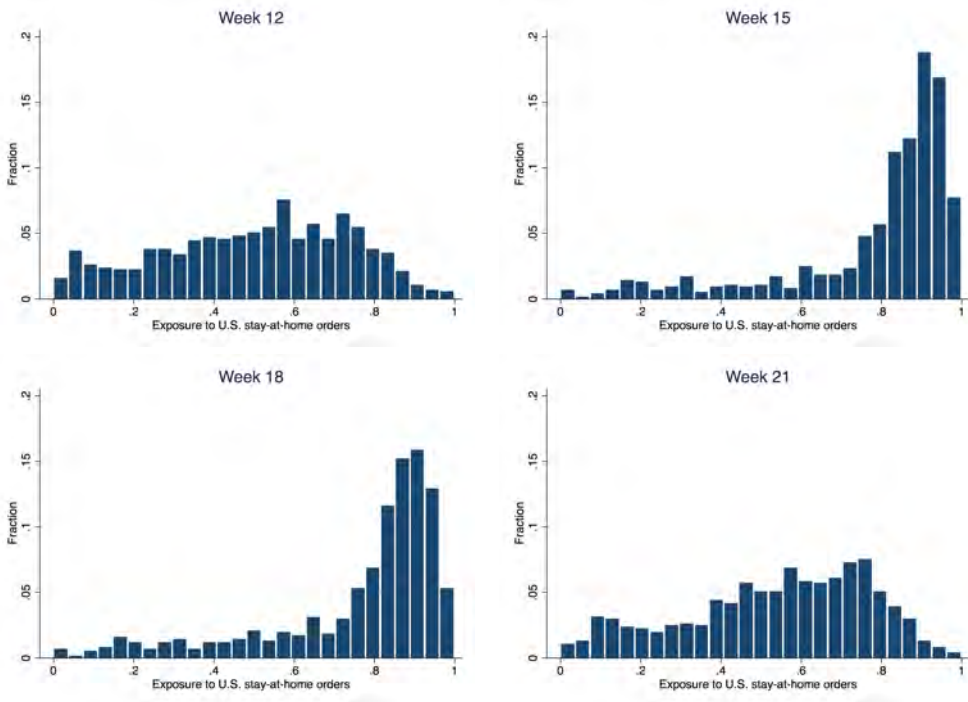
Note: This figure shows the share of U.S. states that have stay-at-home orders on a particular date. All 50 states and the District of Columbia are included. Stay-at-home or shelter-in-place orders only include directives and orders, but not guidance, and the order must apply to the entire states. According to this definition, 11 states never enacted the order, including: Arkansas, Connecticut, Iowa, Kentucky, Nebraska, North Dakota, Oklahoma, South Dakota, Texas, Utah, and Wyoming. Similarly, the end of the order must also apply to the entire state. See details of the definition at Raifman et al. [2020].

Covid Economics 54, 29 October 2020: 62-120



Figure 15 shows the distribution of the exposure to U.S. stay-at-home orders faced by Mexican *municipios* in Week 12, 15, 18, and 21. There is a great deal of variation in exposure to U.S. stay-at-home orders across *municipios* and over time.

Figure 15: Variation in exposure to U.S. stay-at-home orders



Note: This figure shows the distribution of Mexican municipios' exposure to stay-at-home orders in the United States, in Week 12, 15, 18, and 21. The sample is restricted to municipios that have at least one covid case by the end of Week 21 and have non-missing measures of social distancing in the corresponding week.

Table 25 performs a similar analysis to the one presented in Table 5, but at the U.S. county level rather than the *municipio* level, demonstrating the determinants of social distancing behavior in the U.S.

Table 25: Predicting the social distancing in the U.S., Week 9 to Week 21, around 2,570 counties

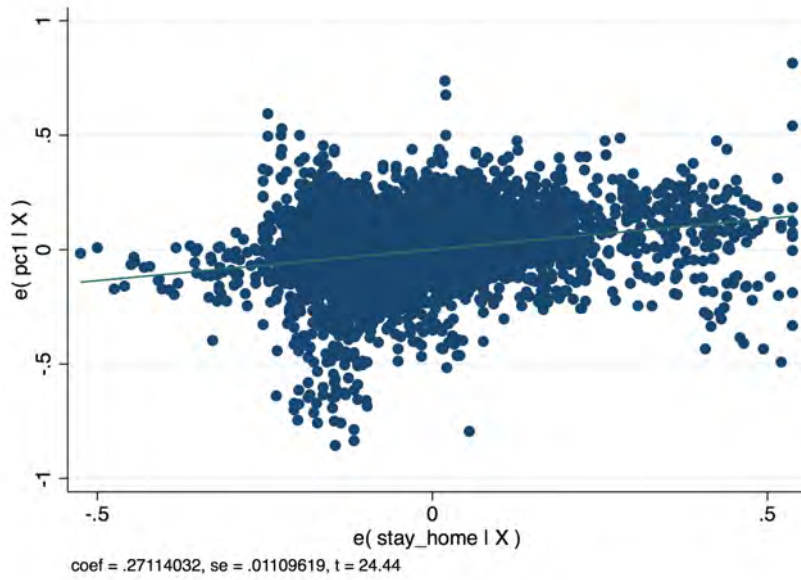
	(1)	(2)	(3)	(4)	(5)	(6)
Outcome: U.S. soc distancing	Facebook measure			Unacast measure		
Shelter-in-place order in the state, U.S.	0.14*** (0.001)	0.13*** (0.001)	0.13*** (0.001)	0.17*** (0.01)	0.16*** (0.001)	0.16*** (0.001)
Employment share controls	No	Yes	Yes	No	Yes	Yes
Commuting share controls	No	No	Yes	No	No	Yes
Observations	33,411	33,411	33,411	33,411	33,411	33,411
R-squared	0.27	0.33	0.33	0.26	0.32	0.33

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. The thirteen industries include: agriculture, construction, manufacturing, wholesale, retail, transportation, information, finance, professional, education, arts, public administration, and other services. The means of transportation include: (1) car, truck, or van; (2) public (excluding taxicab); (3) taxicab; (4) motorcycle; (5) bicycle; (6) walked; and (7) worked at home.

Covid Economics 54, 29 October 2020: 62-120

Figure 16 shows the first-stage residual plot corresponding to Column (1) of Table 5, showing a strong positive relationship between exposure to U.S. social distancing (1) and the stay-at-home exposure instrument (4).

Figure 16: First-stage residual plot of Table 5 Column (1)



Note: This figure is the residual plot of Table 5 Column (1).

Table 26 presents reduced-form regressions using the stay-at-home exposure instrument (4), showing a positive relationship between exposure to U.S. stay-at-home orders and Mexican social distancing.

Table 26: Reduced form evidence: exposure to U.S. stay-at-home orders positively affected Mexican social distancing

Outcome: Mexican social distancing	(1)	(2)	(3)	(4)
Exposure to U.S. stay-at-home orders	0.013** (0.005)	0.013** (0.005)	0.011** (0.005)	0.011** (0.005)
Exposure to log U.S. cum. cases		-0.010*** (0.003)		-0.012*** (0.003)
Log cum. cases Mexican muni.			0.016*** (0.001)	0.016*** (0.001)
Observations	10,051	10,051	10,051	10,051
R squared	0.91	0.91	0.92	0.92

Note: Robust standard errors in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ . Week fixed effects and *municipio* fixed effects are controlled in all columns. The mean (s.d.) of the exposure to U.S. stay-at-home orders is 0.54 (0.36), and the mean (s.d.) of Mexican social distancing is 0.21 (0.15). The sample is the Week-9-to-21 panel of *municipios* with at least one Covid-19 case by the end of Week 21.

Table 27 replicates Table 5 in our instrumental variables analysis including controls for stay-at-home orders imposed in Mexican states that differ from those imposed by the federal government. The tables show that the first stage results are robust to the inclusion of these controls.

Table 27: First stage: *municipios* with larger exposure to U.S. stay-at-home policies were also more exposed to U.S. social distancing

Outcome: Exposure to U.S. social distancing	(1)	(2)	(3)	(4)
Exposure to U.S. stay-at-home orders	0.274*** (0.017)	0.274*** (0.016)	0.274*** (0.017)	0.274*** (0.016)
Exposure to log U.S. cumulative cases		0.051*** (0.014)		0.050*** (0.014)
Log cum. cases Mexican muni.			0.006*** (0.002)	0.006*** (0.002)
Mexico state-level stay-at-home orders	-0.039*** (0.004)	-0.044*** (0.004)	-0.040*** (0.004)	0.006*** (0.004)
Constant	-0.157*** (0.009)	-0.449*** (0.078)	-0.165*** (0.009)	-0.452*** (0.078)
Observations	10,051	10,051	10,051	10,051
R-squared	0.993	0.993	0.993	0.993
First-stage F-statistic	614	622	613	621

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. This table replicates Table 5 by including controls for Mexican state-level stay-at-home orders. Week fixed effects and *municipio* fixed effects are included in all columns. The mean (s.d.) of exposure to U.S. social distancing is -0.02 (1.4), and the mean (s.d.) of the exposure to U.S. stay-at-home orders is 0.54 (0.36). The sample is the Week-9-to-21 panel of *municipios* with at least one Covid-19 case by the end of Week 21.

Table 28 replicates Table 6 including controls for stay-at-home orders imposed in Mexican states that differ from those imposed by the federal government (See Appendix B). The results are nearly identical to those in Table 6.

Table 28: IV results confirm main findings in Table 3.

Outcome: Mexican social distancing	(1)	(2)	(3)	(4)
Exposure to U.S. social distancing	0.046** (0.019)	0.045** (0.019)	0.042** (0.019)	0.0421** (0.019)
Exposure to log U.S. cum. cases		-0.012*** (0.003)		-0.014*** (0.003)
Log cum. cases Mexican muni.			0.015*** (0.001)	0.016*** (0.001)
Mexico state-level stay-at-home orders	0.003 (0.002)	0.004** (0.002)	-0.001 (0.002)	0.001 (0.002)
Observations	10,051	10,051	10,051	10,051

Note: Robust standard errors in parentheses. \*\*\* p<0.01, \*\* p<0.05, \* p<0.1. This table replicates Table 6 by including controls for Mexican state-level stay-at-home orders. Week fixed effects and *municipio* fixed effects are included in all columns. The sample is the Week-9-to-21 panel of *municipios* with at least one Covid-19 case by the end of Week 21. The exposure to U.S. social distancing is instrumented with the exposure to U.S. stay-at-home orders in all columns.

# Urban flight seeded the COVID-19 pandemic across the United States

Joshua Coven,<sup>1</sup> Arpit Gupta<sup>2</sup> and Iris Yao<sup>3</sup>

Date submitted: 15 October 2020; Date accepted: 15 October 2020

*We document large-scale urban flight in the United States in the wake of the COVID-19 pandemic. Populations that flee are disproportionately younger, whiter, and wealthier. Regions that saw migrant influx experience greater subsequent COVID-19 case growth, suggesting that urban flight was a vector of disease spread. Urban residents fled to socially connected areas, consistent with the notion that individuals were sheltering with friends and family or in second homes. The association of migration and subsequent case growth persists when instrumenting for migration with social networks, pointing to a causal association.*

1 PhD Student, NYU Stern School of Business.

2 Assistant Professor, NYU Stern School of Business.

3 PhD Student, NYU Stern School of Business.

Copyright: Joshua Coven, Arpit Gupta and Iris Yao

## I INTRODUCTION

*“Rumors that cholera was moving west and not south from Canada could not stem the growing panic; mass exodus from the city had already begun. A hyperbolic and sarcastic observer remarked later that Sunday had seen ‘fifty thousand stout hearted’ New Yorkers scampering ‘away in steamboats, stages, carts, and wheelbarrows.’ ”*

— The Cholera Years: The United States in 1832, 1849, and 1866 by Charles E. Rosenberg

The density and international connections of cities foster human interaction and economic activity. These same associations, however, have historically made urban areas uniquely vulnerable to contagious disease. The role of urban proximity in pandemics has seen renewed attention in the context of novel coronavirus disease 2019 (COVID-19). However, the channels through which communities employ resources to mitigate their own risk, and the spillovers of these actions on broader community transmission, remain unclear.

This paper quantifies the extent of urban flight in response to COVID-19 in its initial phase in the United States, and documents how this migratory behavior seeded the pandemic in the rest of the country. We take advantage of mobile phone geolocation data, which allows for much higher frequency analysis than has been possible in prior studies of migratory behavior. We find that as much as 15–20% of the population of some high-income urban regions, such as Manhattan, fled in response to COVID-19. Regions that saw greater flight were generally richer, whiter, and younger than other areas, pointing to important disparities in the availability of migration as a risk-mitigating technique during the pandemic. We use Facebook friendship data to establish that migration was especially high between socially connected regions, consistent with the idea that urban flight led to sheltering with friends and family, or in second homes.

We then document the impact of this urban flight on increases in COVID-19 cases in the destination counties. Our instrumental variable strategy leverages the social connections between countries to causally identify the relationship between migratory flows and increased cases. We find that a one standard deviation increase in SCI-instrumented inflow is associated with a 0.5 standard deviation increase in case growth. Alternatively, an increase in inflows by an additional 100 residents raises local cases by about 20. The median county sees an average instrumented inflow of 120 residents, and over 10% of all counties see instrumented inflows of over 200 residents during our sample period. Our estimates are substantial and point to urban migration as an important vector of COVID-19 spread across the United States.

Our results are also relevant to both questions of local government public finance shocks, as well as the long-term future of cities. Because the population which left was disproportionately rich, their flight deprives cities of valuable tax revenue in the short-run. To the extent that urban migration remains persistent, cities may also face long-term challenges around budget shortfalls, real estate prices, and population size. While our results end in July 2020, and so we are limited in addressing these long-term questions, we quantify important urban disruptions in the wake of COVID-19.

Our results additionally have implications on the possible value of travel restrictions. Virtually every country has placed severe restrictions on international travel, against the initial advice of the World Health Organization. Many countries have additionally placed restrictions on regional travel within countries. These include complete bans on interstate travel, as between Victoria and New South Wales in Australia<sup>1</sup>, as well as requested isolation orders for out-of-state visitors, as prevailed in some U.S. states in our period, such as Rhode Island.<sup>2</sup> However, the potential value of these travel restrictions remains very unclear. Our results suggest that intra-US travel was a main vector spread, and therefore

<sup>1</sup>See <https://www.npr.org/2020/07/06/887659557/australia-closes-interstate-border-because-of-coronavirus-outbreak>.

<sup>2</sup>See Rhode Island Government official press release from April 7: <https://www.ri.gov/press/view/38091>.



the existing level of travel restrictions and self-suggested quarantine orders were insufficient to prevent out-of-state migration from impacting local spread. As a consequence, our work points to the value of potential policies related to travel restrictions in curbing spread.<sup>3</sup>

Our work is closely linked to a rapidly growing literature using mobile phone geolocation data to assess the spread of COVID-19. Most closely related is work by [Chiou and Tucker \(2020\)](#), which finds shelter-in-place effects vary by income. This paper differs by considering the role of leaving the city, connecting mobility with actual COVID exposure, and incorporating analysis of other demographic groups. [Glaeser, Gorback, and Redding \(2020\)](#) examines the change in mobility within regions, while we examine the flight response across areas. Other work ([Allcott, Boxell, Conway, Gentzkow, Thaler, and Yang, 2020](#); [Barrios and Hochberg, 2020](#); [Engle, Stromme, and Zhou, 2020](#); [Painter and Qiu, 2020](#); [Andersen, 2020](#)) has looked at political partisanship and COVID-19 responses.

Prior work, such as [Athey, Ferguson, Gentzkow, and Schmidt \(2019\)](#), [Chen, Haggag, Pope, and Rohla \(2019\)](#), and [Chen and Rohla \(2018\)](#), has used mobile phone geolocation data to examine segregation, racial disparities in voting waiting times, and partisanship. Another use of individual ping-level geolocation data includes [Chen, Chevalier, and Long \(2020\)](#), who examine nursing home networks in the wake of COVID-19. [Holtz, Zhao, Benzell, Cao, Rahimian, Yang, Allen, Collis, Moehring, Sowrirajan, Ghosh, Zhang, Dhillon, Nicolaides, Eckles, and Aral \(2020\)](#) also uses Facebook data to shows similarity in social distancing responses among regions connected through friendship links. Our work highlights a direct migration linkage between socially connected regions.

Our work also relates to a literature that uses geotyping of strains to establish both international chains of transmission in Europe ([Pybus, Rambaut, COG-UK-Consortium,](#)

---

<sup>3</sup>[Chandrasekhar, Goldsmith-Pinkham, Jackson, and Thau \(2020\)](#) also highlights the importance of regional spillovers and network interactions. We contribute to this work by quantifying the role of the migration channel in contributing to case growth at an early stage of the COVID-19 pandemic. [Lee, Mahmud, Morduch, Ravindran, and Shonchoy \(2020\)](#) also finds an association of migration and cases in the context of COVID-19 in South Asia.

et al., 2020) as well as domestic transmission in the United States (Fauver, Petrone, Hodcroft, Shioda, Ehrlich, Watts, Vogels, Brito, Alpert, Muyombwe, et al., 2020). While this work establishes the strain similarities across regions in the United States points to a role for domestic transmission, we provide additional evidence on the physical migration behind this regional transmission.

## II DATA AND SPECIFICATION

### *II.A Data*

Mobile location data was sourced from VenPath—a holistic global provider of compliant smartphone data. We obtain unique data on smart phone GPS signals. Our data provider aggregates information from approximately 120 million smart phone users across the United States. GPS data were combined across applications for a given user to produce pings corresponding to time stamp-location pairs. Ping data include both background pings (location data provided while the application is running in the background) and foreground pings (activated while users are actively using the application). Our sample period covers the period February 1 to July 13, 2020.

We supplement our mobility data with county-level coronavirus case counts from the COVID-19 Data Repository by the Center for Systems Science and Engineering at Johns Hopkins University.<sup>4</sup> We join this with a county-to-county Social Connectivity Index (SCI) measures from Facebook as discussed in Bailey, Cao, Kuchler, Stroebel, and Wong (2018) and applied in reference to COVID-19 in Kuchler, Russel, and Stroebel (2020). We also include demographic data from the Census ACS, and urban-rural county classifications from the National Center for Health Statistics.

<sup>4</sup>Drawn from <https://coronavirus.jhu.edu/>. We also incorporate nursing home data from the Centers for Medicare and Medicaid Services and <https://www.cms.gov/>.

