CEPR PRESS

COVID ECONOMICS
VETTED AND REAL-TIME PAPERS

ISSUE 77
30 APRIL 2021

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.

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

Editorial Board

Ethics

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

Submission to professional journals

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

American Economic Journal, Applied Economics
American Economic Journal, Economic Policy
American Economic Journal, Macroeconomics
American Economic Journal, Microeconomics
American Economic Review
American Economic Review, Insights
American Journal of Health Economics
Canadian Journal of Economics
Econometrica*
Economic Journal
Economics of Disasters and Climate Change
International Economic Review
Journal of Development Economics
Journal of Econometrics*
Journal of Economic Growth

Journal of Economic Theory
Journal of the European Economic Association*
Journal of Finance
Journal of Financial Economics
Journal of Health Economics
Journal of International Economics
Journal of Labor Economics*
Journal of Monetary Economics
Journal of Public Economics
Journal of Public Finance and Public Choice
Journal of Political Economy
Journal of Population Economics
Quarterly Journal of Economics
Review of Corporate Finance Studies*
Review of Economics and Statistics
Review of Economic Studies*
Review of Financial Studies

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

Covid Economics
Vetted and Real-Time Papers

Issue 77, 30 April 2021

Contents

All or nothing? Partial business shutdowns and COVID-19 fatality growth[1]

Matthew Spiegel[2] and Heather Tookes[3]

Using a hand-collected database of partial business closures for all U.S. counties from March through December 2020, we examine the impact of capacity restrictions on fatality growth due to COVID-19. For the restaurant and bar sector, we find that several combinations of partial capacity restrictions are as effective as full shutdowns. Point estimates indicate that, for the average county, limiting restaurants to 25% of capacity and bars to outdoor service reduces the fatality growth rate six weeks ahead by approximately 41%, while completely closing them reduces fatality growth by about 32%. For gyms, we find that, while full closures reduce the COVID-19 fatality growth rate, partial closures may be counterproductive relative to leaving capacity unrestricted. For salons and other personal services, we find mixed evidence that limiting them to 25% of capacity reduces fatalities. However, other constraints are either ineffective or even counterproductive.

The White House declared COVID-19 a national emergency on March 13, 2020. At the same time, state and local government entities began to respond to the public health crisis with a wide range of business shutdowns and restrictions on the general population. Many of these early policies tended to be binary. Establishments were either fully open or completely closed. Given this early response, most of the literature on policy effectiveness has focused on full shutdowns and their impact (e.g., Ling, Wang and Zhou (2020), Chaudhry et al. (2020), Courtemanche et al. (2020), Atkeson et al. (2020), and Spiegel, and Tookes (2020)). Total closures are costly. For example, states that kept restaurants closed longer saw higher unemployment in the hospitality industry in 2020.[1] As the year progressed and cases began to decline in some areas, governments often responded by letting reopen with capacity limits that they often changed over time. This adaptive approach of "loosening and tightening the spigot" seeks to balance policy-makers' desires to protect public health while doing less damage to local economies. But the approach is not without critics.[2,3] How effective are partial shutdowns compared to full closures? This paper attempts to answer that question by analyzing the impact of capacity constraints on restaurants, bars, gyms, and spas on the growth of COVID-19 fatalities.

Overall, we find several combinations of capacity restrictions on restaurants and bars that appear to slow the growth in COVID-related deaths over time. Some of these partial closures appear to be at least as effective as full shutdowns. Among the set of partial restrictions governments have

---

[1] We calculate a correlation of -0.1 between the population weighted number of weeks each state closed restaurants and the 12 month change in employment in the hospitality industry.

[2] For critiques in the popular press, see e.g., Hamblin, J. The Atlantic, (2020, September 9). Paging Dr. Hamblin: Why Didn't America's Shutdowns Work? Many places only half-heartedly gave lockdowns a shot in the spring—and now might have to try them again. See also Board, T. E. (2021, February 10). *Governors Are Easing Restrictions at Exactly the Wrong Time*. https://www.nytimes.com/2021/02/10/opinion/covid-cuomo-indoor-dining.html.