Table I summarizes some key statistics—in particular, we note that the distribution of caseloads across counties is highly skewed, with several key counties exhibiting extremely high growth in new cases while approximately half of the counties in our sample had fewer than 1 new case per day. Secondly, the difference between the SCI and its inflow-weighted counterpart shows that net inflows between counties are occurring most frequently between highly connected counties, giving some preliminary evidence in support of the SCI as a relevant instrument for county-county flows.

TABLE I: Summary Statistics

	1st Quartile	Median	Mean	3rd Quartile
Daily New Cases	0.00	0.00	8.19	2.00
Daily New Cases Per 1000 People	0.00	0.00	0.04	0.04
Net Inflow	31.00	67.00	152.73	153.00
SCI-Instrumented Inflow	107.18	123.70	152.67	153.05
Social Connectedness Index	31.68	63.07	111.19	121.10
Social Connectedness Index (inflow-weighted)	62.64	171.03	757.78	515.25
Distance	311.35	453.03	520.15	635.14
Distance (inflow-weighted)	182.20	250.02	301.96	349.29
Number of Nursing Homes	2.00	5.00	7.67	8.00

*Notes:* Summary statistics across all counties in our sample, with caseload and inflow data between 3/1/2020 and 7/13/2020. SCI and distance are calculated as the raw arithmetic average between the destination county and other counties from which there is a positive inflow. The inflow-weighted counterpart of these measures are calculated by using inflows from a particular county over total inflows as the weight in the weighted average.

We isolate the migration behavior of users in the US by identifying each user's modal census tract each night (6pm - 8am) if they ping in it three or more times. We do this for each night in a month. If a user has the same tract as their modal night tract on at least five nights in that tract, we define it as their "home tract:" the Census tract that they spend the most time in during night hours. We repeat this process each month from February to June to analyze mobility from March to July. We use only one month of data at a time to identify residents' home tracts. We then analyze their data in the month immediately following the month that was used to identify their home locations. The resulting sample includes a population of 9–11 million unique users per month for our base analysis across the United States. In New York City, we observe a 0.89 correlation between the population

of each ZIP Code and our observed mobile phone population in that area.

To calculate county to county flow or ZIP Code to ZIP Code flow, we observe the count of users spending the night in a given census tract and aggregate up to the county or ZIP level each date. We aggregate tracts to counties, and link tracts to ZIP codes using a crosswalk provided by the Department of Housing and Urban Development.<sup>5</sup> In cases where tracts map to multiple ZIP Codes, we select the ZIP Code with the highest number of residents. We aggregate resident counts to the home geography, current geography, and date level to see where people from a given geography are spending the night on each date. The resulting measures estimate flows from home geographies to new geographies on each date.

## II.B Empirical Specification

Our core empirical specifications examine the determinants and consequences of urban flight in the context of the COVID-19 pandemic in the United States. After aggregating mobile phone migration data into county-day information, we sum, for a given day  $t$ , all net inflows into county  $i$ . Our core OLS specification measures case growth in a destination county as a function of gross inflow into that county:

$$\begin{aligned} \text{New Cases}_{i,t} = & \beta_0 \cdot \text{Inflow}_{i,t} + \beta_1 \cdot \mathbb{1}(\text{High Cases in Originating Counties})_{i,t} + \\ & + \beta_2 \cdot \mathbb{1}(\text{Far})_i + \beta_3 \cdot \text{Inflow}_{i,t} \times \mathbb{1}(\text{Far})_i + \gamma_1 \cdot X_i + \delta_{s,m,p} + \varepsilon_{i,t}. \end{aligned} \quad (1)$$

We are primarily interested in the  $\beta_0$  coefficient, which measures the effect of inflows on cases. We measure cases both in levels as well as per capita. In addition, we test whether inflow from counties which experience high case counts and inflow from more distant counties have a differential impact on case growth. The indicator for high cases

<sup>5</sup>See: [https://www.huduser.gov/portal/datasets/usps\\_crosswalk.html](https://www.huduser.gov/portal/datasets/usps_crosswalk.html) for the crosswalk.

in originating counties equal to 1 for counties where the inflow-weighted cases from incoming counties fall within the top quartile in the current month. We construct  $\mathbb{1}(\text{Far})$  by assigning 1 to counties where the inflow-weighted distance is at least 500km (roughly equal to the distance between NYC and Pittsburgh).

We further include a number of county-level control variables in our regression specification to account for sources of coronavirus case growth orthogonal to inflows from outside counties. Controls include the distance between the home and destination counties, mean household income, population density, the NHS urban-rural classification, share of the population above 60 years of age, the share of essential workers, and the number of nursing homes. Finally, we sort counties into deciles by population and include state by month by population decile fixed effects to isolate the effect of migration from unobservable heterogeneities across city size.

While our specification includes a number of plausible control variables, an important potential identification concern with equation 1 is the endogenous nature of migration. If counties that receive higher domestic migration are also more likely to be susceptible to case growth for other reasons, a positive  $\beta_0$  may simply reflect spurious correlation, rather than measuring the causal effect of migration on COVID-19 case growth. We develop an identification strategy to address endogenous migration decisions based on Facebook connectivity. We draw on prior research, as mentioned in [Bailey, Cao, Kuchler, Stroebel, and Wong \(2018\)](#), that suggests social connectivity is a driver of migration decisions when measured at annual frequencies. Our analysis establishes that social connections explain the high-frequency pandemic-driven migration observed during COVID-19. To examine the relationship between migration and social connectivity, we first run a first stage regression of migration against social connectivity between regions:

$$\text{Inflow}_{i,t} = \gamma_0 \text{SCI}_{i,t} + \sum_{d=2}^{10} \gamma_d \cdot \text{Distance Decile}_{i,t} + \gamma_1 \cdot X_i + \delta_{s,m} + v_{i,t}. \quad (2)$$

The SCI between two counties or ZIP Codes  $i$  and  $j$  measures the strength of social connections between them, and is defined (as in [Kuchler, Russel, and Stroebel \(2020\)](#)) based on the number of users in two regions  $i$  and  $j$  as well as the friendship links between them:

$$SCI_{ij} = \frac{FB\_Connections_{ij}}{FB\_Users_i \times FB\_Users_j}. \quad (3)$$

Our primary specification examines social connectivity at the county level, which measures the total degree of connectivity across the country. We also examine connectivity between ZIP pairs. The coefficient  $\gamma_0$  measures the strength of our first stage—the predictiveness of social connectivity in a gravity regression on migration, controlling for the decile of physical distance and other factors.

In order to isolate the effect of county inflows due urban flight, we instrument county-county flows with the county SCI measure. The identifying assumption is that Facebook connections between county  $i$  and other counties does not correlate with the trajectory of case growth, *except* through the inflow of people into the county. Our instrumental variables specification first instruments for inflow using equation 2, and uses predicted inflow instead of realized inflow as a covariate in equation 2. We conduct our main analysis at the County level, where we have case data nation-wide, but are also able to establish the relationship between migration flows and SCI at the ZIP Code level.

### III RESULTS

#### *III.A Temporal and Spatial Patterns of Urban Flight*

We begin by highlighting the demographic characteristics of individuals who leave urban areas in the wake of the COVID-19 pandemic. Our initial focus is New York City, which

was the epicenter of the pandemic in the initial stage, and we emphasize the scale of the migration response. However, subsequent results analyze urban flight across the country.

In Panel A of Figure I, we show responses of individuals in leaving New York City by borough. We observe stark patterns in the response of individuals along this extensive margin: residents of Manhattan are substantially more likely to leave the city after the crisis, as are other wealthy parts of the city in Brooklyn. We find that as many as 10–15% of the population of Manhattan, formerly resident in the city in February, leaves the city by April 15. By contrast, residents in Queens—the epicenter of the COVID-19 pandemic in New York City—Brooklyn, and the Bronx are overwhelmingly more likely to stay in the city. This urban exodus continues through the end of our sample. By July 10, we observe about 17% of the population of Manhattan remains absent in the city.

These shifts in leaving the city, however, are concentrated in the higher-income Census tracts, suggesting that richer New York City residents were disproportionately able to take advantage of the option to flee the city and escape physical COVID-19 exposure in the city.

We confirm the role of income as a factor in explaining moves away from the city in Panel B of Figure I, which shows a heatmap of responses by tract and date among tracts with median household incomes in excess of \$100k. We find a large breakpoint in our sample in March 14, as reflected in the sharp changes in colors beginning on that date in a number of tracts, corresponding to a sharp rise in the increase of former New York City inhabitants leaving the city. This break comes just before Mayor Bill Di Blasio ordered schools, restaurants, bars, cafes, entertainment venues, and gyms in the city closed on March 16.<sup>6</sup> We observe very high flight behavior in the highest-income Census tracts after that date.

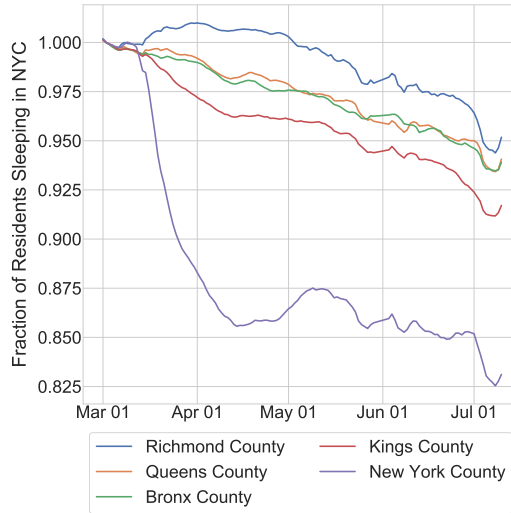
The fact that we observe migration response prior to the introduction of nonpharmaceutical interventions such as lockdowns is also consistent with Goolsbee and Syverson

---

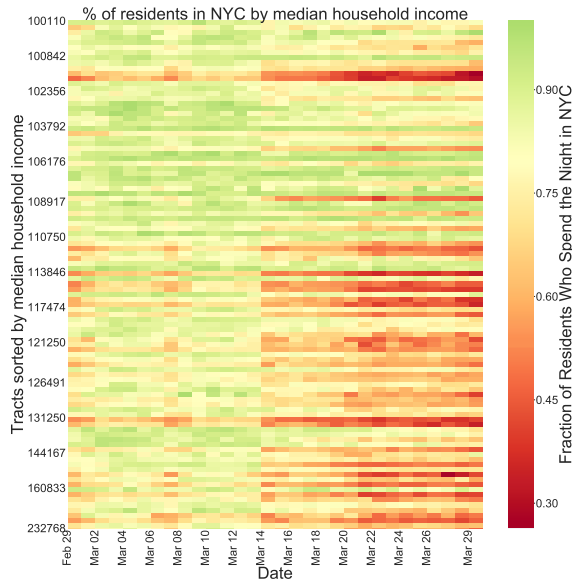
<sup>6</sup>See <https://www1.nyc.gov/assets/home/downloads/pdf/executive-orders/2020/eo-100.pdf>.

FIGURE I: Propensity to Leave NYC by Tract

Panel A: Propensity to Leave NYC by Borough



Panel B: Propensity to Leave NYC by ZIP Income, \$100k > tracts





(2020), and points to important behavioral responses separate from policy responses. Unlike other mobility changes, such as sheltering at home—migration entails possible externalities on the exposure of others which we explore.

We examine the spatial distribution of flight patterns in Figure II, which shows the fraction of residents who leave across six cities (New York, San Francisco, Los Angeles, Washington DC, and Boston). We measure the pre-existing urban population in each Zip Code, and plot the fraction that has left by March 29, 2020. New York City experiences extremely high flight concentrated in Manhattan, with several Zip Codes seeing over 50% of the resident population leaving. Flight is concentrated in the Downtown and Midtown regions, though we also observe extensive urban flight in the Upper West Side, Upper East Side, and the wealthier regions of Brooklyn. We also observe distinctive patterns of urban flight in San Francisco (concentrated in downtown regions) as well as Boston (high in Cambridge and downtown Boston). These maps suggest that large-scale urban flight in response to the COVID-19 pandemic, with responses concentrated in the richer parts of several major metropolitan areas.

### *III.B Demographic Associations of Urban Flight*

We examine the demographic associations of urban flight in Figure III, which focuses on the propensity to remain in our six city sample. We plot background demographic associations at the ZIP Code level against the fraction of ZIP Code population that stays. Background dots show all data points, while binscatter dots plot the average population within 25 quantiles. We find that the fraction of residents who remain in cities is strongly decreasing in tract income, decreasing in the fraction of the tract that is White, and the fraction of residents aged 18–45. These results are large in magnitude and statistically significant—moving from a tract at the bottom quartile of income to the top quartile raises the fraction of New York City residents who leave, for instance, by about 3 percent, or about the same as the unconditional average of the number of New York City residents

FIGURE II: Propensity to Leave Cities

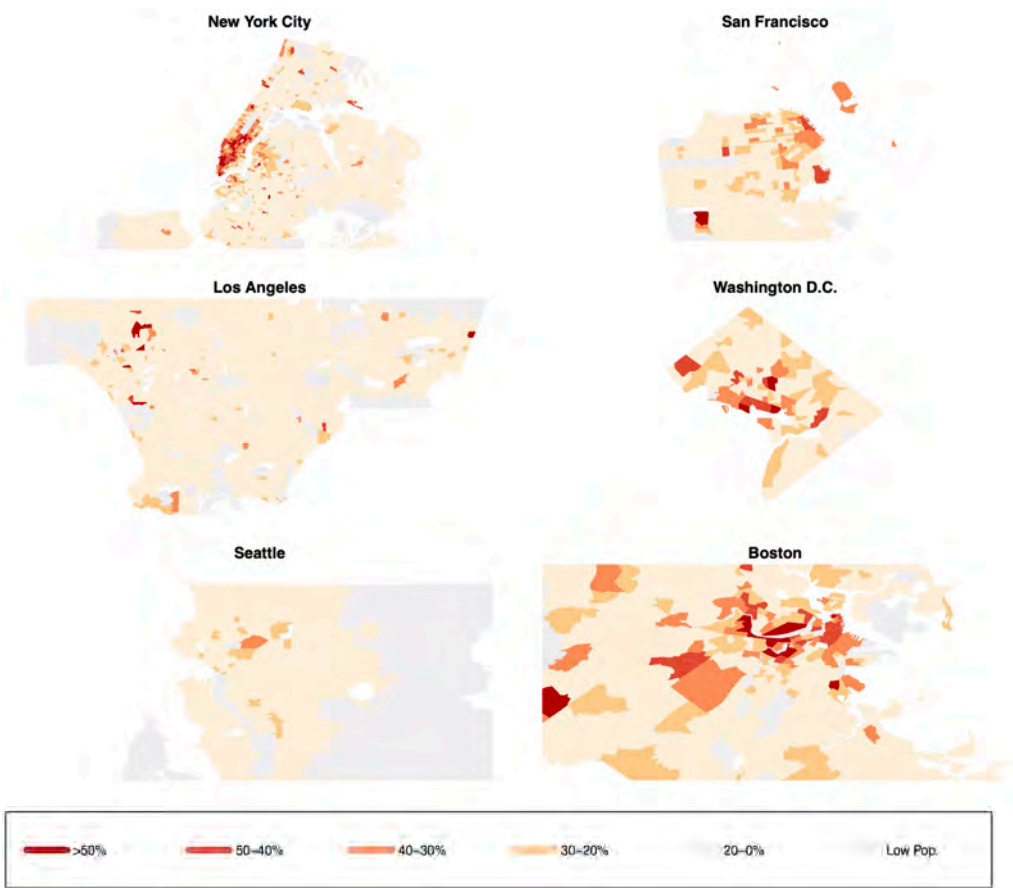




FIGURE III: Fraction Change in Out of Town Visits by Demographics


Panel A: Change in Out of Town Visits by ZIP Income



Panel B: Change in Out of Town Visits by Fraction White



Panel C: Change in Out of Town Visits by Fraction Aged 18-45



Notes: Panel A shows average household income against the change in the fraction of out of town visits. Panel B repeats the exercise for the fraction who are white, and Panel C for the fraction between the ages of 18-45. Data is tract-level and shown only for New York City, San Francisco, Los Angeles, Washington DC, Seattle, and Boston. Charts on the left represent the entire sample; charts on the right winsorize the fraction of out of town visits at the 5% level. All demographic variables are left un-winsorized, with the exception of the fraction between 18-45 in Panel C. Light gray points show all the data points; for each demographic variable, the data is divided into 25 quantiles and each dark blue dot represents the average fraction of the population sleeping at home and average demographic variable within each quantile. Income data is drawn from the IRS SOI Tax Statistics at <https://www.irs.gov/statistics/soi-tax-stats-individual-income-tax-statistics-zip-code-data-soi>, and demographic data are drawn from the ACS.

who leave. The income and demographic associations are even stronger in the panels on the right of this Figure, which winsorize the fraction of out of town visits at a 5% level.

### *III.C Associations of National Migration*

We first begin by descriptively analyzing the nation-wide nature of migration in the context of COVID-19. Figure IV documents the net flow and outflow of residents across counties in the United States as of the end of each month.<sup>7</sup> Map colors indicate the fraction of residents who left or entered the county, while the size of the circles indicate the size of the flow.

By the end of March, we document substantial flight out of New York City as well as several other metropolitan areas (including Boston, Los Angeles, San Francisco, and Phoenix). This inflow went to a mix of interior locations including rural areas across the country as well as Southern urban areas. Several cities in the Sunbelt, in particular Atlanta, Houston, Charlotte, and Austin saw substantial net in-migration during this period. Some other cities in North, such as Des Moines, Chicago, Detroit, Kansas City, and St. Louis, also saw substantial inflow. We also observe substantial inflow to numerous smaller counties in the vicinity of New York City, in the Hamptons and Hudson Valley. Broadly, the pattern of migration reflects flight away from the initial waves of the pandemic which hit the coasts more strongly, towards the national interior.


We observe continued urban flight from New York City, as well as additional flight away from Phoenix, Florida, and some Californian and Texan cities in April. By May, we observe substantial inflow into coastal regions for vacation purposes. Because our analysis ends in mid-July, we are limited in our ability to observe further migratory events associated with subsequent case increases across the Sunbelt regions of the United States over the period beginning in June. Our focus largely remains the substantial migratory

---

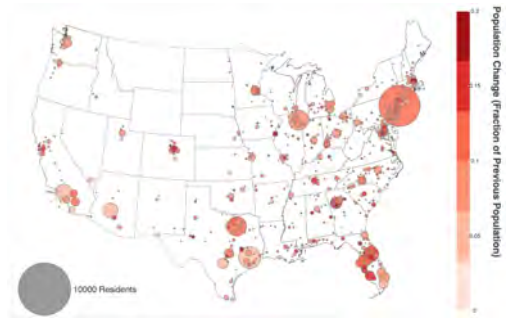
<sup>7</sup>Appendix Figure A1 examines inflow.