[3] For advocates of partial reopenings see e.g., Ip, G. Wall Street Journal, (2020, August 24). New Thinking on Covid Lockdowns: They're Overly Blunt and Costly Blanket business shutdowns—which the U.S. never tried before this pandemic—led to a deep recession. Economists and health experts say there may be a better way. See also Serkez, Y., New York Times (2020, December 16). The magic number for reducing infections and keeping businesses open, https://www.nytimes.com/interactive/2020/12/16/opinion/coronavirus-shutdown-strategies.html. Serkez argues for 20% capacity limits.

imposed, letting restaurants open at 25% of full capacity while limiting bars to outdoor service appears to be the most effective policy combination that also passes all of our statistical tests. Baseline estimates imply that if restaurants are limited at 25% capacity and bars to outdoor service, a county will see 6-week ahead fatality growth that is 3.2% lower than a county that keeps both restaurants and bars open without any restrictions. This magnitude is meaningful. During the sample period, counties saw a mean growth rate in the weekly fatalities of 7.83% or more. Based on these estimates, limiting restaurants to 25% of capacity and bars to outdoor service will reduce that rate by 41%. By comparison, the estimates imply that closing both types of establishments reduces 6-week ahead fatalities by approximately 32%.

In the case of gyms, the evidence suggests that closing them helps reduce fatality growth relative allowing them to them to open beyond 50% of their capacity. However, intermediate capacity restrictions set at 25% or 50% are either unhelpful or counterproductive (i.e., they are associated with higher future fatality growth relative to allowing them to open at more than 50% capacity, including full capacity). For barbershops, salons, and other personal services (which we define as "spas"), multiple tests produce weak evidence that limiting them to 25% of capacity offers some benefits. However, closing them altogether appears counterproductive, as do capacity limits between 25% and 50%. Overall, our estimates indicate spa capacity restrictions are unhelpful.

It may seem surprising that some capacity restrictions fail to reduce fatality growth due to COVID-19, or that some partial restrictions do more to reduce it than full shutdowns. However, these findings may reflect unintended consequences, where rules meant to reduce risks have the opposite effect. For example, it could be that the population substitutes into riskier activities when restrictions are imposed, or it could be that having rules in place give a false sense of security and that causes citizens to exercise less caution. These types of patterns have also been documented in other settings. Safely guidelines are one example. Risa (2001) finds that mandatory seatbelt legislation increases the rate at which other road users (e.g. pedestrians and bikers) are injured in urban areas. Jones and

Tomcheck (2000) find that pedestrian crosswalks in Los Angles increase the rate at which pedestrians are involved in accidents. In athletics, a meta-analysis by Schneider et al. (2017) finds that protective equipment like headgear and face shields do not affect concussion risk. The mechanism driving this result is unclear, but it may be that athletes engage in riskier play because they feel more protected with safety equipment. In our context, a finding that some of the intermediate capacity restrictions are counterproductive is consistent with the population feeling comfortable visiting these establishments when capacity limits are in place. Conversely, they may be more cautious and endogenously choose to stay at home when no such limits exist. Of course, it is also possible that some results reflect uncontrolled for endogeneity; however, we offer a wide array of controls and tests to try to minimize this issue.

The data that we use in this study are much more granular than in other studies of the effectiveness of non-pharmaceutical policy interventions in the U.S. We begin with the Spiegel-Tookes (2020) database of U.S. county-level policy restrictions for the period March 1, 2020 through December 31, 2020. The database captures full closures and lockdowns, including: general business closures; specific closures targeting bars, restaurants, gyms and spas; no visitation policies at nursing homes; mandatory mask orders; park and beach closures; and limits on the size of gatherings. To examine the impact of partial closures, we add a range of capacity limits on restaurants, bars, gyms and spas. These are the businesses for which capacity restrictions are common and where understanding whether it is possible to limit the spread of COVID-19 while keeping businesses at least partially open is of particular interest. (This policy question is important, even as of late April 2021, because it is likely to take time for sufficiently widespread distribution and uptake of effective vaccines.) For each county or state government order, we record a restriction level and a start date. For restaurants and bars, we categorize capacity restrictions as follows:

- Closed. Indoor and outdoor dining shut; or takeout only;

- Outdoor. Outdoor service only, no indoor service or consumption;
- 25% capacity. Indoor service allowed at up to 25% of capacity;
- 50% capacity. Indoor service allowed at up to 50% of capacity;
- >50% capacity. Indoor service allowed, set at more than 50% of capacity, including completely open (100% capacity).