FIGURE IV: Nationwide Migration


Panel A: March Net Flow



Panel B: March Outflow




Panel C: April Net Flow



Panel D: April Outflow



Panel E: May Net Flow



Panel F: May Outflow



response in the wake of the first migration event around the beginning of the pandemic in March, 2020.

To further understand this unprecedented pattern of domestic migration, we begin by examining role of social connections in determining where individuals flee to in Figure V. In Panel A of this figure, we plot pre-existing social connections measured using the Facebook SCI variable and migration between March 1 and July 13 at the ZIP Code level across the entire United States. Background light points show a 1/100<sup>th</sup> random sample, while dark points show a binscatter of the 25th quintiles. We find a very strong positive association between higher social connectivity between ZIP Codes and migration over this period. We winsorize both SCI and inflows at a 1% level in the left plots, and at a 5% level in plots on the right and find similar results across both specifications. The strong relationship between inflow and SCI suggests that individuals with the ability to leave disproportionately went to areas where they had pre-existing social networks, and could take refuge with friends and family.


We then examine the relationship between migratory inflow and subsequent case growth in Panel B of V. In these plots, we move to the county-level in which we have COVID-19 case information. We plot the daily growth in cases against daily inflow for all counties over our entire sample period (March 1 to July 13). Binscatter dots show the 25 quantiles of the distribution, and suggest a strong relationship between migratory inflows and case growth. Our graphical evidence suggests that urban migration, directed towards socially connected regions, had spillover effects on destination regions in increasing COVID-19 case counts for destination counties.

### *III.D Impact of Urban Flight on Nationwide Case Growth*


Having established the nature of urban flight over the course of the COVID-19 pandemic, we turn next to our main analysis on the implications of this flight on destination regions.

FIGURE V: SCI-Related Mobility and Case Growth

Panel A: ZIP Inflows against SCI



Panel B: Inflows against Case Growth



Notes: Panel A show inflows per 1000 people between pairs of zip codes against SCI between the two zip codes. Light gray points show 1/100th of the entire sample. Panel B shows daily inflows against new cases at the county level. Light gray points show one-tenth of the entire sample. In the left column, SCI and daily inflow are winsorized at the 1% and 99% levels; in the right column, SCI and daily inflow are winsorized at the 5% and 95% levels. Dark blue dots are obtained by dividing the data into 25 quantiles and calculating the mean x-axis and y-axis value in each quantile.

Because coronavirus is a predominantly respiratory disease spread in close contact, direct exposure with individuals formerly living in high-risk areas is a plausible vector for the spread of the disease. While urban areas with international connections (particularly Seattle and New York City) appear to have been the initial hotspots for spread; the disease appears to have quickly spread from there to outlying regions through the travel patterns of affected individuals. We explore the idea that this pattern of extraordinary urban flight, by individuals perhaps avoiding the risk of contagion in urban areas, may have seeded the pandemic through the rest of the country.


To illustrate our key mechanism and hypothesis in graphical form, we begin by illustrating the relationship in Figure VI. This figure highlights the impact of increased migration on increased cases in the destination counties and focuses in flight form New York City. The city is central to our analysis both due to the size of its urban flight, as well the early presence of COVID-19 in the urban area. We first begin by separating our analysis into different urban categories based on attributes of the destination regions. Regions differ in their exposure to infectious disease spread based on urban form, so we analyze separately the impact of inflows of New York City residents contrasting the large and medium sized metros (NCHS category 1, 2 & 3 Panel A), compared with micropolitan and non-core areas (NCHS categories 4, 5 & 6, Panel B). Within each category, we compare case growth among counties that received the highest quartile of inflow of New Yorkers, compared with counties that saw the lowest quartile of inflow. Left panels show per capita cases in logs, while right panels show total cases.

We find sizable impacts of urban inflow from New York on COVID-19 cases across regions. In the largest urban areas, we find that case growth starts to increase for counties that receive high inflow from New York City beginning in March. We plot the raw seven day average in case growth for the counties receiving highest and lowest quartile of New York City inbound residents, and plot the difference between these regions in grey bars the background. Regions that saw high inbound migration see the greatest relative difference in cases in the beginning of April, a difference that declines over time. The timing of




FIGURE VI: March Inflows From NYC vs. Case Growth

Panel A: Large Central, Fringe, and Medium Metropolitan Areas



Panel B: Metropolitan and Non-Core Areas



Notes: Counties are split into quartiles based on total inflows from New York City during the month of March. The two charts on the left show log(total cases per capita) over time for counties in the fourth quartile and the first quartile of inflows. The two charts on the right show log(total cases). Urban classification based on the NCHS urban-rural classification scheme: large central metros, large fringe and medium metropolitan areas (categories 1, 2, & 3) and micropolitan and non-core areas (categories 4, 5 & 6).

case growth matches the period of influx of New York City residents with a lag, consistent with a channel of direct infection.

The difference between areas with high inflow of New York City residents and areas that see low inflow starts to decline in July. We show similar plots that plot case growth in Figure VII, which confirm that areas that initially saw higher New York City inflow see negative relative case growth into the summer. This result suggests that New York City inflow brought forward some cases which may have counterfactually been experienced later in the course of the pandemic. Urban flight would still be a quite important, even it only accelerates case growth, because of steady improvements in treatment and expanded supply of personal protective equipment such as masks over the pandemic, which lowered mortality rates among those infected later in the course of the pandemic.<sup>8</sup>

We observe that the impacts of New York City inflow are increasing in city size. Urban areas which saw higher inflow of New York City residents saw the greatest increase in new cases (a representative example would be Atlanta, which saw large case growth in the first wave of the pandemic). Large fringe and medium metropolitan areas also see substantial increase in early cases as a result of New York City inflow, but to a lesser degree than the largest urban areas. However, micropolitan and non-core areas see substantially weaker effects, which also turn negative around mid-May. Urban inflow could be most related to subsequent case growth in the largest urban areas due to greater realized population density and possibility for individuals to interact in the crowded, indoor environments which are most conducive to COVID-19 spread.

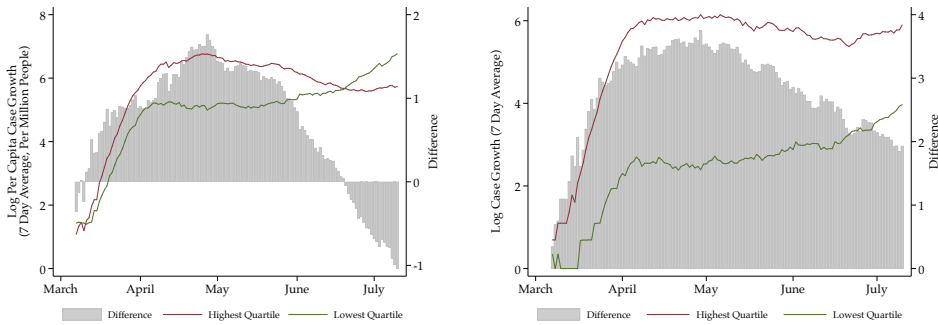
Overall, we find evidence across the cross-section and time-series consistent with urban migration being a substantial vector of spread for COVID-19 across the country. Areas with greater inflow of New York City residents see greater case growth exactly in the

---

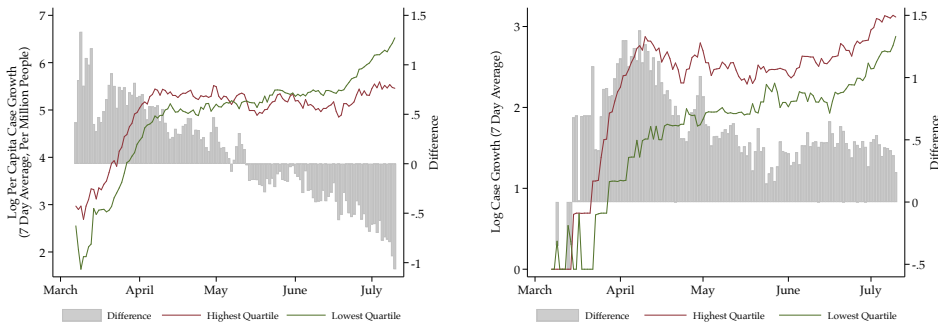
<sup>8</sup>See Horwitz, Jones, Cerfolio, Francois, Greco, Rudy, and Petrilli (2020) and Ciceri, Ruggeri, Lembo, Puglisi, Landoni, Zangrillo, and on behalf of the COVID-BioB Study Group (2020) on improved mortality and Gandhi and Rutherford (2020) who connects the improved mortality to increased mask adherence. Other explanations for greater mortality at the early stage of the pandemic including crowding at hospitals, learning-by-doing in medical care, and improved treatments over time.

FIGURE VII: March Inflows From NYC vs. Log Case Growth

Panel A: Large Central, Fringe, and Medium Metropolitan Areas



Panel B: Micropolitan and Non-Core Areas



Notes: Counties are split into quartiles based on total inflows from New York City during the month of March. The two charts on the left show median log(total cases per capita growth) over time for counties in the fourth quartile and the first quartile of inflows. The two charts on the right show log(total case growth). Urban classification based on the NCHS urban-rural classification scheme: large central metros, large fringe and medium metropolitan areas (categories 1, 2, & 3) and micropolitan and non-core areas (categories 4, 5 & 6). This figure differs from VI by considering case growth, rather than totals.

time period when it would be most expected. Our effects are largest for urban areas that receive New York City residents, and partially reverse over the course of the pandemic—suggesting that this migration brought forward additional cases. We next move on to regression analysis, which allows us expand our focus to migration across the entire United States, control for additional factors, and account for endogeneity in the migration decision by using the social connectivity of regions as an instrumental variable.

Our core regression specification in Table II follows our primary specification in equation 2 and covers inflows and case growth from March 1 to July 13. We first show that inflows lead to large increases in county caseloads, in our OLS specifications (columns 1–6). We find that for every additional 100 people who enter a county, case growth is increased by 0.8.<sup>9</sup> We also see that counties which experience inflows originating from counties which fall within the top quartile of case growth in a given month also see higher case growth on average. This large and statistically significant relationship parallels our graphical results in highlighting a role for case influx on cases. We also find evidence that migration from areas with higher case loads, and influx from areas further away lead to higher infection rates. These results are consistent with long-distance and interstate migration, especially from New York City, contributing to greater spread around the country.

---

<sup>9</sup>Coefficients in this table other than 4–6 are scaled by  $1 \times 10^3$ , and so correspond to the case increase resulting from an additional influx of 1000 people. Columns 4–6 are scaled by  $1 \times 10^6$ .

TABLE II: Impact of Migration on COVID-19 Cases

Panel A: COVID-19 Case Growth Against Inflow												
	OLS						IV					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Inflow	92.585*** (2.745)	12.022*** (2.556)	8.288*** (2.644)				267.821*** (9.468)	197.785*** (12.621)	194.605*** (12.551)			
Per Capita Inflow				-771.807*** (82.935)	-151.848* (90.625)	-199.698** (94.784)				-2969.855*** (931.814)	261.882*** (31.698)	197.772*** (20.711)
High Incoming Cases	4779.249*** (113.903)	4563.025*** (134.450)	4595.383*** (134.190)				5444.178*** (187.469)	6661.104*** (235.459)	7040.705*** (264.511)			
High Incoming Cases Per Capita				48.924*** (0.438)	49.372*** (0.443)	49.372*** (0.443)				-0.090* (0.050)	0.063*** (0.002)	0.060*** (0.002)
Far Indicator	3622.632*** (471.740)	2333.182*** (691.956)	-1362.127 (948.833)	0.427 (0.541)	-0.119 (1.114)	-1.102 (1.279)	3214.693*** (491.789)	1842.587** (792.740)	-18806.066*** (2961.841)	0.317*** (0.111)	0.008 (0.006)	-0.074*** (0.012)
Far Indicator × Inflow			20.448*** (4.690)									
Far Indicator × Per Capita Inflow						369.540 (247.614)						30.254*** (4.678)
Controls	N	Y	Y	N	Y	Y	N	Y	Y	N	Y	Y
State × Month × Population Decile FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
N	414,847	399,083	399,083	414,847	399,083	399,083	414,847	399,083	399,083	414,847	399,083	399,083

Panel B: First Stage Estimates of Inflow on SCI												
Weighted SCI							42947.436*** (443.427)	29688.738*** (456.984)	32749.326*** (526.249)	-0.002*** (0.001)	0.006*** (0.001)	0.003*** (0.001)
F							9,381	2,651	1,346	10	106	115

Columns 1-3 shows our main regression specification 1. Columns 4-6 repeat the exercise using inflow per capita as the outcome variable. Column 7-12 repeats the exercise for columns 1-6, where inflow is instrumented with the weighted SCI, as in 2. The sample period is March 1 through July 13. Standard errors are in parentheses, and \* denotes 10% significance, \*\* denotes 5% significance, \*\*\* denotes 1% significance. Note that all coefficients and standard errors in Panel A are scaled up by  $1 \times 10^3$ , with the exception of columns 4-6, where coefficients and standard errors are scaled up by  $1 \times 10^6$ . All coefficients and standard errors in Panel B are scaled up by  $1 \times 10^6$ .

We then instrument for migration patterns with Facebook friendship linkages (columns 6–10) to confirm a causal relationship between migration and subsequent case growth. Our IV estimates suggest statistically significant and economically substantial estimates. A one standard deviation increase in inflows instrumented for by social connections (roughly 116 people) causes a 0.5 standard deviation in daily case growth (roughly 23 cases). To put this into context, an increase in instrumented inflows by one standard deviation is akin to the difference in inflows experienced by Essex County in upstate New York (home to Lake Placid) as compared to Mercer County, NJ (home to the state capital). In this particular example, we also see that our back of the envelope calculations for mobility-linked case-growth is significant, as Mercer County's average case growth was 58 per day over our sample period, as compared to Essex County's daily case growth of 0.39. Our IV results suggest that migration accounts for slightly less than half of the difference in case growth across these two counties.

Our IV estimates are larger than our OLS estimates since raw inflow tends to overweight areas which exhibit lower case growth relative to the SCI-instrumented inflow. This appears likely because a substantial component of urban flight was motivated by fleeing to geographically remote areas where case growth is more likely to be low. These remote regions tend to be places where travellers do not have many existing social connections (e.g. renting a temporary property in upstate New York, or visiting a second home in a region where individuals do not know many others). Our IV, by contrast, identifies a LATE based on migration towards socially connected regions. Migration towards these areas, which is instead highlighting the flow based on the migration towards friends and families, appears to be more conducive to COVID-19 case transmission. Our IV, additionally, cleans up potential measurement error in our measurement of migration.

To provide a descriptive sense for the differing changes in sample and provide further support for why we see larger IV estimates, we provide some evidence in Figures VIII and IX. We show that instrumented inflow tends to be higher in regions with lower case

growth—these regions tend to have lower population density and a higher proportion of seasonal homes. In short, our OLS estimate underestimates the impact of inflow on case growth, because inflow is positively correlated with a variable which drives lower case growth (a characteristic which naturally makes these regions more desirable destinations for those fleeing urban areas). Additionally, we contrast the raw inflows against the instrumented inflows from the SCI measure. Our focus on the SCI measure effectively picks a different set of regions across the United States based on predicted inflows due to social connections, rather than realized migration activity. Raw inflows tend to pick up coastal and rural areas, in particular, relative to more connected urban areas in the SCI measure.

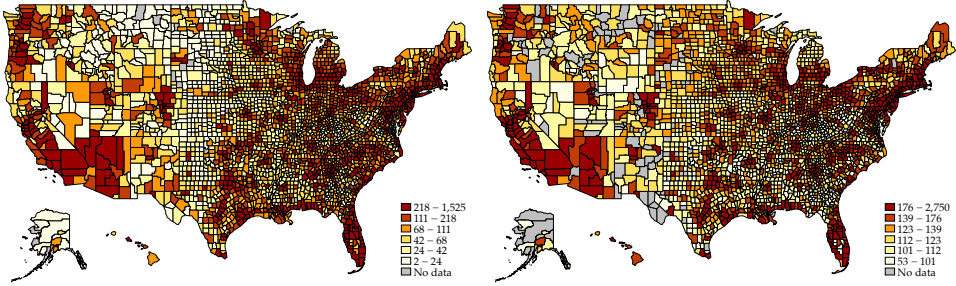
Finally, we find a strong and statistically significant first stage, and report the coefficient and F-statistic from the first stage regression 2 in Panel B for the IV regressions.

We perform several key robustness tests on our primary sample. In Appendix Table III, we include additional lags to account for past inflows. These additional lagged controls account for the incubation period between contracting the virus, exhibiting symptoms, and receiving a positive test result. These results do not suggest meaningfully different impacts from lagged inflows as compared to concurrent inflows.

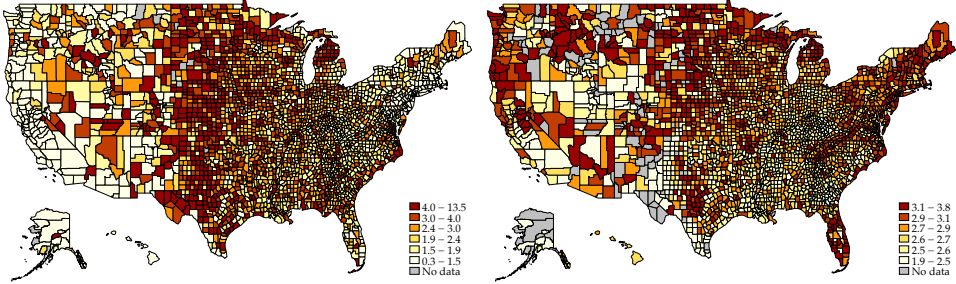
We also examine specifications that restrict on flight from New York City specifically. The large exodus from New York City, and the substantial case load in the city overall, make it a key focus of our analysis. In Table IV, we regress local cases against an indicator for counties in the top quartile of counties receiving inflows from NYC in March, which is most comparable with our graphical evidence. In Table V, we repeat our full analysis but subset to just inflows from New York City specifically. Both specifications find large and statistically significant effects. These results suggest that urban flight, and specifically the large urban exodus from New York City, was a central feature of the spread of the COVID-19 pandemic across the United States. An additional robustness table, Table VI, clusters standard errors at the Commuting Zone-level. While we lose significance in our per capita results, our core inflow measure remains statistically significant.

FIGURE VIII: Raw vs. SCI-Fitted Inflows

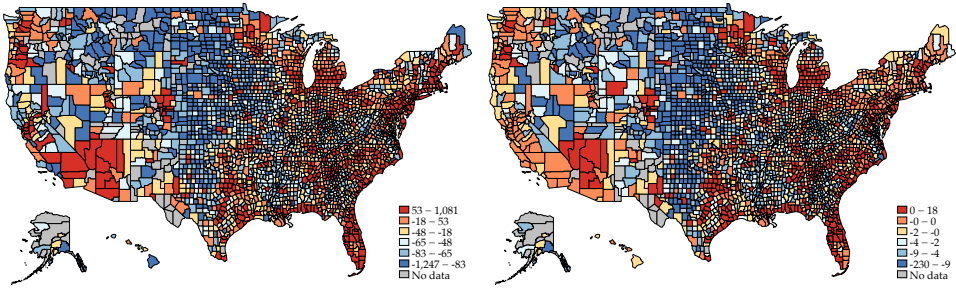
Panel A: Total Inflow



Panel B: Inflow Per Capita



Panel C: Difference Between Raw and SCI-Instrumented Inflows



Notes: Panels A & B show raw inflows and raw per capita inflows (left column) and SCI-fitted inflows and inflows per capita (right column). Darker red represent higher values; each map is colored by splitting the data into sextiles, resulting in different cut-off values for different colors vary across maps. Panel C shows the difference between raw and SCI-instrumented inflows (inflow - SCI-instrumented inflows); red counties represents regions where raw inflow is higher than instrumented inflow, and blue counties represent regions where raw inflow is lower than instrumented inflow.



TABLE III: Impact of Migration on COVID-19 Cases: Lagged Inflows


Panel A: COVID-19 Case Growth Against Lagged Inflows								
	OLS				IV			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Inflow (1 Week Lag)	23.299*** (2.721)				193.162*** (12.468)			
Inflow (1 Month Lag)		27.190*** (3.357)				197.702*** (13.490)		
Per Capita Inflow (1 Week lag)			722.257*** (101.890)				175.463*** (16.038)	
Per Capita Inflow (1 Month Lag)				521.679*** (122.954)				201.144*** (20.681)
High Incoming Cases (1 Week Lag)	3747.654*** (139.241)				4436.659*** (196.018)			
High Incoming Cases (1 Month Lag)		2224.591*** (175.153)				4113.848*** (301.818)		
High Incoming Cases Per Capita (1 Week Lag)			32.141*** (0.483)				0.026*** (0.001)	
High Incoming Cases Per Capita (1 Month Lag)				21.342*** (0.517)				0.028*** (0.002)
Far Indicator (1 Week Lag)	-2808.387*** (589.952)		-2.731*** (0.918)		-1.41e+04*** (2397.465)		-0.045*** (0.009)	
Far Indicator (1 Month Lag)		-1089.620 (662.542)		0.089 (1.016)		-1.48e+04*** (2785.041)		-0.060*** (0.011)
Far Indicator × Inflow (1 Week Lag)	27.131*** (5.066)				101.291*** (16.842)			
Far Indicator × Inflow (1 Month Lag)		23.243*** (5.981)				110.110*** (19.830)		
Far Indicator × Per Capita Inflow (1 Month Lag)			83.988 (264.241)				25.628*** (3.840)	
Far Indicator × Per Capita Inflow (1 Month Lag)				12.958 (313.248)				33.710*** (5.178)
Controls	Y	Y	Y	Y	Y	Y	Y	Y
State × Month × Population Decile FE	Y	Y	Y	Y	Y	Y	Y	Y
N	377,556	335,587	377,556	335,587	377,556	335,587	377,556	335,587

Panel B: First Stage Estimates of Lagged Inflows on SCI							
Weighted SCI				33012.669*** (516.184)	32419.067*** (534.261)	0.006*** (0.001)	0.004*** (0.001)
F				1,442	1,297	133	117

This table replicates table II, using 1-week and 1-month lagged explanatory variables. Columns 1-2 shows the results for total inflow and case growth; columns 3-4 shows results for per capita inflow and case growth. Columns 5-9 is the IV analogue of columns 1-4, where inflow and inflows per capita are instrumented with weighted SCI values, as in 2. The sample period is March 1 through July 13. Standard errors are in parentheses, and \* denotes 10% significance, \*\* denotes 5% significance, \*\*\* denotes 1% significance. Note that all coefficients and standard errors in Panel A are scaled up by  $1 \times 10^3$ , with the exception of columns 4-6, where coefficients and standard errors are scaled up by  $1 \times 10^6$ . All coefficients and standard errors in Panel B are scaled up by  $1 \times 10^6$ .

FIGURE IX: Differences in IV and OLS Estimates Due to Differences in Instrumented and Raw Inflow



Notes: The left binscatter shows case growth tends to be lower in regions with a greater proportion of seasonal homes. The right binscatter shows the unexplained portion of raw inflows (after partialling out the impact of the control variables we use in our baseline specification) tends to be higher than the unexplained portion of instrumented inflows, precisely in regions with more seasonal homes, and hence, regions with lower case growth. We point to seasonal homes as an imperfect measure of a possible omitted variable, which is positively correlated with the raw (not instrumented) inflow and negatively correlated with case growth.

TABLE IV: March Inflows from NYC vs. Case Growth

Panel A: Log(New Cases) Against Inflow from NYC				
	OLS		IV	
	(1)	(2)	(3)	(4)
High NYC Inflow	0.800*** (0.016)	0.299*** (0.011)	5.901*** (0.213)	5.262*** (0.315)
Controls	Y	Y	Y	Y
State × Month × Population Decile FE	Y	Y	Y	Y
N	64,287	63,228	64,287	63,228
Panel B: First Stage Estimates				
Weighted SCI			0.014 (0.0004)	0.142 (0.0077)
F			1,364	334

This table shows our baseline regression, where the outcome is the log(New Cases) and the inflow variable is replaced with an indicator equal to 1 if a county was in the top quartile of all counties which received inflows from NYC in March. The sample period is March 1 through July 13. Standard errors are in parentheses, and \* denotes 10% significance, \*\* denotes 5% significance, \*\*\* denotes 1% significance.

TABLE V: Total Inflows from NYC vs. Case Growth

Panel A: COVID-19 Case Growth Against Inflow from New York City												
	OLS						IV					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Inflow from NYC	201.893*** (29.120)	145.378*** (22.988)	146.173*** (23.024)				1701.690*** (90.574)	576.055*** (40.851)	525.346*** (38.314)			
Per Capita Inflow From NYC				222385.775*** (12297.432)	221671.430*** (12610.849)	267572.009*** (13777.545)				7750.832*** (306.637)	3010.973*** (195.477)	2728.955*** (190.074)
High Incoming Cases	4119.744*** (129.971)	4283.269*** (137.438)	4271.320*** (138.531)				1830.249*** (369.907)	3547.601*** (187.022)	3445.972*** (197.225)			
High Incoming Cases Per Capita				48.556*** (0.438)	49.001*** (0.443)	48.990*** (0.443)				0.035*** (0.001)	0.044*** (0.001)	0.043*** (0.001)
Far Indicator	4240.410*** (491.841)	2509.743*** (684.991)	2012.740*** (708.570)	-0.101 (0.541)	0.009 (1.114)	1.325 (1.121)	7228.572*** (610.141)	2624.269*** (680.683)	-6226.212*** (1437.595)	-0.015*** (0.001)	-0.000 (0.001)	-0.026*** (0.005)
Far Indicator × Inflow from NYC			498.949* (288.965)						8867.827*** (1296.926)			
Far Indicator × Per Capita Inflow From NYC						-305462.360*** (30595.187)						5932.824*** (1112.899)
Controls	N	Y	Y	N	Y	Y	N	Y	Y	N	Y	Y
State × Month × Population Decile FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
N	414,847	399,083	399,083	414,847	399,083	399,083	414,847	399,083	399,083	414,847	399,083	399,083

Panel B: First Stage Estimates of Inflow on SCI												
Weighted SCI							6759.300*** (268.328)	8362.031*** (363.540)	9188.484*** (392.011)	0.001*** (0.00002)	0.001*** (0.00002)	0.001*** (0.00003)
F							635	530	289	1,264	930	467

Columns 1-3 shows our main regression specification 1, with inflows from the five NYC boroughs as the explanatory variable rather than total inflows. Columns 4-5 repeat the exercise using inflow per capita as the outcome variable. Column 7-12 repeats the exercise for columns 1-6, where inflow from New York City is instrumented with the weighted SCI, as in 2. The sample period is March 1 through July 13. Standard errors are in parentheses, and \* denotes 10% significance, \*\* denotes 5% significance, \*\*\* denotes 1% significance. Note that all coefficients and standard errors in Panel A are scaled up by  $1 \times 10^5$ , with the exception of columns 4-6, where coefficients and standard errors are scaled up by  $1 \times 10^6$ . All coefficients and standard errors in Panel B are scaled up by  $1 \times 10^6$ .

TABLE VI: Impact of Migration on COVID-19 Cases with Clustering at the Commuting Zone Level

Panel A: COVID-19 Case Growth Against Inflow												
	OLS						IV					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Inflow	92.585*** (17.086)	12.022* (6.678)	8.288 (8.082)				267.821*** (36.036)	197.785** (80.534)	194.605** (85.408)			
Per Capita Inflow				-771.807*** (207.943)	-151.848 (224.998)	-199.698 (227.826)				-2969.855 (11562.496)	261.882 (223.374)	197.772 (147.866)
High Incoming Cases	4779.249*** (757.515)	4563.025*** (831.002)	4595.383*** (830.135)				5444.178*** (894.825)	6661.104*** (1463.364)	7040.705*** (1628.642)			
High Incoming Cases Per Capita				48.924*** (1.637)	49.372*** (1.649)	49.372*** (1.649)				-0.090 (0.543)	0.063*** (0.013)	0.060*** (0.009)
Far Indicator	3622.632 (3400.618)	2333.182 (1724.929)	-1362.127 (2869.619)	0.427 (2.024)	-0.119 (1.520)	-1.102 (2.756)	3214.693 (3336.820)	1842.587 (1797.050)	-18806.066* (10336.236)	0.317 (1.271)	0.008 (0.013)	-0.074 (0.060)
Far Indicator × Inflow			20.448 (16.700)						113.999** (56.528)			
Far Indicator × Per Capita Inflow						369.540 (856.118)						30.254 (24.184)
Controls	N	Y	Y	N	Y	Y	N	Y	Y	N	Y	Y
State × Month × Population Decile FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
N	414,847	399,083	399,083	414,847	399,083	399,083	414,847	399,083	399,083	414,847	399,083	399,083

Panel B: First Stage Estimates of Inflow on SCI											
Weighted SCI						42947.436*** (443.427)	29688.738*** (456.984)	32749.326*** (526.249)	-0.002 (0.006)	0.006 (0.005)	0.003 (0.005)
F						112	85	74	0	2	3

Columns 1-3 shows our main regression specification 1. Columns 4-6 repeat the exercise using inflow per capita as the outcome variable. Column 7-12 repeats the exercise for columns 1-6, where inflow is instrumented with the weighted SCI, as in 2. The sample period is March 1 through July 13. Standard errors are in parentheses, and \* denotes 10% significance, \*\* denotes 5% significance, \*\*\* denotes 1% significance. Note that all coefficients and standard errors in Panel A are scaled up by  $1 \times 10^3$ , with the exception of columns 4-6, where coefficients and standard errors are scaled up by  $1 \times 10^6$ . All coefficients and standard errors in Panel B are scaled up by  $1 \times 10^6$ . Estimates in this table have standard errors clustered at the Commuting Zone level.

Overall, our findings also potentially help explain the result in [Kuchler, Russel, and Stroebel \(2020\)](#), that greater social connections with Westchester (another pandemic hub) helps to predict subsequent COVID-19 deaths. A plausible transmission mechanism, which we explore here, is the refugee behavior of New York City residents into these socially connected regions, which has been frequently cited by local officials as a possible transmission mechanism of the disease.<sup>10</sup>

## IV CONCLUSION

We document substantial urban flight in the wake of the COVID-19 pandemic and find large effects of this migration on the growth in COVID-19 cases elsewhere in the country. We observe large migration responses by individuals living in major urban areas. In New York City, for instance, as much as 15-20% of Manhattan had fled by the middle of the summer in 2020. These individuals came from areas which were disproportionately wealthy, white, and young. These individuals appear to have left for regions with a high degree of social connections to New York City, suggesting that individuals are taking shelter with friends and family.

We then use the social networks structure to develop a causal estimate of the impact of migration on case growth. We find that instrumented migration patterns predict subsequent rise in cases in destination counties, suggesting that urban flight helped to seed the pandemic from an initially urban disease to a more widespread nation-wide pandemic.

Our work has implications for public policy in the wake of the disease. First, it highlights an important phenomenon of urban flight. Wealthy individuals, who contribute disproportionately to the local revenue and tax base of cities, are more likely to flee cities.

---

<sup>10</sup>For instance, the Governor of Rhode Island mandated a 14-day quarantine for New York City residents entering the state, punishable with a fine or arrest, see <https://www.nytimes.com/2020/03/28/us/coronavirus-rhode-island-checkpoint.html>

This finding points to important challenges for municipal finances in the wake of the pandemic, and has implications for the future of cities. Second, our work highlights the role of domestic migration in spreading the pandemic. As such, our work suggests the possible value of travel restrictions—in the form of bans, quarantine periods, or requiring testing as a precondition for entry—to help curb the spread of COVID-19.

## REFERENCES

- Allcott, H., L. Boxell, J. C. Conway, M. Gentzkow, M. Thaler, and D. Y. Yang, 2020, "Polarization and Public Health: Partisan Differences in Social Distancing during the Coronavirus Pandemic," Working Paper 26946, National Bureau of Economic Research.
- Andersen, M. S., 2020, "Early evidence on social distancing in response to Covid-19 in the United States," *Working Paper*.
- Athey, S., B. Ferguson, M. Gentzkow, and T. Schmidt, 2019, "Experienced Segregation," Working paper.
- Bailey, M., R. Cao, T. Kuchler, J. Stroebel, and A. Wong, 2018, "Social connectedness: Measurement, determinants, and effects," *Journal of Economic Perspectives*, 32(3), 259–80.
- Barrios, J. M., and Y. V. Hochberg, 2020, "Risk Perception Through the Lens of Politics in the Time of the COVID-19 Pandemic," *Working Paper*.
- Chandrasekhar, A. G., P. S. Goldsmith-Pinkham, M. O. Jackson, and S. Thau, 2020, "Interacting Regional Policies in Containing a Disease," .
- Chen, M. K., J. A. Chevalier, and E. F. Long, 2020, "Nursing Home Staff Networks and COVID-19," Working Paper 27608, National Bureau of Economic Research.
- Chen, M. K., K. Haggag, D. G. Pope, and R. Rohla, 2019, "Racial Disparities in Voting Wait Times: Evidence from Smartphone Data," Working Paper 26487, National Bureau of Economic Research.
- Chen, M. K., and R. Rohla, 2018, "The effect of partisanship and political advertising on close family ties," *Science*, 360(6392), 1020–1024.

- Chiou, L., and C. Tucker, 2020, "Social Distancing, Internet Access and Inequality," Working Paper 26982, National Bureau of Economic Research.
- Ciceri, F., A. Ruggeri, R. Lembo, R. Puglisi, G. Landoni, A. Zangrillo, and on behalf of the COVID-BioB Study Group, 2020, "Decreased in-hospital mortality in patients with COVID-19 pneumonia," *Pathogens and Global Health*, 114(6), 281–282, PMID: 32584660.
- Engle, S., J. Stromme, and A. Zhou, 2020, "Staying at Home: Mobility Effects of COVID-19," *Working Paper*.
- Fauver, J. R., M. E. Petrone, E. B. Hodcroft, K. Shioda, H. Y. Ehrlich, A. G. Watts, C. B. Vogels, A. F. Brito, T. Alpert, A. Muyombwe, et al., 2020, "Coast-to-coast spread of SARS-CoV-2 during the early epidemic in the United States," *Cell*.
- Gandhi, M., and G. W. Rutherford, 2020, "Facial Masking for Covid-19 Potential for "Variolation" as We Await a Vaccine," *New England Journal of Medicine*.
- Glaeser, E. L., C. S. Gorbach, and S. J. Redding, 2020, "How Much does COVID-19 Increase with Mobility? Evidence from New York and Four Other U.S. Cities," Working Paper 27519, National Bureau of Economic Research.
- Goolsbee, A., and C. Syverson, 2020, "Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020," Working Paper 27432, National Bureau of Economic Research.
- Holtz, D., M. Zhao, S. G. Benzell, C. Y. Cao, M. A. Rahimian, J. Yang, J. Allen, A. Collis, A. Moehring, T. Sowrirajan, D. Ghosh, Y. Zhang, P. S. Dhillon, C. Nicolaidis, D. Eckles, and S. Aral, 2020, "Interdependence and the cost of uncoordinated responses to COVID-19," *Proceedings of the National Academy of Sciences*.



- Horwitz, L., S. A. Jones, R. J. Cerfolio, F. Francois, J. Greco, B. Rudy, and C. M. Petrilli, 2020, "Trends in Covid-19 risk-adjusted mortality rates in a single health system," *medRxiv*.
- Kuchler, T., D. Russel, and J. Stroebel, 2020, "The Geographic Spread of COVID-19 Correlates with Structure of Social Networks as Measured by Facebook," Working Paper 26990, National Bureau of Economic Research.
- Lee, J., M. Mahmud, J. Morduch, S. Ravindran, and A. Shonchoy, 2020, "Migration, Externalities, and the Diffusion of COVID-19 in South Asia," .
- Painter, M. O., and T. Qiu, 2020, "Political Beliefs affect Compliance with COVID-19 Social Distancing Orders," *Working Paper*.
- Pybus, O., A. Rambaut, COG-UK-Consortium, et al., 2020, "Preliminary analysis of SARS-CoV-2 importation and establishment of UK transmission lineages," *Preprint at <https://virological.org/t/preliminary-analysis-of-sars-cov-2-importation-establishment-of-uktransmission-lineages/507>*.
- Rosenberg, C. E., 1968, *The Cholera Years: The United States in 1832, 1849, and 1866*. University of Chicago Press.