Gyms and spas have an analogous list of opening levels, but we do not track outdoor regulations (which are rare for these businesses). Instead, in the rare instances in which these businesses are restricted to outdoor service only, we classify them as closed.

The restaurant and bar regulations deserve special mention. During our sample period, restrictions on bars are always as strict as or stricter than those on restaurants. As such, we can only study the impact of restricting bar capacity, given a specific restriction on restaurants. To account for this, we focus the analysis on restaurant-bar restriction pairs.[4]

Potential false positives and false negatives are natural concerns when evaluating the type of policy data examined in this paper. For example, suppose government entities impose a rule near the natural peak of fatality growth. In this case, an ineffective or even somewhat counterproductive restriction can appear to be beneficial – a false positive. False negatives are also possible when effective policies are introduced when fatalities are accelerating. Government officials may introduce policies in response to local factors that accelerate the growth rate in fatalities that are poorly captured by our data or empirical model. For example, suppose that a government agency imposes a restriction using information superior to what our model captures as a means of reducing COVID-19 deaths. If that restriction dampens, but does not reverse, the rate of acceleration in fatalities per week, it might appear counterproductive (because fatalities continue to grow beyond what the empirical model predicts) even though it has actually helped. To help combat these twin problems, this paper takes several approaches.

---

[4] For example, we can compare the impact of limiting bars to outdoor service while restaurants are restricted to 50% capacity to the impact of limiting bars to 25% with restaurant capacity set at 50%. It is never the case that bars are limited to 50% capacity while restaurant capacity is set to 25%.

First, we include a wide array of controls. These include county demographic variables, weather, and a variety of non-business restrictions such as limits on gathering sizes, park closures, and mask policies. Earlier studies have shown that these variables are closely associated with COVID-19's spread. In addition, every test controls for six lags of prior weekly fatality growth, the total number of fatalities to date, the time since the first reported fatality in the county, and the time since March 1, 2020. These all help to identify which policies alter the trajectory of transmission and death versus those that happen to be implemented during its natural rise or fall within a community.

Second, we employ a range of forecast models and compare them to the trend in fatalities at the time a policy is enacted. [5] We begin with a set of policies as of date t as regressors on the COVID-19 fatality growth rate k weeks from now, where k is set to either 4 or 6. We only consider a policy effective or counterproductive if it passes two hurdles: First, within a table, all of the estimated coefficients on the policy of interest have the same sign and at least half are statistically significant. Second, the trend in fatality growth at the time of policy introduction (i.e., during the first and second week after a policy's enactment) are either both insignificantly different from zero or, if significant, the trend at the time of policy introduction has a sign that is the opposite sign of the forecast horizon effect.

Third, we conduct analyses in which we remove each state's most populous counties. The basic idea is that state governments are likely to focus their attention on their more populated areas, leaving less populous areas subject to policies initiative by state governments that may not reflect their current situations. We have seen several cases in the media that are consistent with this view. Counties in Arizona (Ruelas and Oxford, 2020), California (Engelberg and Kaur, 2020), Colorado (Axelrod, 2020) and New York (Sykes, 2020) and others have all objected to state-ordered restrictions on their populations. To the degree that these counties fall under orders optimized in some way for other more populous

---

[5] The forecasts in this paper are in-sample and use the entire dataset. As such, they are not true forecasts. However, since academic studies typically call regressions with a dependent variable k periods out from the regressors forecast regressions adopt the common usage here.

parts of the state, they allow us to estimate the impact of "out-of-equilibrium" policies on future fatalities, thereby mitigating the chances of picking up false positives (e.g., policy introduction at the county's natural peak of fatalities).

Fourth, we conduct near-border tests that are similar in spirit to the standard nearest neighbor analysis. Instead of examining matched pairs of counties that lie on state borders, we require at least a one-county buffer between a near-border county and a state line. This design helps to mitigate potential interpretation problems from spillover effects. Given the way that COVID-19 spreads, a policy that affects transmission in one county is likely to have impact on an adjacent one.[6] The near-border tests help address this issue while also allowing us to control for the possibility that policies simply reflect trends in the trajectory of the virus.