# INTERNET APPENDIX

FIGURE A1: Propensity to Leave Cities: Inflow

Panel A: March

Panel B: April



Panel C: May



# COVID-19 and the future of US fertility: What can we learn from Google?<sup>1</sup>

Joshua Wilde,<sup>2</sup> Wei Chen,<sup>3</sup> and Sophie Lohmann<sup>4</sup>

Date submitted: 6 October 2020; Date accepted: 8 October 2020

*We use data from Google Trends to predict the effect of the COVID-19 pandemic on future births in the United States. First, we show that periods of above-normal search volume for Google keywords relating to conception and pregnancy in US states are associated with higher numbers of births in the following months. Excess searches for unemployment keywords have the opposite effect. Second, by employing simple statistical learning techniques, we demonstrate that including information on keyword search volumes in prediction models significantly improves forecast accuracy over a number of cross-validation criteria. Third, we use data on Google searches during the COVID-19 pandemic to predict changes in aggregate fertility rates in the United States at the state level through February 2021. Our analysis suggests that between November 2020 and February 2021, monthly US births will drop sharply by approximately 15%. For context, this would be a 50% larger decline than that following the Great Recession of 2008-2009, and similar in magnitude to the declines following the Spanish Flu pandemic of 1918-1919 and the Great Depression. Finally, we find heterogeneous effects of the COVID-19 pandemic across different types of mothers. Women with less than a college education, as well as Black or African American women, are predicted to have larger declines in fertility due to COVID-19. This finding is consistent with elevated caseloads of COVID-19 in low-income and minority neighborhoods, as well as with evidence suggesting larger economic impacts of the crisis among such households.*

- 1 This research was improved by discussions with (in alphabetical order) Padmaja Ayyagari, Kieron Barclay, Kasey Buckles, Jennifer Caputo, Audrey Dorelien, Peter Eibich, Willa Friedman, Kathryn Grace, Melanie Guldi, Mahesh Karra, Mine Kühn, Giulia La Mattina, Susie Lee, Mikko Myrskylä, Daniela Negraia, Natalie Nitsche, Anna Oksuzyan, Lucie Schmidt, Alyson van Raalte, David Weil, and Emilio Zagheni. Author Contributions according to the CREDIT taxonomy: Joshua Wilde – Conceptualization, Methodology, Formal Analysis, Writing - Original Draft; Wei Chen – Investigation, Data Curation, Writing - Review & Editing; Sophie Lohmann – Software, Data Curation, Writing - Review & Editing. The authors declare no competing interest
- 2 Max Planck Institute for Demographic Research, and IZA Institute of Labor Economics.
- 3 Department of Economics, Fordham University.
- 4 Max Planck Institute for Demographic Research.

Copyright: Joshua Wilde, Wei Chen, and Sophie Lohmann

# 1 Introduction

The COVID-19 pandemic has had significant consequences for human mortality. At their peak, all-age weekly excess deaths exceeded 14% in Germany, 45% in the US, 90% in Italy, and 154% in Spain[1]. While most attention is focused on the current mortality and economic consequences of the pandemic, its effects on fertility are currently unknown[2]. In spite of this uncertainty, some in the popular media have suggested the pandemic will result in a “baby boom” as couples spend more time together. Such pronouncements are viewed with skepticism by many demographers, citing evidence on the short-term fertility effects from other mortality crises, including natural disasters [3] [4] [5] [6] and previous pandemics [7] [8]. These studies generally show that mortality spikes are followed by reductions in births within a year, with some evidence of fertility rebounds after several years [9]. Theoretically, these declines are explained by couples’ aversion to childbearing during periods of economic uncertainty, in an unhealthy environment, or by increased spontaneous pregnancy termination due to stress or disease [10]. While some rebounds have been even larger than the initial decline and led to long-run increases in fertility, these usually occur when mortality spikes are concentrated in children as parents attempt to replace deceased offspring — a mechanism which should be less relevant for the COVID-19 pandemic, since it primarily affects mortality among older individuals [5] [6].

Fertility change has significant economic and social consequences[11] [12]. Reductions in birthrates accelerate population ageing, increasing dependency ratios in populations already far below replacement fertility, such as in Southern Europe and East Asia[13]. This can lower economic growth by reducing the fraction of working age population, while simultaneously increasing the burden of caring for the elderly, both through public social programs as well as private channels[14] [15]. However, these social and economic effects are mainly affected by long-run fertility change. If the mortality crises only leads to a postponement of childbearing and leaves lifetime births per woman unchanged, the long-run economic effects from postponement should be minimal. However, recent post-crisis fertility declines — including the 2008-2009 financial crisis – did not experience a rebound and led to permanently lower fertility rates. Since the exact nature of the effect of COVID-19 on the future of human fertility is unclear, the economic and social affects of the crisis due to demographic change is also unknown.

Human gestation takes on average 268 days, so there is a natural delay from the onset of these crises and their effect on full-term births [16]. For example, full-term births from conceptions realized during the rapid onset of the pandemic in February or March of 2020 would not appear until November or December. Unfortunately, this delay in understanding the effect of the pandemic on fertility is compounded since natality data generally does not become available for analysis instantaneously. For example, the US Natality File birth microdata from the National Center for Health Statistics (NCHS) are generally released at least 6-9 months

after the end of the year in which those births occurred. Therefore, the earliest analyses of the pandemic's beginnings on birth outcomes using these traditional data will not be available until late 2021, with more detailed analyses including births from early 2021 not becoming available until late 2022. Even in countries which have faster data releases, significant delays between the advent of births and the release of data hamper the ability of researchers to make timely analyses of the relationship between COVID-19 and fertility rates.

An emerging literature suggests that data on Google searches can be used to monitor a number of social and biological phenomena in the absence of more reliable or timely data. Such data have been used for real-time analyses of disease outbreaks such as the seasonal flu and Dengue [17] [18], as well as studies on well-being [19], tourism [20], financial trading behavior [21], and demographic processes such as fertility [22], migration [23], sexual behavior [24] [25] [26], and mortality [27]. Moving beyond now-casting with Google data is generally difficult due to complexity or uncertainty surrounding the long-term processes which govern such phenomena. Fortunately, behavior and information-seeking surrounding human gestation takes place in predictable phases, and with well-known time lags, making Google search data particularly appealing for prediction. As an example, morning sickness generally only occurs during the first trimester, so an increase in searches for morning sickness and its symptoms may help pinpoint an increase in the fertility rate over half a year in the future, and with an accuracy of just a few months.

In addition to uncertainty surrounding the pandemic's effect on fertility, there may be important heterogeneous effects across sub-national regions, or between socioeconomic groups. For example, COVID-19 incidence and mortality have been elevated among the Black or African American community in the US, and the economic impacts have been particularly acute for workers with lower levels of education [28] [29]. Additionally, while planned births may fall as the economic fallout of the pandemic induces couples to delay childbearing, reduced contraceptive access may lead to an increase in unintended pregnancies. This effect is particularly acute for areas with historically poor contraceptive access: A 2020 UNFPA report noted that COVID-19 is already exacerbating unmet family planning needs due to a variety of reasons, including decreased demand for health facility visits, unavailability of trained medical staff, and supply chain disruptions for contraceptive commodities [30]. Analyzing differential changes in Google search volumes across regions with varying proportions of ethnic and other socioeconomic groups may yield deeper insights into the potential effects of COVID-19 on fertility.

In addition, heterogeneous changes in the types of keywords searched across regions may help pinpoint specific mechanisms by which the pandemic will affect fertility. For example, information-seeking on Google regarding emergency contraception may indicate possible changes in unplanned pregnancies, while searches for miscarriages may indicate fetal loss. Finally, inasmuch as behavior regarding conception and sexual behavior is considered by many to be socially taboo, some individuals may be more willing to search for such information on the

internet than to discuss their behavior in person [26]. Therefore, although Google data is only an imperfect reflection of the behaviors which affect fertility, it may more accurately reflect actual behavioral change than a direct self-report of those behaviors.

In this study, we use high-frequency data Google search term volume for keywords related to fertility to predict the direction, magnitude, and timing of fertility change expected from the pandemic in the United States at the state level. We do this in four steps. First, we show that periods of above-normal search volume for Google keywords relating to conception and pregnancy are associated with higher numbers of births in the following months. Excess searches for unemployment keywords has the opposite effect. Second, by employing simple statistical learning techniques, we demonstrate that including information on keyword search volumes in prediction models significantly improves forecast accuracy over a number of cross-validation criteria. Third, we use data on Google searches during the COVID-19 pandemic to predict changes in aggregate fertility rates in the United States at the state level through February 2021. Finally, we test for heterogeneous effects of the COVID-19 pandemic on fertility by education, ethnic identity, and age of mother, and by parity of birth.

## 2 Data and Methods

### 2.A Google Data

Our data on keyword search frequency comes from Google Trends (<http://trends.google.com>), a website which allows users to access data on Google keyword search frequency, stratified by geographic regions ranging from as large as a country to as small as a city. Data for smaller geographic areas can be more difficult to use, since the data are suppressed unless the overall search volume reaches a minimum threshold. For this article, we use data from the state level in the United States to ensure a wide selection of available search terms, yet still preserve geographic variation in search frequency.

### 2.B Natality Data

Data on births by state and month are from the National Vital Statistics System (NVSS) which is part of the National Center for Health Statistics (NCHS). Monthly birth counts for each US state and the District of Columbia from 2004-2018 were used, since 2004 was the year the Google search data began, and 2018 is the most recent year of data on births is available. This yields 15 years of usable data across 51 geographic regions, or 9,180 possible state-month-year observations.

## 3 Methods

### 3.A Keyword Selection

Keyword sets were created in a multi-step process. Initially, keywords were chosen by brainstorming possible words which we believed to be predictive, specific, and common enough for use in forecasting. We then sought input from other researchers to inform ourselves of any important topics or keywords we might be missing. Due to the necessarily arbitrary initial keyword set, all keywords were pre-selected before looking at Google data in order to avoid data-mining. Details on each selected keyword or topic and their meaning are discussed in the Supplemental Information. The list of initial keywords or topics are: Baby Heartbeat, BabyCenter, Birth Control, Clearblue, Conceive, Cytotec, Dilation and Curettage, Divorce, Emergency Contraception, File Unemployment How, Folate, Furlough, Human chorionic gonadotropin (hCG), Intra-uterine Device (IUD), In-vitro Fertilization (IVF), Layoff, Medical Abortion, Midwife, Miscarriage, Misoprostol, Missed Period, Morning After Pill, Morning Sickness, Obygn, Online Dating, Ovulation, Ovulation Test, Plan B, Porn, Pregnancy, Pregnancy Bleeding, Pregnancy Symptoms, Pregnancy Test, Pregnant, Prenatal Vitamins, Sexually Transmitted Infection (STI), Ultrasound, Unemployment, Unemployment Office, Unprotected Sex, and a control set. We then subjected each keyword to a series of three screens in order to test goodness of fit and predictive power. These screens are mentioned in the main text, and described more fully later in this section and in the Supplemental Information. In addition, details on our statistical learning method for selecting the most predictive keywords for the MSPE Reduction keyword set are also outlined in the Supplemental Information.

### 3.B Prediction Model

Our baseline prediction model is an OLS regression as follows:

$$Y_{sm y} = \alpha_{sm} + \sum_{s=1}^2 \gamma_s * t^s + \sum_w^W \sum_{l=t_0}^T \beta_{w,t-l} I_{sm y}^w + \epsilon_{sm y} \quad (1)$$

where  $s$ ,  $m$ , and  $y$  index state, month, and year respectively,  $\alpha_{sm}$  is a state-month fixed effect,  $\gamma_s * t$  is a state specific linear and quadratic time trend. The independent variables  $I_{sm y}^w$  are the natural log of the normalized search volume as given by the google data and whose construction are outlined in the supplemental information. The double summation represents a series of  $\beta$  coefficients for the effects of different keywords search volumes at a number of monthly time lags. The dependent variable  $Y$  is the natural log of births, implying the interpretation of the  $\beta$ s are an elasticity — the percentage change in births from each percentage change in monthly search volume. We use  $t_0$  of 7 and a  $T$  of 12, representing monthly time lags from 7-12 months before birth. Huber–White standard errors are utilized for estimation.

### 3.C Data Screens and Cross-Validation Techniques

We use three data screens to reduce our initial keyword set. We first omit any keyword for which more than 1/3 of the 9,180 state-month observations were missing. Second, we further omitted any keywords that were not jointly or individually significant predictors of births for the lagged months 7-12. We call this the Significance screen, which selected 14 keywords for inclusion, listed in alphabetical order: ClearBlue, Divorce, HCG, IVF, Layoff, Morning Sickness, Ovulation, Porn, Pregnancy Test, Pregnancy, Pregnant, STI, Ultrasound, and Unemployment.

Next, the MSPE Reduction screen used a statistical learning procedure called forward step-wise selection to select the keywords from the Significance screen which lead to the largest reduction in mean squared prediction error (MPSE). To find this MSPE reduction, we used a variation of the "Leave-One-Out" cross-validation technique in which for each tested keyword set, we omit two years of data, train the data using the remaining 13 years, and then use the estimated associations to predict births in the omitted two years. We then calculate the prediction error as the absolute value — in percentage terms — of the difference between these predictions and the actual births for every state-month. We then repeat for every possible combination of two years of data, resulting in 105 model runs. We then average the mean squared prediction error across the runs to find the MSPE for that keyword set.

The forward step-wise selection learning procedure proceeds as follows. First, the base model without any keywords is estimated. Then, we add each keyword one at a time to the model (and a control keyword) and employ the MSPE calculation procedure described above for each word. The word which minimizes the MSPE the most is selected, and becomes part of the base model. Then a second round begins, where each remaining word is selected to be added to this new base model one at a time, and the word which minimizes the MSPE is selected. This procedure continues until the additional reduction in MSPE is less than one percentage point. This screen selected five keywords, listed in order of selection: Unemployment, HCG, Clearblue, Unemployment Office, and Ovulation. (See Supplemental Information for more details.)

Additionally, two extra topical screens were utilized to explore heterogeneity between word types as described in the article — the Early indicators screen (including the words Clearblue, HCG, Morning Sickness, Ovulation, Pregnancy Test, Pregnancy, Pregnant), and Unemployment screen (including the words Unemployment, Unemployment Office, and Layoff).

## 4 Results

In our first step, we tested associations between 40 fertility-related keywords and monthly births at the state level in the United States between 2004 and 2018. Our keywords are listed in the Material and Methods section, with further descriptions and keyword selection methodologies are found in the Supplemental Information. Our estimation methodology utilized an Ordinary



Least Squares (OLS) regression which controlled for seasonality in both births and keyword searches, which controls for time-independent state effects and state-specific time polynomials. We also included a set of control keywords as a proxy for overall internet use.

We estimated the month-specific effect of each keyword on births for each month between 1 to 12 months before the observed state-level births, and plot these results for a subset of keywords in Figure 1. Specifically, we show the results for seven of the most predictive keywords (ClearBlue, Morning Sickness, IVF, Ovulation, Pregnancy, Ultrasound, and Unemployment) and our control keyword set. These findings confirm that there are intuitive associations between fertility-related keywords and future births. For example, excess searches for "ClearBlue" — the name of a popular brand of pregnancy test in the US — is significantly associated with more births 5-8 months later, with the largest effect 7 and 8 months later. This is consistent with the timing when many pregnant women are experiencing their first pregnancy symptoms at 4-8 weeks after conception. Similarly, excess searches for "Morning Sickness" are associated with more births 6-8 months later, corresponding with the first trimester when morning sickness generally occurs.

For our second step, we conduct model cross-validation to show that including subsets of pregnancy related keywords in our estimation model reliably improve the predictive power beyond our baseline model. This is important, since simple associations between individual keywords may be spurious, or a result of model over-fitting. We use well-accepted statistical learning methodologies and divide our data into a training data set and a test data set – the first of which is used to estimate the associations between keywords and births, and the second to test the predictive power of the keywords out-of-sample. We find that including these keywords increases forecast model precision by approximately 25% over a number of criteria. Descriptions of these cross-validation techniques and their results are given in the Supplemental Information section.

Third, we used these estimated associations between future births and pregnancy-related keyword searches to forecast state-level births through February 2021 using Google search volumes up to July 2020. These results are displayed in Figure 2, in which state-specific predictions using our model with Google searches are aggregated to the national level, and then shown relative to the same predictions using the model without searches. Displaying our results in terms of this relative index has the benefit of showing the predicted deviation in births between the models in percentage terms relative to what would be otherwise expected in a comparable month. We compared four different keyword sets against this baseline, where the selection criteria for each are fully explained in the Supplemental Information. Briefly however, the Significance keyword set contains those keywords which were both jointly significant across all monthly lags included in our model and individually significant for at least one monthly lag; the MSPE keyword set was the set of words which minimized the Mean Square Prediction Error (MSPE) using a forward step-wise keyword selection criteria; the Early keyword set only included words which were statistically significant and concerned early pregnancy keywords

Figure 1: Fertility Keyword Searches and Later Births. These figures show regression coefficients between births in a state and Google keyword search volume for the preceding 12 months. Coefficients are elasticities, and can be interpreted as the change in births due to a doubling of keyword search volume in a given month. Dashed lines represent 95% confidence intervals.