The existing literature on government interventions and COVID-19 fatalities is rapidly expanding. Chaudhry et al. (2020) find that countries that imposed stay-at-home orders were able to reduce future fatalities. In light of evidence that specific businesses like full service restaurants and gyms are tied to the spread of the virus (Chang et al. (2020)), several papers examine the impact of business restrictions. The evidence is that some of these restrictions have aided in the control of Covid-19 (e.g., Courtemanche et al., 2020; Karaivanov et al. (2020); Spiegel and Tookes (2021)) and that restrictions result in reduced population movement (Dave et al. (2020a) and Nguyen et al. (2020)). Still, Atkeson et al. (2020) caution that some empirical approaches may not allow for clear interpretation because of the virus's natural progression. Many recent papers focus on policies introduced at the state level (e.g., Abouck and Heydari, 2020; Friedson et al., 2020; Dave et al., 2020) or they rely on cross-country evidence (e.g. Askitas, Tatsiramos and Verheyden, 2020), where social norms, healthcare infrastructure,

---

[6] For example: different regulations for dining in between Queens County and Nassau County in New York caused some patrons to dine in Nassau County. https://nypost.com/2020/08/31/indoor-dining-rules-send-queens-patrons-over-to-nassau-lawsuit/; differential restrictions near Kansas City state borders may have caused some patrons to go to Kansas https://nypost.com/2020/08/31/indoor-dining-rules-send-queens-patrons-over-to-nassau-lawsuit/. Holtz et al. (2020) find that stay-at-home orders in one area affected movement in others.

and demographics are likely to vary widely. Similar to Spiegel and Tookes (2021), we analyze counties rather than states (or countries) so that we can exploit the granularity of the available fatality data as well as county location and relative size within a state (to improve the overall interpretation of our findings). Relative to existing work, ours is one of the few to study partial restrictions that have been imposed by state and county governments.[7] The analysis in this paper takes a step towards informing policy-makers working to keep the economy at least partially flowing while limiting new deaths.

1   Data

We obtain data from a variety of sources, including those that track COVID-19 fatalities, local weather conditions, and population characteristics. We also collect data on various restrictions that county and state government entities have imposed in efforts to slow the spread of COVID-19. While it might help to include hospitalizations and case counts in the analysis, there are important data limitations. First, systematic hospitalization data for all U.S. counties are generally unavailable for our sample period. Also, while case data are more complete, case counts are a function of testing capacity and availability. These vary significantly, both across regions and over time through at least the fall of 2020. Uneven and incomplete testing is especially problematic in low income and rural areas.[8] Overall, while the fatality data are not perfect (some causes of death are misclassified), it does seem to be the most accurate and consistent measure of COVID-19's spread within local communities that we have.

1.1   COVID-19 Fatality Rate

---

[7] Research on the impact of actual policies that only partially closed establishments is sparse, but several papers run simulations to examine the potential effect of partial closures (e.g., Baqee (2020), Bruaner et al. (2021), Dickens et al. (2020), and Domenico et al. (2021); Schonberger et al. (2020)) or they look at case studies to try to infer best practices, as in Couzin-Frankel (2020). Although it is not the focus on their study, Karaivanov et al. (2020) incorporate data on partial restrictions in Canada.

[8] See, Barry-Jester, Hart and Bluth (2020), Chapman (2020), Najmabadi and Platoff (2020) and Mervosh and Fernandez (2020).

The website USAFacts.org makes available daily COVID-19 fatalities for each county in the U.S. We aggregate these data to create a Wednesday-to-Wednesday series of weekly fatality growth, by county. We use Wednesday fatality counts since weekend reporting delays can produce lower weekend totals and thus higher early week totals (Kissane, 2020) that might be inconsistent from week to week. We convert the weekly fatality counts to weekly growth rates, written as

$$G_{i,t} = \ln\left(\frac{D_{i,t}}{D_{i,t-1}} \right) \qquad (1)$$

where $G_{i,t}$ is the natural log of the growth in COVID-19 deaths in county i in week t and $D_{i,t}$ equals the total deaths during the week. The variable $G_{i,t}$ is the dependent variable in all regressions.

The regressions also include the following controls: six lags of weekly fatality growth; total deaths to date; the time since the county's first reported death; and the number of days since March 1, 2020. Lagged growth controls for serial correlation in the fatality rate. Total deaths to date indicate population's likely level of immunity from those that were infected but survived. We also interact this measure with the lagged growth rates, in case the level of fatalities itself influences the degree which past fatality growth predicts future fatality growth. The days since the first county fatality controls for the total time the virus has been circulating. County populations subject to the virus' impact for longer may behave differently than populations newly experiencing its effects. Finally, the days since March 1, 2020 allows for advancements in medical care that might lower the rate of growth in the number of COVID-19 deaths.

1.2   Demographics

Using the most recently available data from the U.S. Census, we include several population characteristics as controls. These include the fraction of the population that identifies itself as Black, Hispanic, Asian, Native American and Other than White. Age-related variables control for the fraction of

the population over 65 years, the fraction over 85 years, and the fraction of total population residing in nursing homes. Housing and population density are in units per square mile. We also control for county per capita income, as reported by the Bureau of Economic Analysis. Since physical health is correlated with severe disease, the regressions include the fraction of the population that smokes, is obese, or has diabetes. These health-related variables are from the County Health Rankings organization.

1.3 Weather

Local weather conditions may influence COVID-19's spread. Nice weather encourages people to spend more time outdoors, where there is less risk of transmission of COVID-19 (Baker et al. (2020), Carlson et al. (2020) and Quian et al. (2020)). We obtain weather data for every weather station in the National Climatic Data Center for 2020. We use reports from the three stations closest to the county's population centroid and average them to produce estimates for temperature, dew point and rainfall. We include five weather related controls in the regressions. These are: average temperature; hot and humid weekdays; hot and humid weekends; cold weekdays; and cold weekends. Weekdays and weekends are separated to allow for the possibility that weekday weather impacts behavior differently from weekend weather. A day is considered *hot and humid* if the average temperature exceeds 80 degrees and the dew point exceeds 60 degrees Fahrenheit. A day is considered as *cold* if the temperature is below 60 degrees. Using these measures, weekdays have between 0 and 5 hot and humid or cold days. Weekends have 0, 1 or 2 such days.

1.4 Policy Data

We hand-collect collect policy data on a range of actions taken by all state and county governments in the United States. These are the same data in Spiegel-Tookes (2020), but we supplement them to account for partial openings of restaurants, bars, gyms and spas. We introduce five categories of restrictions for restaurants and bars: completely closed, outdoor dining only, indoor up to 25% capacity,

indoor greater than 25% up to 50% capacity, and indoor over 50% of capacity. Gyms and spas have four categories: closed, indoor up to 25% capacity, indoor up to 50% capacity and indoor over 50% capacity. In some cases, governments use population limits (e.g. no more than 50 people indoors at a time) or limits on the number people per square feet. When this was the case, we tried to convert the restrictions into capacity percentages by combining data on average facility size and capacity limits.

As noted in the Introduction, the data on bars and restaurants present a unique challenge. For most business types, the relative strictness of different policies differed substantially across time and counties. Spas in one jurisdiction may have seen their capacities limited to 25% while gyms could fill to 50%. Simultaneously, another jurisdiction might have limited gyms to 25% of capacity while allowing spas to fill to 50%. That was never the case for bars and restaurants. No county ever restricted restaurant capacity more than bars. Thus, all of the tests focus on the impact of a restriction on bar capacity, given a particular restriction on restaurant capacity. Table 1 contains a list of keywords used throughout the paper to describe the data.

1.5 Data Filters

The total database has 178,467 observations and covers the period March 1, 2020 to December 31, 2020. Since prior rates of growth in COVID-19 fatalities are likely predictive of the current rate, we include six lags of weekly fatality growth in all of our regressions. That requires us to drop all date-county observations until six weeks after a county records its first fatality. Although this reduces the database to 67,535 observations (the Baseline Data"), it guarantees that tests of whether a policy alters the trajectory of deaths due to COVID-19 are conducted on areas where the virus is actually present.

In any study of government policy, reverse causality is a potential concern. Do changes in current and projected COVID-19 fatalities lead government officials to impose policies that ultimately do little? Alternatively, do the policies alter the number of future fatalities? We employ tests based on two

sets of filters to help address this issue. The first test drops the five most populous counties in each state from the sample. We call this the Low Population dataset and it has 45,824 observations. This sample is of particular interest because, although all of our analyses focus on policies at the county level, many restrictions come from State Governors' orders. In these cases, elected officials are likely to focus their policy efforts based on concerns about their state's more populated areas. If so, then dates that policies are introduced or lifted will be untethered to projections of COVID-19's spread in a state's more rural counties. Given the number of lawsuits and general objections to state rules by rural officials, this assumption seems reasonable.[9] From the perspective of our tests, removing the state's most populous counties increases the likelihood that, if a policy's enactment is then followed by reduced fatality growth in the low population dataset, the reduction is due to the policy rather than politicians reacting to future forecasts.