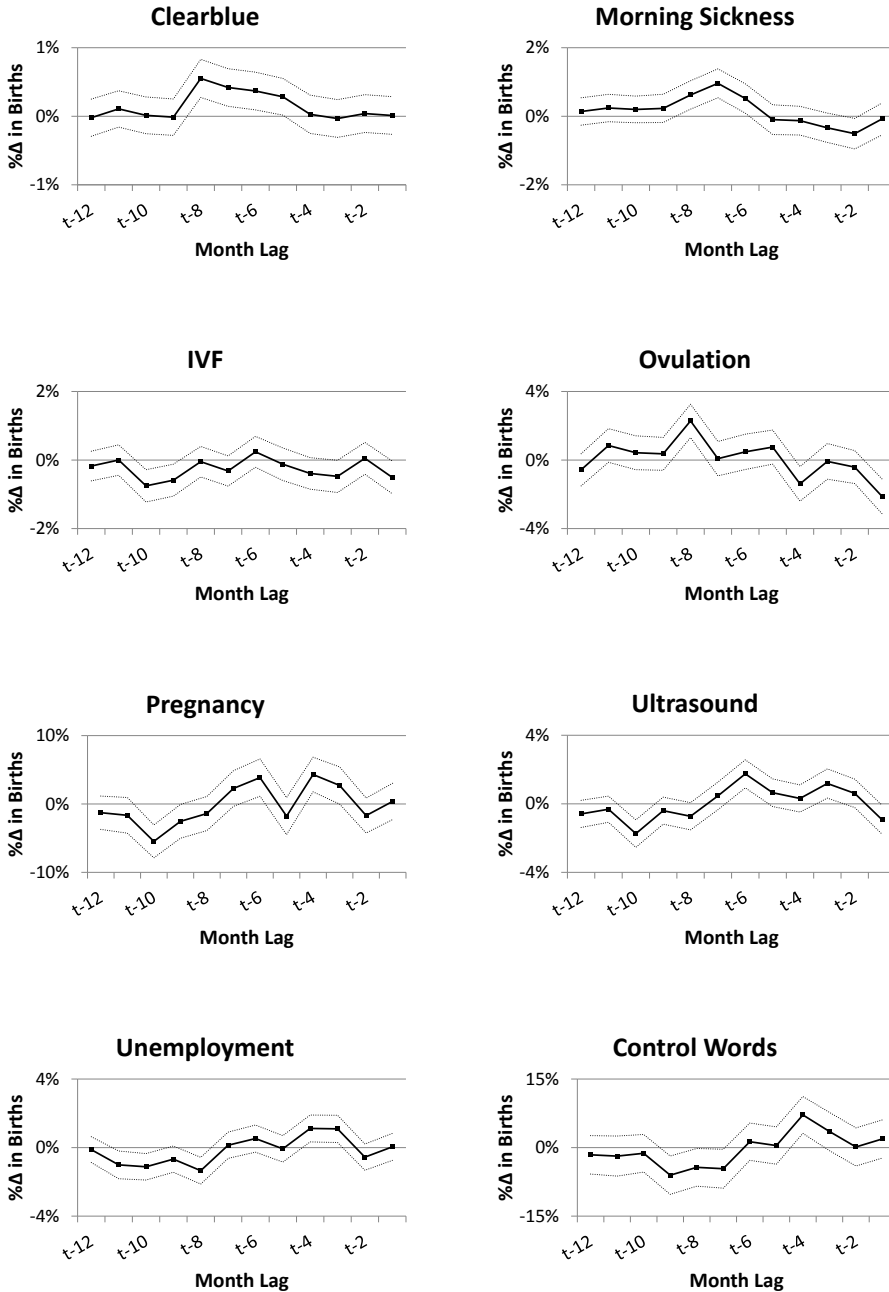



Figure 2: Predicted 2020 US Births by Month. National predicted births, for various keyword sets, relative to the corresponding prediction without keyword searches, and normalized to 100 in October 2020. The four keyword sets are: Significance, including keywords whose associations meet statistical significance thresholds; MSPE, including keywords selected from a mean-squared prediction error minimizing forward step-wise learning selection criteria; Early, including keywords related to early pregnancy; and Unemployment, including keywords related to economic conditions.



unrelated to health services; an Unemployment keyword set which only contained statistically significant keywords related to unemployment.

Each of the keyword sets — with the exception of Early pregnancy indicators — show similarly striking results: beginning in November 2020, fertility is predicted to fall by several percentage points each month until February 2021. For example, using the MSPE reduction keyword set, we predict the number of births in February 2021 will be just 84.3% of what they would have been for a normal February. Estimates for the Unemployment keyword set are almost identical to those for the MSPE Reduction keyword set. The Significance keyword set predicts a smaller — yet still very large — decline to just 88.1% of expected births. All of these declines are highly statistically significant at the 0.1% level. However, the Early indicators keyword set shows almost no decline at all — fertility is predicted to be a statistically insignificant 98.8% of its normal level.

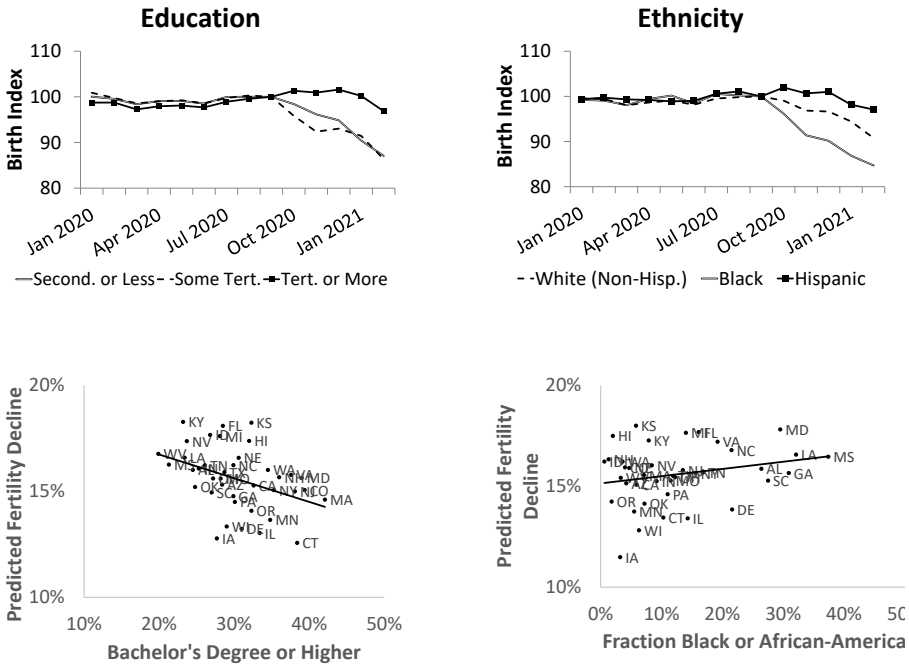
Covid Economics 54, 29 October 2020: 158-190

These disparate estimates can be explained by examining changes in Google search volume for specific keywords as a result of the pandemic. In the Supplemental Information, we show Google keyword search volume for the keywords shown in Figure 1, after having been corrected for seasonality. Several patterns emerge in these figures. First, whereas searches for many words fell dramatically around the first state-wide shutdown, many fertility-related search terms have remained surprisingly stable throughout the crisis. The words which did decline have a unifying commonality: they all relate to prenatal or conception services which must be experienced in a clinical setting. For example, searches related to “Obgyn”, “Ultrasound”, and “IVF” (In-vitro fertilization) all fell sharply. However, it is unlikely that searches for ultrasounds — which generally are associated in our data with pregnancies in the late 1st and 2nd trimester — fell sharply due to significant declines in conceptions in December and January. The abrupt decline in searches for these words likely reflected an interruption in clinical services due to state-wide shutdowns and social distancing behavior. Conversely, highly predictive words which do not involve prenatal services such as “ClearBlue”, “Morning Sickness”, “Ovulation”, and “Pregnancy” were essentially unaffected by the coronavirus pandemic. Therefore, in spite of their predictive power, forecasts based solely on these keywords — such as our Early indicators keyword set — imply no large negative effect of the pandemic on fertility.

However, the results are reversed for any keyword set which includes words related to unemployment. It is well documented that fertility is sensitive to economic conditions, and generally follows a counter-cyclical pattern[10]. As seen in Figure 1, excess searches for unemployment around and immediately before the time of conception (8-11 months before birth) are strongly associated with lower birthrates. This is presumably due to some couples timing their births to avoid periods of financial hardship or economic uncertainty[10]. In addition, searches for unemployment jumped to 20 times their normal level by April – a month after the first statewide lockdown – and stabilized by the beginning of June to approximately half this amount. The combination of large changes in search volumes for unemployment and large coefficient values from our prediction model for this keyword imply that the pandemic and its subsequent economic fallout will strongly decrease fertility, as shown in the Unemployment, Significance, and MSPE keyword set estimates in Figure 2.

Finally, we explore whether our predicted effects vary across different states, or by type of birth. Among the latter, we consider four sources of heterogeneity – maternal education, age, and ethnicity, and by parity of birth. We find sizable differences across US states: the state with the highest predicted decline between October 2020 and February 2021 using the MSPE reduction set is Hawaii with reduction of 23.6%, while the lowest decline is found in Connecticut of only 11.9%. As shown in the Supplemental Information, the largest predicted declines are generally in the Southern US and the Rust Belt, with the exception of Hawaii. In Figure 3, we only present our results based on differences in maternal education and ethnicity – those for parity and maternal age are found in the Supplemental Information, and demonstrate

Figure 3: Predicted Fertility Decline by Various Mother Characteristics. The top panel shows national predicted births, by group, for the Significance keyword set, relative to the corresponding prediction without keyword searches, and normalized to 100 in October 2020. The bottom panel shows a scatter-plot relating the predicted fertility decline between Oct. 2020 and Feb. 2021 – after controlling for COVID-19 cases per capita and population density – by the fraction of the population in the respective subgroup in that state. Data source: American Community Survey 2013-2017, authors calculations.



somewhat smaller declines for first births and no differences by maternal age.

In Figure 3, we find that mothers with less of education or who identify as Black or African American are predicted to have sharper declines in births. We do this in two ways. First, we change the dependent variable in our prediction model from all births in a state to only births for each educational or ethnic subgroup, and present the same relative birth index as in Figure 2, but now broken out by subgroup. For comparison, we do this analysis using the Significance keyword set, which predicted an average decline in fertility of 11.9%. These results are found in the top row of Figure 3. Second, in the bottom row of Figure 3 we plot the overall predicted decline, by state, between October 2020 and February 2021 — after controlling for COVID-19 cases per capita and population density — against the fraction of the population in each state belonging to various ethnic or educational groups. We find that for those who have completed a Bachelor’s degree or higher, the COVID-19 pandemic is predicted to have very little effect on

Covid Economics 54, 29 October 2020: 158-190

fertility. Specifically, by Feb 2021 we only predict a statistically insignificant 3.1% decrease in fertility for this group. However, for both women with only some college, and those with only a high school degree, we predict a 13% decline. Similarly, states with a higher fraction of women with at least a Bachelor's degree have smaller predicted fertility decline than states with lower Bachelor's degree attainment. By ethnic group, we find very small effects for Hispanics (2.9% decline) relative to non-Hispanic Whites (9.2% decline) and Black or African Americans (15.3% decline). We also find that states with larger Black or African American populations have larger predicted declines in fertility.

## 5 Discussion

A 15% collapse in births over a four-month period would be unprecedented in the modern United States. Therefore, it may be useful to explore the plausibility of this prediction. To do so, we consider three major crises in the US with similarities to the current pandemic: the Spanish Flu pandemic of 1918-1919, the Great Depression of 1929-1933, and the financial crisis of 2008-2009. In doing so, it is important to keep in mind that the context of the current pandemic is fundamentally different from each of these crises, and therefore these comparisons should be taken with a measure of caution. However, as the coronavirus pandemic itself is in many ways unprecedented, these three crises arguably serve as the best comparisons at our disposal.

The first crisis we consider is the 1918–1919 H1N1 influenza A pandemic — commonly known as the Spanish Flu. Birth rates fell 13% from 1918 to 1919 in the United States, with a small rebound in 1920 and 1921[8]. This decline is strikingly similar to that predicted for the current pandemic in this study. The fertility effects of the Great Depression are also very similar in magnitude to the Spanish Flu and our prediction for the COVID-19 pandemic. Between 1929 and 1933, birthrates fell by 15.2%[31]. Finally, the financial crisis of 2008-2009 caused a smaller decline in fertility, albeit in response to a smaller economic crisis. From May 2008 to October 2010, births fell by 9.3% and failed to rebound thereafter.

The evolution of births over these three crises suggests that a 15% decline in the fertility rate in response to the COVID-19 pandemic is not unreasonable. Using the above information on changes in unemployment and changes in births, the response of predicted births to unemployment for the current pandemic is firmly between the response from the financial crisis and the Great Depression. Calculating the elasticity of births to unemployment — defined as the ratio between the percentage change in births and the percentage change in unemployment — yields an elasticity of -0.050 for the current pandemic, firmly between that of the financial crisis (-0.109) and the Great Depression (-0.027). Importantly, if there is a rebound in births later in the year, the decline in fertility could be significantly less than 15%, lowering this elasticity substantially. Therefore, the main difference between the prediction of this study and the evolution of fertility during these three crises is not the magnitude of the decline, but rather the

speed. Since the current pandemic has led to historically fast increases in unemployment, this suggests that similarly rapid decreases in fertility may be expected.

Our finding that individuals identifying as Black or African American, or with lower educational attainment, have larger predicted declines in fertility can be interpreted in two ways. First, areas with high concentrations of African Americans have been more impacted by COVID-19: incidence among Black individuals was 2.6 times higher than for non-Hispanic whites, and mortality was 2.1 times higher[28]. Since individuals of lower socioeconomic status are disproportionately affected by the virus, this could manifest in higher predicted fertility declines[29]. Second, it may be that the heterogeneity in the predicted fertility effect is not caused by the differential impact of the virus itself, but rather from the uneven economic fallout of the pandemic. Historically, during recessions employment losses are often concentrated among minority groups – and among Black individuals in particular – and the current economic downturn is no exception. Between February and April of 2020, the unemployment rate among non-Hispanic White Americans rose from 3.1% to 14.2%, while for Black or African Americans it rose from 5.8% to 16.7%. While these initial increases were similar, the recovery has been much slower for Black or African Americans: by August of 2020 non-Hispanic White unemployment had fallen back to 7.3%, yet Black unemployment remained elevated at 13.0%. The patterns in unemployment between educated and uneducated workers were even more striking, peaking at only 8.4% for those with a Bachelor's degree or higher, compared with 17.3% and 21.2% for non-college graduates with and without a high school diploma respectively[32].

There are several important caveats to this study which should be noted. First, the evidence presented here can be interpreted in two ways. According to the first interpretation, since google search volume related to pregnancy behavior and symptoms alone – and the corresponding fertility prediction – did not change significantly over the course of the pandemic, there is little reason to think that there will be a large decrease in births as a result. In this view, the effect of the large implied changes in unemployment can be ignored, because if high unemployment searches truly implied fewer future births, we would expect fewer searches for morning sickness, pregnancy tests, and other fertility keywords as a result. We do not observe such reductions. However, the second interpretation concludes that because unemployment searches are strongly associated with future births, and because these searches have spiked since the beginning of the pandemic, fertility will decline by approximately 15%. Since our methodology shows that searches related to unemployment are far more predictive than the rest of our keywords combined, an objective data science perspective would favor this interpretation. Importantly however, neither interpretation includes the possibility of a large increase in births, notwithstanding recent speculation.

The second important caveat is that this method can only predict fertility change up to 9-12 months in the future, leaving the long-run effects of the pandemic unknown. If the decline in fertility is simply due to a postponement in births, completed lifetime fertility may remain relatively unchanged. However, even if these predicted declines are primarily due to couples

postponing children they still intend to have, this may still negatively affect long-run fertility. Since the 1970s there has been a shift towards births at later ages as women postpone childbearing for educational or career motives[33]. Since fecundity declines rapidly towards the end of a woman's childbearing years, women in their late 30s and early 40s who further delay childbearing due to the pandemic may face unexpected difficulties trying to conceive, leading to unintentionally lower completed lifetime fertility [34]. If this happens, the resulting lower completed fertility and accompanying shift to an older population age structure has significant social and economic consequences, particularly for regions already well below replacement fertility. Decreasing tax revenue, greater costs for social programs for the elderly, pension liabilities, increased health care costs, reduced economic growth, and increased burdens of caring for the elderly are all important policy issues surrounding falling fertility rates and the resulting aging of the population. [33] [15] [11] [12].



## References

- [1] Centers for Disease Control. Excess deaths associated with covid-19, 2020.
- [2] A. Aassve, N. Cavalli, L. Mencarini, S. Plach, and M. Livi Bacci. The COVID-19 pandemic and human fertility. *Science*, 369(6502):370–371, July 2020.
- [3] Jocelyn E. Finlay. Fertility Response To Natural Disasters: The Case Of Three High Mortality Earthquakes. *POLICY RESEARCH WORKING PAPERS*, apr 2009.
- [4] Richard W. Evans, Yingyao Hu, and Zhong Zhao. The fertility effect of catastrophe: U.S. hurricane births. *Journal of Population Economics*, 23(1):1–36, nov 2009.
- [5] Jenna Nobles, Elizabeth Frankenberg, and Duncan Thomas. The Effects of Mortality on Fertility: Population Dynamics After a Natural Disaster. *Demography*, 52(1):15–38, jan 2015.
- [6] Julia Andrea Behrman and Abigail Weitzman. Effects of the 2010 Haiti Earthquake on Women’s Reproductive Health. *Studies in Family Planning*, 47(1):3–17, mar 2016.
- [7] Jacques Bertillon. *Annuaire statistique de la ville de Paris*. Imprimerie nationale, 1895.
- [8] Siddharth Chandra, Julia Christensen, Svenn-Erik Mamelund, and Nigel Paneth. Short-term birth sequelae of the 1918–1920 influenza pandemic in the United States: state-level analysis. *American Journal of Epidemiology*, 187(12):2585–2595, 2018.
- [9] Alberto Palloni. On the Role of Crises in Historical Perspective: An Exchange: Comment. *Population and Development Review*, 14(1):145–158, 1988.
- [10] Tomáš Sobotka, Vegard Skirbekk, and Dimiter Philipov. Economic recession and fertility in the developed world, jun 2011.
- [11] Quamrul H. Ashraf, David N. Weil, and Joshua Wilde. The effect of fertility reduction on economic growth. *Population and Development Review*, 39(1):97–130, mar 2013.
- [12] Mahesh Karra, David Canning, and Joshua Wilde. The Effect of Fertility Decline on Economic Growth in Africa: A Macrosimulation Model. *Population and Development Review*, 43:237–263, may 2017.
- [13] Population Reference Bureau. 2020 world population data sheet, 2020.
- [14] Nicole Maestas, Kathleen J. Mullen, and David Powell. The effect of population aging on economic growth, the labor force and productivity. *NBER Working Paper No. 22452*, 2016.