The second filter focuses the analysis on a set of matched counties that we use for a modified version of the standard nearest neighbor analysis. A typical nearest neighbor analysis compares county pairs in different states that border each other. One county acts as the treatment area (with the policy) and the other as a control. However, in the setting of virus transmission, analyses of adjacent counties on state borders can produce misleading results. If a policy reduces the rate of spread of COVID-19 in one county, it will likely reduce it in the county next door as well. This can lead to false negatives. Alternatively, a policy imposed in an area with a high rate of infection may induce people to travel to the nearby county to avoid the restriction. That may result in false positives. To mitigate this problem, our neighbor samples only include counties that are nearby, but do not lie along a state's border. We refer

---

[9] Many government officials in the less populated counties across the U.S. view the policies enacted at the state level in this light. See e.g., County of Butler, et al. v. Thomas W. Wolf, et al. (Civil Action No. 2:20-cv-677); "Newsom threatens California counties that defy coronavirus rules as cases spike" San Francisco Chronicle 6/24/2020; "North Texas counties are declaring themselves 100% open for business despite Covid limits," Fort-Worth Star Telegram 6/19/2020; "Governor to Henderson County: Decision to reopen economy will be based on experts, officials", Times-News, 5/4/2020; and "Commissioners to Hogan: Let us reopen" Herald Mail Media, 5/22/2020.

to these as interior counties. For an interior county to enter the database its population centroid must lie within 100 or 200 miles (depending on the version of the filter imposed) of another interior county's population centroid. If there are multiple possible matches, we select the paring that is closest in characteristic space, based on a Euclidean measure. We refer to these as the Neighbor 100 and Neighbor 200 miles samples. These samples have 24,550 and 38,728 observations with average distances between county centroids of 85 and 127 miles respectively.

The distance function that we use to select among possible pairs is

$$d_{ij} = \sqrt{\sum_{k=1}^{n}\left(\frac{h_{ki} - h_{kj}}{\sigma_k}\right)^2} \qquad (2)$$

Where the h_{ki} represent county i's hedonic measure k and σ_k is the standard deviation of the hedonic measure across counties. This implies that a one standard deviation difference in hedonic k between the target county i and another county j is coded as one. The hedonics used in (2) are: per capita income, the fraction of the population over age 85, population density, housing density, weekly temperature, and rain. Matching on distance, along with demographics and weather should produce county pairs with similar infection transmission rates. Also note, that by matching on both hedonics and distance greatly attenuates the degree to which selecting a 100 or 200 mile radius limit matters. As one might expect, counties closer in distance are more likely to be closer in their hedonic attributes as well, As a result, increasing the distance limit from 100 to 200 miles only increases the average distance between counties and their matched pair by about 40 miles.

1.6 Data Summary

Table 2 shows the frequency of restaurant, bar, gym and spa restrictions. The Total column describes the data prior to imposing any filters. The Baseline Data column shows data used in the baseline regressions. Comparing the Total and Baseline Data columns and focusing on the rows labeled

"Closed," the restaurant-bar closed pair goes from about 14% of the data to just over 3%, gyms closures fall from about 18% to under 9% and spas closures drop from about 15% to 5%. Since the only difference between the Total and Baseline data is that the latter only includes the date-county pairs six weeks after a county records its first fatality, it is clear that many businesses were shut down very early in the pandemic, prior to when the first deaths were reported in many areas. As time went by, areas began lifting restrictions and eventually also recorded their first COVID-19 related death. At that point, they enter the Baseline data, with a Table 2 entry in a row other than one labeled Closed.

For whatever reason, locales hardly every restricted bars to outside service while allowing restaurants to fill beyond 50% of capacity at the same time. In the Total, Baseline and Low Population databases there are only eight such observations and in the near neighbors just four. Given this paucity of data, this policy variable does not enter into any of the paper's subsequent tests. At the other extreme, the most common policy response to the pandemic seems to have been letting facilities open to 50% of capacity or more.