- [15] John R. Beard and David E. Bloom. Towards a comprehensive public health response to population ageing. *The Lancet*, 385(9968):658–661, February 2015.
- [16] A.M. Jukic, D.D. Baird, C.R. Weinberg, D.R. McConaughy, and A.J. Wilcox. Length of human pregnancy and contributors to its natural variation. *Human Reproduction (Oxford, England)*, 28(10):2848–2855, October 2013.
- [17] Jeremy Ginsberg, Matthew Mohebbi, Rajan Patel, Lynnette Brammer, Mark Smolinski, and Larry Brilliant. Detecting influenza epidemics using search engine query data. *Nature*, 457:1012–1014, 2009.
- [18] Herman Anthony Carneiro and Eleftherios Mylonakis. Google trends: a web-based tool for real-time surveillance of disease outbreaks. *Clinical Infectious Diseases: An Official Publication of the Infectious Diseases Society of America*, 49(10):1557–1564, November 2009.
- [19] Abel Brodeur, Andrew E Clark, Sarah Flèche, and Nattavudh Powdthavee. COVID-19, Lockdowns and Well-Being: Evidence from Google Trends. page 55, 2020.
- [20] Boriss Siliverstovs and Daniel S. Wochner. Google Trends and reality: Do the proportions match? *Journal of Economic Behavior & Organization*, 145(C):1–23, 2018.
- [21] Tobias Preis, Helen Susannah Moat, and H. Eugene Stanley. Quantifying Trading Behavior in Financial Markets Using Google Trends. *Scientific Reports*, 3(1):1684, April 2013.
- [22] Francesco Billari, Francesco D’Amuri, and Juri Marcucci. Forecasting Births Using Google. In *Proceedings of the 1st International Conference on Advanced Research Methods and Analytics*. Universitat Politècnica València, July 2016.
- [23] Emilio Zagheni and Ingmar Weber. You are where you e-mail: using e-mail data to estimate international migration rates. In *WebSci ’12*, 2012.
- [24] Patrick M. Markey and Charlotte N. Markey. Seasonal Variation in Internet Keyword Searches: A Proxy Assessment of Sex Mating Behaviors. *Archives of Sexual Behavior*, 42(4):515–521, May 2013.
- [25] Joshua Wilde, Bénédicte H. Apouey, and Toni Jung. The effect of ambient temperature shocks during conception and early pregnancy on later life outcomes. *European Economic Review*, 97:87–107, August 2017.
- [26] Seth Stephens-Davidowitz. *Everybody Lies: Big Data, New Data, and What the Internet Can Tell Us About Who We Really Are*. Harper-Collins, 2017.

- [27] J. K. Tamgno, R. M. Faye, and C. Lishou. Verbal autopsies, mobile data collection for monitoring and warning causes of deaths. *2013 15th International Conference on Advanced Communications Technology (ICACT)*, 2013.
- [28] Centers for Disease Control. Covid-19 hospitalization and death by race/ethnicity, 2020.
- [29] W. Holmes Finch and Maria E. Hernández Finch. Poverty and covid-19: Rates of incidence and deaths in the united states during the first 10 weeks of the pandemic. *Frontiers in Sociology*, 5:47, 2020.
- [30] UNFPA - United Nations Population Fund. Impact of the COVID-19 Pandemic on Family Planning and Ending Gender-based Violence, Female Genital Mutilation and Child Marriage. Technical report, 2020.
- [31] Robert Heuser. *Fertility tables for birth cohorts by color*. Number DHEW Publication No. (HRA)76-1152. U.S. Department of Health, Education, and Welfare.
- [32] Bureau of Labor Statistics. Unemployment rate - series Ins14000003, Ins14000006, Ins14027662, Ins14027660, and Ins14027659, 2020.
- [33] GBD 2017 Population and Fertility Collaborators. Population and fertility by age and sex for 195 countries and territories, 1950–2017: a systematic analysis for the Global Burden of Disease Study 2017. *The Lancet*, 392(10159):1995–2051, 2018.
- [34] Emmanuel Attali and Yariv Yogev. The impact of advanced maternal age on pregnancy outcome. *Best Practice Research Clinical Obstetrics Gynaecology*, 2020.

# Supplemental Information

## Word Selection Methodology

Choosing a set of keywords with which to analyze search frequency is not straightforward for several reasons. First, there is no *a priori* agreed upon set of keywords which should be both theoretically associated with — and empirically predictive of — pregnancy. Second, words associated with pregnancy generally have multiple meanings, which meanings can also change over time.<sup>1</sup> Third, different keyword sets may have different associations for different types of pregnancy behaviors.<sup>2</sup> Finally, the same keywords may theoretically have different correlations with births at various pregnancy timeframes.<sup>3</sup>

The lack of a defined, agreed upon pregnancy keyword set may lead to several statistical and ethical problems when determining which keywords to include. For example, an unscrupulous researcher may datamine keywords to find those which are correlated with pregnancy, ignore those which don't, and present the highly predictive nature of these keywords as evidence that Google searches are excellent at predicting births. Even consciously ethical researchers may be tempted to engage in unconscious motivated reasoning to include or exclude words once they see the results on how predictive they actually are. Unfortunately, without a predetermined word set, part of the word selection process is necessarily at least partially arbitrary.

## Pregnancy and Birth Keywords

The first step in our word selection process was to consult a number of sources to identify possible words. We consulted the literature on Google keywords and any aspect of fertility – unfortunately, these are very scarce. The only paper we could find specifically on pregnancy and Google searches was Billari et al (2013), which utilized three keywords – ovulation, pregnancy, and maternity. Other papers utilized keyword sets which, while related to pregnancy, were not pregnancy specific. For example, Markey and Markey (2013) developed a keyword set for pornographic, sexual, and mate-seeking keywords, which was subsequently used in Wilde, Apouey, and Jung (2017) and Wilde, Lohmann, and Chen (2020). Additionally, Lohmann and

<sup>1</sup>For example, the term “Plan B” – a popular brand-name emergency contraceptive in the United States which can reduce unplanned pregnancies – could also refer to generic back-up plans completely unrelated to pregnancy. Another example is the word “Tinder” – before 2012, it referred to a type of firework, yet after 2012 the majority of searches referred to an online dating app.

<sup>2</sup>For example, searches for emergency contraceptives may indicate a higher incidence of risky or unplanned sexual behavior, which may lead to more births. However, the same searches for emergency contraceptive may be an indication that for the same amount of risky sexual behavior, a higher proportion of those acts are accompanied by actions to reduce unplanned pregnancy, leading to fewer pregnancies.

<sup>3</sup>For example, searches regarding pregnancy or ovulation tests may be positively associated with pregnancy 9 to 11 months later as it suggests more individuals may be trying to get pregnant. However, a high frequency of searches for these same tests will likely be negatively associated with births 1 to 5 months later, since if a large fraction of women in a population are currently pregnant, there should be a lower fraction of women trying to get pregnant.

Albarracin (2019) developed an exhaustive keyword set for sexually transmitted infections. We then consulted online sources such as pregnancy and maternity websites, dictionaries, and thesauruses. Finally, we conducted a number of interviews with women which had experiences with different aspects of pregnancy, such as miscarriage, unexpected pregnancies, abortion, etc.

After this initial round, we reduced our keyword set to those 40 keywords which we believed would be 1) specific enough to be highly correlated with births, 2) broad enough to have enough data to run our analysis, and 3) which would sufficiently cover each of the following seven aspects of pregnancy: Unplanned Pregnancies, Pregnancy Intentions, Pregnancy Symptoms, Prenatal Services, Miscarriage, Economic Indicators, and Other. Importantly, we determined this final keyword set *before looking at any data*. These words and word groupings can be found in Table 1 in this supplement.

Once fixing this initial keyword set, we employed three systematic screens to allow us to objectively choose keyword subsets for analysis without conscious or subconscious bias in word selection. All three of these screens take a data-oriented approach. The first screen regarded sample size – if the search frequency for a given keyword was missing for more than 1/3 of the total state-month observations, it did not pass screen one. Of our 40 initial words, only 6 did not pass this screen: Morning After Pill, Unprotected Sex, Baby Heartbeat, Cytotec, File Unemployment How, and Furlough. Screens two and three were data-driven relevance screens, are referred to as the Significance and MPSE screens in the body of the paper, and will be described in detail in the Methodology section of this supplement. In short, a word passed screen two if it was jointly partially correlated with births at specific time lags in our main regression specification, and passed screen 3 if it was selected by a commonly used variable selection algorithm.

## Control Keywords

Our method for determining control keywords was similar to that which determined pregnancy keywords, in that we began by selecting a large set of possible control words. We combined a set of two keyword groups: commonly searched words as reported by Google, and common – but not specific – words in the English language. Specifically, we created an initial keyword set of the 25 most searched words reported by Google in the United States from 2004 to the present, in addition to 40 of the most commonly used words in the English language: the nominative and accusative personal pronouns (me, you, I, us, we, they, she, he, it); the major interrogative pronouns and other “w”-words (who, what, where, why, how, which, whose, whom); definite and indefinite articles (the, a, an); common conjunctions (and, but, or); superlatives (best, worst); various conjugations of common verbs (be, is, are, were, do, does, did, have, has, had, can, could, would, should); and “of”. Of these 65 words, the top ten words with the highest search volumes are: the, of, how, a, is, what, and, me, you, and do. Searches for these ten

Table 1: Keyword Selection Criteria.

Word	Topic or Keyword	Screen		
		1	2	3
<b>Controls</b>				
Control Set*	Keyword	X	X	X
<b>Unplanned Pregnancy</b>				
Emergency Contraception	Topic	X		
“Morning After Pill”	Keyword			
Plan B	Keyword	X		
STI	Topic	X	X	
“Unprotected Sex”	Keyword			
Online Dating*	Keyword	X		
<b>Pregnancy Intention</b>				
Clearblue	Keyword	X	X	3
Conceive	Keyword	X		
HCG - Weight - Diet	Keyword	X	X	2
IVF	Topic	X	X	
Ovulation	Topic	X	X	5
Ovulation Test	Keyword	X		
Pregnancy Test	Topic	X	X	
<b>Pregnancy Symptoms</b>				
Missed Period	Keyword	X		
Morning Sickness	Topic	X	X	
Pregnancy Symptoms	Keyword	X		
<b>Prenatal Services</b>				
BabyCenter	Topic	X		
Folate	Topic	X		
Midwife	Topic	X		
Obygn	Keyword	X		
Prenatal Vitamins	Topic	X		
Ultrasound	Keyword	X	X	
<b>Pregnancy Termination</b>				
Baby Heartbeat	Keyword			
Cytotec	Keyword			
Dilation and Curettage	Topic	X		
Medical Abortion	Topic	X		
Miscarriage	Topic	X		
Misoprostol	Topic	X		
Pregnancy Bleeding	Keyword	X		
<b>Unemployment</b>				
File Unemployment How	Keyword			
Furlough	Keyword			
Layoff	Topic	X	X	
Unemployment	Keyword	X	X	1
“Unemployment Office”	Keyword	X		4
<b>Other</b>				
Birth Control	Topic	X		
Divorce	Topic	X	X	
IUD	Topic	X		
Porn	Keyword	X	X	
Pregnancy	Topic	X	X	
Pregnant	Keyword	X	X	

\*See Google Data section for extra details.

words comprise over 50% of the combined search volume for the larger 65 word control set. For simplicity, we use these top ten as our control set. In practice, our control search variable is a single aggregate search index for any search involving any of these 10 words, standardized in the same manner as described below in the Google Data section.

## Data

### Google Data

Our data on keyword search frequency comes from Google Trends (<http://trends.google.com>), a website which allows users to access data on Google keyword search frequency, stratified by geographic regions ranging from as large as a country to as small as a city. Data for smaller geographic areas can be more difficult to use, since the data are suppressed unless the overall search volume reaches a minimum threshold. For this paper, we use data from the state level in the United States to ensure a wide selection of available search terms, yet still preserve geographic variation in search frequency.

Google Trends does not provide information on the actual number of keyword searches. Rather, it provides a relative search index which reports aggregate searches as a percentage of the most searched term and time period, which is given the value of 100. Therefore, the search index may vary depending on the time frame requested, as well as the set of keywords in a given query. In order to standardize the data, we requested the data for the entire 15-year sample one keyword and state at a time. We then divided the keyword frequency for each word by the state-specific mean, giving us a search frequency index for each state and word with a mean value of 1. Therefore, each search frequency observation can be interpreted as the percentage of the average search frequency for a given word in a given state. For example, in the descriptive statistics in Table 2, the maximum search frequency for the term “unemployment” is reported as 3.72. This means that for that specific month and year in that state, searches were 3.72 times higher than the monthly state average for that word over the 15-year period. One benefit of specifying the search frequency in this fashion is that it is easy to interpret – coefficients in our regression will represent the effect of a doubling of search frequency from the mean on the outcome of interest. Missing, incomplete, or relatively low frequency search data are reported as 0, which we set to missing.

Search data in Google can be specified as either a keyword or topic, which are measured differently. For keywords, Google reports searches with an exact keyword match, while for topic searches, Google utilizes an algorithm which attempts to include other searches which have to do with the keyword, but may not include the actual keyword itself. Unfortunately, this algorithm is proprietary to Google, and therefore is somewhat of a black box.

In Table 1, we report each of our search terms, whether we utilized the keyword or topic search data, and whether the keyword passed a given screen. Topical searches were priori-

Table 2: Descriptive Statistics

Variable	Obs	Mean	S.d.	Min	Max
<b>Nativity Variables</b>					
Births	9,180	6,602	7,776	359	51,747
Ln(Births)	9,180	8.27	1.04	5.94	10.85
<b>Keyword/Topics</b>					
Control Set*	9,180	1	0.16	0.64	1.35
BabyCenter	8,744	1	0.62	0.07	8.22
Birth Control	9,153	1	0.24	0.24	3.88
Clearblue	7,097	1	0.71	0.15	9.03
Conceive	7,916	1	0.60	0.17	12.95
Divorce	9,173	1	0.22	0.21	4.64
Dilation & Curettage	7,917	1	0.65	0.17	11.84
Emergency Contra.	8,725	1	0.43	0.14	8.30
Folate	8,633	1	0.54	0.17	7.90
HCG*	8,641	1	0.51	0.15	6.53
IUD	8,712	1	0.38	0.20	4.83
IVF	8,627	1	0.52	0.12	9.31
Layoff	8,575	1	0.75	0.15	11.00
Medical Abortion	7,286	1	0.61	0.15	9.60
Midwife	8,498	1	0.60	0.14	7.24
Miscarriage	8,927	1	0.35	0.23	6.20
Misoprostol	6,693	1	0.76	0.18	10.81
Missed Period	7,931	1	0.51	0.17	6.83
Morning Sickness	8,424	1	0.49	0.20	7.04
Obgyn	8,618	1	0.41	0.11	4.06
Online Dating*	9,162	1	0.42	0.12	2.31
Ovulation	9,086	1	0.26	0.22	4.02
Ovulation Test	7,198	1	0.59	0.18	9.72
Plan B	8,398	1	0.40	0.11	6.14
Porn	9,180	1	0.35	0.25	2.29
Pregnancy	9,180	1	0.18	0.35	1.96
Pregnancy Bleeding	7,954	1	0.61	0.20	8.90
Pregnancy Symptoms	8,945	1	0.38	0.19	6.28
Pregnancy Test	8,975	1	0.32	0.21	4.80
Pregnant	9,177	1	0.37	0.12	2.12
Prenatal Vitamins	7,838	1	0.55	0.19	6.43
STI	9,064	1	0.54	0.18	8.07
Ultrasound	9,069	1	0.35	0.16	5.43
Unemployment Office	8,283	1	0.61	0.13	9.09
Unemployment	9,155	1	0.54	0.09	3.81

\*See Google Data section for extra details.



tized over keywords when possible. In order to understand these terms, a note on the logic behind the Google Trends search algorithm is necessary. Certain delimiters, such as ", -, and + allow users to change the combinations of keywords searched. A search for a single keyword will yield the search frequency index counting all searches that contain that keyword, including searches which contain other words. For example, the reported search volume for keyword "pregnant" will contain searches for "am I pregnant". If more than one keyword is entered, the resulting Google search volume will contain only searches which contain both words in the same search. In essence, the space delimiter serves as the Boolean logic operator "and". The Boolean "or" operator is "+". For example, the search Emergency Contraception will contain only searches which contain both "Emergency" and "Contraception", while the search Emergency+Contraception will contain all searches which either contain "Emergency" or "Contraception". The "-" operator removes searches containing a specified keyword. For example, one of our search terms mentioned in the paper was "HCG" short for Human chorionic gonadotropin, a hormone produced by a woman's body for the maternal recognition of pregnancy. Testing for the presence of this hormone in urine serves as the basis for many pregnancy tests. However, in the early 2010s, a fad diet called "The HCG Diet" rose then fell in popularity, leading to increases in searches for HCG unrelated to pregnancy. Therefore, we specified our search query as HCG - Weight - Diet, which kept searches for HCG, but removed any search which contained HCG but also included either the word Weight or Diet. Finally, keyword searches in quotations require an exact string match. Returning to our example, the search Emergency Contraception will include all searches which include both the words "Emergency" and "Contraception", independent of where those words appear in the query, whereas the search "Emergency contraception" will only include searches which that exact string. Therefore, a search for "Where do I find contraception in an emergency" would be counted in the former search, but not the latter.

Returning to Table 1, the keywords are listed exactly how they were queried by the system, including the ", -, and + operators. There are two exceptions. The first is the Control Set, which as explained in the Control Keywords section, was specified as the + of + how + a + is + what + and + me + you + do. The other exception is Online Dating. For this query, we found that both the keyword and the topic yielded unbelievably low search volume given the popularity of online dating. We concluded that a better query would be a composite term including combined searches for popular dating sites. Therefore, for our online dating keyword, we used the search term match.com + tinder + okcupid + bumble + zoosk + eHarmony + POF + plenty of fish, reflecting a combination of the seven most popular online dating platforms used over our sample period (Note that POF is merely an abbreviation for the site Plenty of Fish, whose url is pof.com).

## Nativity Data

Data on births by state and month are from the National Vital Statistics System (NVSS) which is part of the National Center for Health Statistics (NCHS). From this system, we get monthly birth counts for each US state and the District of Columbia. We only utilize births in years since 2004 as this is the year Google search data begins. As of this writing, the most recent year of data on births is 2018, giving us a total of 15 years of monthly birth data at the state level, or 9,180 state-month-year observations. In addition, this system allows users to aggregate births by various sub-groupings, such as educational attainment, age of mother, etc. We used these data by these sub-groupings in our heterogeneity analysis.

## Methodology

### Model

Our prediction model relies on linear OLS fixed effects regression. Specifically, we estimate the following prediction model:

$$Y_{smly} = \alpha_{sm} + \sum_{s=1}^2 \gamma_s * t^s + \sum_w \sum_{l=t_0}^T \beta_{w,t-l} I_{smly}^w + \epsilon_{smly}$$

where  $s$ ,  $m$ , and  $y$  index state, month, and year respectively,  $\alpha_{sm}$  is a state-month fixed effect,  $\sum_{s=1}^2 \gamma_s * t^s$  is a state specific time polynomial, and the double summation represents a series of  $\beta$  coefficients for the natural log of the normalized search volume for different keywords at a number of monthly time lags. The dependent variable  $Y$  is the natural log of births, implying the interpretation of the  $\beta$ s are an elasticity – the percentage change in births from each percentage change in search volume. We use  $t_0$  of 7 and a  $T$  of 12, representing monthly time lags from 7-12 months before birth. Huber-White standard errors are utilized. Given the state-month fixed effects, this regression effectively controls for state-specific seasonality in both births and keyword search volumes. In essence, it estimates the effect on births of larger than normal search volume for a given month, in a given state, compared with that same month and state across years. It also controls for changes in aggregate births over time specific to the state due to the linear and quadratic time terms.

For the results on the association between births and google searches at different time lags shown in Figure 1 of the main text, we expand the lagged month parameters from between 7 and 12 months before birth to 1 and 12 months in order to capture more time effects for illustrative purposes. However, since our prediction model is about early prediction, including all this information would only allow us to forecast births one month in advance. As a result, in our main prediction model we only include information on google searches 7 months before birth and earlier.

## Assessing Predictive Power

The significant associations between certain keywords and later births presented in Table 1 in the main text do not necessarily imply that these keywords can effectively be used in prediction. For example, over-fitting is a common and well-known problem in forecasting. Conversely, however, if important explanatory variables are excluded, then even highly predictive variables may not be sufficient to accurately forecast. Therefore, there is generally a trade-off between increasing the number of explanatory variables – thereby increasing the information used in prediction — and over-fitting the data.

To explore a range of keyword sets on this trade-off spectrum, we use a set of keyword screens with varying levels of stringency. Our first, most basic screen is to omit any keyword which has more than 1/3 of the 9,180 state-month observations missing. Our next screen, the Significance screen, takes the first screen and omits any keyword for which both of the following conditions are not true: 1) the month-specific estimated associations are jointly significant at the 5% level for the lagged months 7-12 while also controlling for all other included keywords, and 2) there is at least one month among the lagged months 7-12 which is individually significant at the 5% level. We implement this screen in an iterative process, where all Screen 1 variables are initially included in the regression equation. Keywords are removed if they fail to meet either conditions above, and then the model is re-run with the reduced keyword set and the conditions reapplied until the only words remaining are those for which both conditions are met. We report in Table 3 the results of each of four rounds which were necessary to select the final keyword set according to this methodology, and report which words were dropped at each round. This screen selected 14 keywords for inclusion, listed in alphabetical order: ClearBlue, Divorce, HCG, IVF, Layoff, Morning Sickness, Ovulation, Porn, Pregnancy Test, Pregnancy, Pregnant, STI, Ultrasound, and Unemployment.

For our third screen, we employ a statistical learning methodology which uses out-of-sample tests to determine predictive power, in order to select a keyword set which optimizes predictive power while minimizing the number of keywords. In this methodology, our data set is divided into two groups: a training data set, and a test data set. The training data set is used to estimate the associations between the words and the births, while the test data set is used to test the predictive power of the word estimates, as evaluated by changes in goodness-of-fit measures compared with a model which does not include any keywords as predictors. In the interest of computational feasibility, we use a variation on the leave-p-out cross validation methodology which omits two years of observations at a time, which we call leave-2-years-out (L2YO) cross validation. As we have 15 years of data, this creates 105 possible two-year combinations, and hence 105 model runs to test the predictive power of each keyword set.

We apply this L2YO methodology to select an optimal keyword set from among the Screen 1 keywords, using a process of forward step-wise selection, which is carried out as follows. First, the base model without any keywords is estimated. Then, we add each keyword one at a

Table 3: Significance Screen Word Selection

Word	Round 1	Round 2	Round 3	Round 4
Clearblue	X	X	X	X
Divorce	X	X	X	X
HCG	X	X	X	X
IVF	X	X	X	X
Layoff	X	X	X	X
Morning Sickness	X	X	X	X
Ovulation	X	X	X	X
Porn	X	X	X	X
Pregnancy Test	X	X	X	X
Pregnancy	X	X	X	X
Pregnant	X	X	X	X
STI	X	X	X	X
Ultrasound	X	X	X	X
Unemployment	X	X	X	X
Dilation & Curettage	X	X	X	
Miscarriage	X	X		
Obgyn	X	X		
Unemployment Office	X	X		
BabyCenter	X			
Birth Control	X			
Conceive	X			
Emergency Contra.	X			
Folate	X			
IUD	X			
Medical Abortion	X			
Midwife	X			
Misoprostol	X			
Missed Period	X			
Online Dating	X			
Ovulation Test	X			
Plan B	X			
Pregnancy Bleeding	X			
Pregnancy Symptoms	X			
Prenatal Vitamins	X			

Table 4: MSPE Screen Word Selection

Round	Word	Cumulative $\Delta$ Mean MSPE
1	Unemployment	-20.7%
2	HCG	-23.3%
3	Clearblue	-24.6%
4	Unemployment Office	-25.7%
5	Ovulation	-26.7%
6	STI	-27.5%

time to the model (with the control keyword), and calculate the prediction error as the absolute value in percentage terms of the difference between these predictions and the actual births for every state-month. We then repeat for each of 105 combination of two years of data. We then find the mean squared prediction error (MSPE) across the runs for that keyword set. The keyword which minimizes the MSPE is selected, and becomes part of the base model. Then a second round begins, where each remaining word is added to this new base model one at a time, and the new word which minimizes the MSPE is also selected. This procedure continues until the additional reduction in MSPE is less than one percentage point. In Table 4, we report for each round the word selected and the respective cumulative reduction in MSPE. Employing this screen selects five keywords, listed in order of selection: Unemployment, HCG, Clearblue, Unemployment Office, and Ovulation.

We find that both the Significance keyword set and the MSPE keyword set significantly improve forecast accuracy. To demonstrate this, we plot a moving average of the MSPE by month for each of three keyword sets between 2004-2018 – the model without keywords, the model utilizing the Significance keyword set, and the model with the MSPE keyword set. In Figure 4 we show the average MSPE across all the training data sets by year, and in Figure 5 we do the same for the test data sets. In these figures, we see that both the Significance and MSPE keyword sets lower model prediction error relative to a model without keywords. Significantly, the efficacy of using keywords to predict fertility is especially high during crises – prediction error using both keyword sets during the 2008-2009 financial crisis is approximately half as large as the model without keywords. This is a comforting result for the purposes of our exercise, since we are using these keywords to predict fertility during the COVID-19 crisis.

Figure 4: Training Error





Figure 5: Test Error



## Other Results

### Searches and Fertility

In this section we show how keyword search volume changed during the COVID-19 pandemic. Specifically, in Figure 6 we plot adjusted keyword search volumes from January 1st, 2018 to July 31st, 2020 for a series of keywords. On each plot, we include two vertical lines for the reader's reference. The first line is on January 13, 2020, which corresponds with the estimated arrival date of the first COVID-19 case to the United States. The second vertical line corresponds with March 19th, the date of the first state-wide stay-at-home order in California.

Plotting the raw search data can be misleading since keyword search volumes generally follow seasonal patterns. Therefore, we first filter the data in the following manner: First, we estimate a regression model for each keyword with search volume as the dependent variable, and month and week effects as independent variables. Second, we save the residuals from this regression and add them to the average keyword volume over this time period. Third, we normalize the volume such that February 1st, 2020 equals 100, in order to see the effect of the pandemic as a percentage increase in searches 6 weeks before the stay-at-home orders were implemented.

### Other Robustness and Heterogeneous Effects

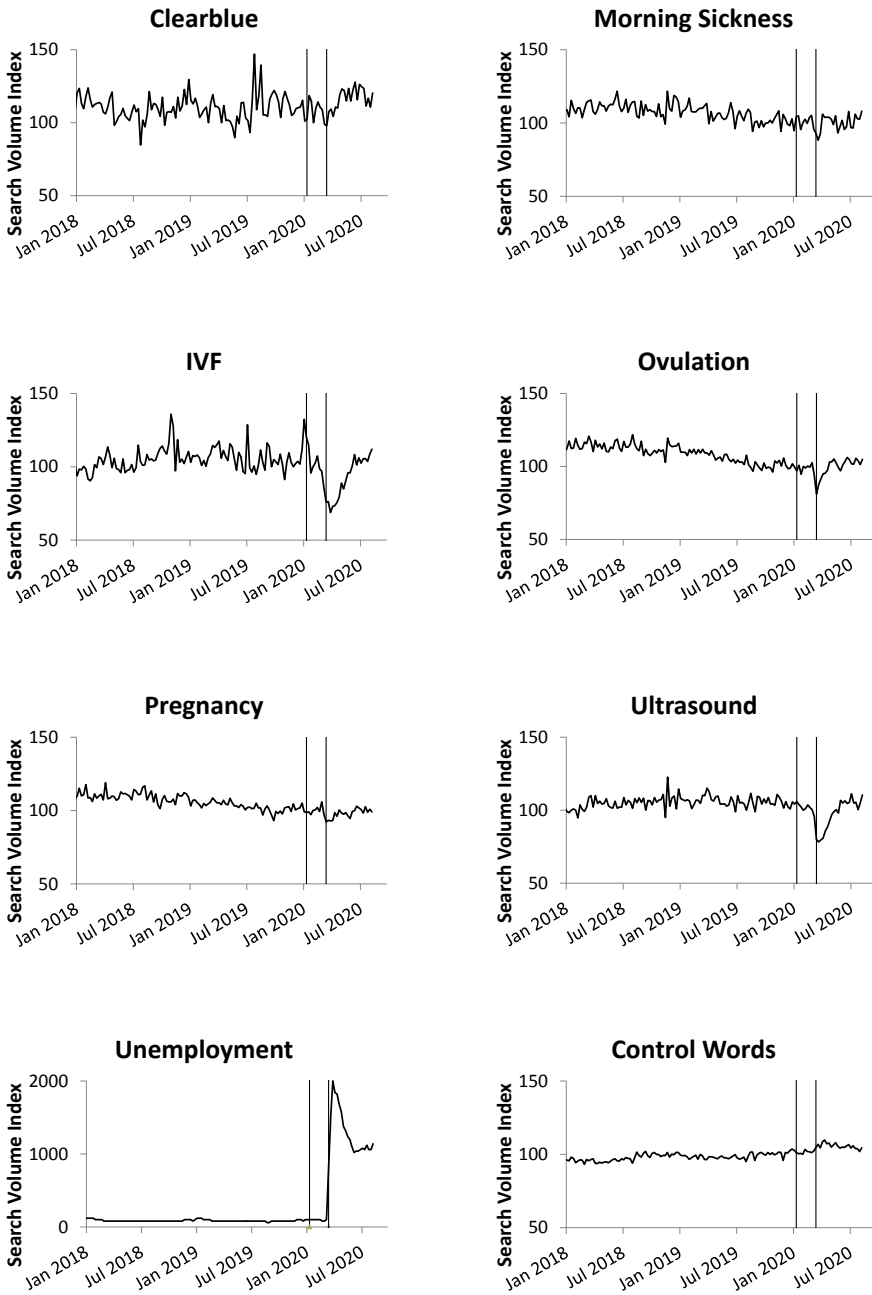
In this section, we present various robustness checks, as well as additional results which were mentioned in the main text. Figure 7 presents our main results across different model lag structures. In the model presented in the main text, we include information on keyword searches from 7 months before birth to 12 months before birth, allowing us to predict births 7 months into the future. In this figure, we show three models, one in which we use 7-12 month keyword lags, another in which we use 8-12 month lags, and finally a third with 9-12 month lags. The model is surprisingly robust across different lag structures. Interestingly, we find that the extended model predicts a small birth rebound beginning in April 2021.

In Figure 3 in the main article, we show heterogeneous birth predictions by educational attainment and ethnic group. In Figure 8 we provide additional predictions by parity and maternal age. We find little evidence of differences by maternal age. However, we do find that first births are predicted to be less affected by the pandemic than higher parity births.

In Figure 9 we show differences in predicted fertility decline across US states. We plot a heat map by state, showing the predicted decline in fertility between October 2020 and February 2021.

In Figure 10 we show a state-level scatter plot of the predicted fertility decline between Oct 2020 and Feb 2021, versus cumulative Covid-19 cases per capita as of Jul 2020. While one might expect that areas with higher caseloads would have larger predicted declines in fertility, we find the opposite. However, this is likely explained by our findings on education – the states

Figure 6: Seasonality-Adjusted Keyword Search Volumes



Covid Economics 54, 29 October 2020: 158-190



Figure 7: Predicted 2020 US Births by Month





Figure 8: Predicted Fertility Decline by Various Mother Characteristics.



Covid Economics 54, 29 October 2020: 158-190

Figure 9: Predicted Fertility Decline between Oct. 2020 and Feb 2021, by State.





Figure 10: Predicted Fertility Decline and Covid-19 Caseload



with the highest COVID-19 incidence were large, urban states, which also have higher levels of education than rural states.

### Estimated Coefficients

Table 5 reports the estimated coefficients from our model with which we predict fertility. The top panel presents the coefficients for the MSPE keyword set, which the bottom panel presents coefficients for the Significance keyword set.

Table 5: Estimated Coefficients Used in Prediction

MSPE Screen Coefficients								
	Unemployment		HCG		Clearblue			
t-7	0.0010	(0.0035)	0.0015	(0.0023)	0.0040***	(0.0014)		
t-8	-0.0134***	(0.0042)	-0.0038*	(0.0022)	0.0058***	(0.0014)		
t-9	-0.0055	(0.0041)	-0.0026	(0.0022)	0.0013	(0.0014)		
t-10	-0.0208***	(0.0041)	-0.0078***	(0.0022)	0.0021	(0.0014)		
t-11	-0.0130***	(0.0040)	-0.0078***	(0.0022)	0.0031**	(0.0014)		
t-12	0.0120***	(0.0035)	-0.0058***	(0.0021)	0.0018	(0.0014)		
	Unemp. Office		Ovulation		Control			
t-7	0.0015	(0.0018)	0.0126***	(0.0044)	0.0472**	(0.0198)		
t-8	-0.0021	(0.0018)	0.0227***	(0.0046)	-0.0373*	(0.0217)		
t-9	-0.0014	(0.0018)	0.0089**	(0.0043)	-0.0320	(0.0206)		
t-10	-0.0040**	(0.0018)	0.0082*	(0.0043)	0.0117	(0.0210)		
t-11	-0.0027	(0.0018)	-0.0095**	(0.0041)	-0.0212	(0.0195)		
t-12	0.0025	(0.0017)	-0.0105**	(0.0043)	0.0219	(0.0165)		
Significance Screen Coefficients								
	Clearblue		Divorce		HCG		IVF	
t-7	0.0041***	(0.0014)	0.0085*	(0.0046)	0.0011	(0.0022)	-0.0010	(0.0022)
t-8	0.0047***	(0.0014)	0.0117***	(0.0045)	-0.0036*	(0.0021)	-0.0002	(0.0021)
t-9	0.0012	(0.0014)	0.0196***	(0.0046)	-0.0042**	(0.0021)	-0.0053**	(0.0023)
t-10	0.0010	(0.0013)	0.0101**	(0.0044)	-0.0077***	(0.0022)	-0.0079***	(0.0022)
t-11	0.0023*	(0.0014)	0.0037	(0.0048)	-0.0066***	(0.0021)	-0.0020	(0.0022)
t-12	0.0004	(0.0014)	0.0123***	(0.0044)	-0.0071***	(0.0021)	0.0003	(0.0020)
	Layoff		Morning Sickness		Ovulation		Porn	
t-7	0.0002	(0.0015)	0.0061***	(0.0020)	0.0137***	(0.0046)	-0.0161	(0.0103)
t-8	0.0005	(0.0016)	0.0052***	(0.0020)	0.0206***	(0.0048)	0.0142	(0.0119)
t-9	0.0023	(0.0016)	0.0017	(0.0020)	0.0087*	(0.0045)	-0.0026	(0.0122)
t-10	-0.0061***	(0.0015)	0.0012	(0.0020)	0.0107**	(0.0045)	0.0361***	(0.0117)
t-11	0.0056***	(0.0015)	0.0004	(0.0019)	-0.0018	(0.0043)	0.0191	(0.0124)
t-12	-0.0040***	(0.0014)	-0.0010	(0.0019)	-0.0052	(0.0043)	-0.0088	(0.0106)
	Preg. Test		Pregnancy		Pregnant		STI	
t-7	0.0040	(0.0039)	0.0317***	(0.0120)	-0.0077	(0.0081)	-0.0064	(0.0041)
t-8	0.0123***	(0.0037)	-0.0022	(0.0121)	-0.0040	(0.0080)	-0.0077*	(0.0041)
t-9	0.0046	(0.0038)	-0.0120	(0.0118)	-0.0037	(0.0077)	-0.0025	(0.0041)
t-10	0.0059	(0.0038)	-0.0468***	(0.0119)	-0.0037	(0.0075)	-0.0027	(0.0041)
t-11	0.0065	(0.0041)	-0.0326***	(0.0123)	-0.0247***	(0.0078)	0.0036	(0.0043)
t-12	0.0086**	(0.0037)	-0.0164	(0.0118)	0.0079	(0.0077)	-0.0001	(0.0037)
	Ultrasound		Unemployment		Control			
t-7	0.0057	(0.0038)	0.0033	(0.0034)	0.0240	(0.0203)		
t-8	-0.0063*	(0.0037)	-0.0150***	(0.0041)	-0.0438**	(0.0221)		
t-9	-0.0083**	(0.0036)	-0.0054	(0.0039)	-0.0377*	(0.0216)		
t-10	-0.0132***	(0.0040)	-0.0213***	(0.0040)	-0.0015	(0.0208)		
t-11	-0.0068*	(0.0035)	-0.0124***	(0.0040)	-0.0426**	(0.0206)		
t-12	-0.0113***	(0.0036)	0.0103***	(0.0036)	0.0232	(0.0184)		

Huber-White Standard errors in parentheses to the right of each estimate. \*\*\* 1%, \*\* 5%, \*10% significance levels.