Policies likely have to be in place for a few weeks to affect the growth rate in COVID-19 fatalities by a measurable amount. At the same time, if policies never change, it may be difficult to discern if one is superior to another. Table 3 tabulates the length of time various policies (or policy pairs in the case of restaurants and bars) are in place during the sample period. The median time is as short as 1.4 weeks for rules that require outdoor service for bars while allowing restaurants to open beyond 50% of capacity and as long as 15 weeks for, spa capacity limits set at more than 50%. With the exception of the policy pair in which bars limited to outside service while restaurants are open beyond 50%, the range of times should be wide enough to provide useful data on every closure type.

2 Empirical Analysis

The basic regression forecasts the *t+k* period rate of fatality growth based on data as of period *t*. This can be written as

$$100G_{c,t+k} = \alpha + \sum_p \beta_{p,k}D_{c,p,t} + controls \tag{3}$$

where *p* is an indicator for the policy restriction put into place by county *c* at date *t*. The *β* are estimated coefficients and the *D* are policy dummies. Each dummy equals 1 if that particular policy is in place at time *t* and is zero otherwise. If multiple restrictions are imposed during a particular week, the restriction covering the most days is set to 1 and the others to 0. For gyms and spas, capacity limits of over 50% act as the omitted variable. For restaurants and bars, the omitted variable is the one indicating that both can open beyond 50% of capacity.

The tables that follow include results for when *j* is set to 4 and 6, corresponding to the growth rate in fatalities *j* weeks after a policy is put in place. For those infected by the COVID-19 virus and do not survive its affects, the CDC reports a median incubation period from exposure to symptom onset is 4-5 days. Among people with severe disease, the median time to ICU admission from the onset of illness or symptoms ranges from 10-12 days.[10] For patients admitted to the hospital and who do not survive, Lewnard et al. (2020) report a median duration of hospital stay of 12.7 days (ranging from 1.6 to 37.7). Based on these studies, it likely takes a policy four weeks to influence the growth rate in COVID-19 related death.

Interpreting the coefficient estimates from (3) is straightforward. Given the area's demographics, weather, past COVID-19 fatality growth and other policies in place as of date *t*. Then *k*

---

[10] https://www.cdc.gov/coronavirus/2019-ncov/hcp/clinical-guidance-management-patients.html.

weeks later if policy p is in place, then one expects the fatality growth rate to differ by $\beta_{p,k}$ relative to an area without the policy in place.

2.1 Trends when Policies are Introduced

If policy-makers systematically introduce restrictions when the growth rate in COVID-19 related deaths are trending down, regressions can produce false positives. To clarify the interpretation, Table 4 reports the results from estimating equation (3), without a particular policy i in place. The residuals from that regression are then averaged over the first week and then the second week after policy i goes into effect. For example, the Gyms 25% row begins by running the model without the Gyms 25% dummy variable. The residuals from this regression are then collected. The 1 Week-Ahead columns then report on the average across all residuals from the first week after gyms are restricted to 25% of capacity. The 2 Week-Ahead column similarly reports the average residual from the second week after gyms are restricted to 25% of capacity. If new policies are introduced when the current set of policies is leading to a particular trend in the fatality growth rate, then the mean values reported in Table 4 will deviate from zero.

The mean columns in Table 4 indicate that several policies are introduced when the growth rate in COVID-19 fatalities are trending, for the most part upwards. The variables "Bars Closed, Rest Out," "Bars Closed, Rest >50%" and "Gyms Closed" all produce positive statistically significant estimates. This is consistent with the idea that policy-makers see fatalities rising at a faster rate than expected given the set of extent restrictions and then respond by adding more. The only policy that generates negative and statistically significant coefficients in both weeks 1 and 2 is "Bars 50%, Rest >50%." Since this is generally associated with a relaxation of restrictions, it may be that, in this case, fatality growth is better (slower) than expected and the government responds by loosening capacity constraints. Overall, the fact that most estimates are either insignificant or positive should help mitigate potential concerns that a

particular policy restriction appears to be helpful simply because it is introduced when fatalities are already trending downward.

Summary of Findings in Table 4. Fatality trends when policies go into effect.

Negative Trend	No Trend	Positive Trend
Bars 25%, Rest 50%; **Bars 50%, Rest >50%**; Gyms 50%; Spas 50%	Bars Closed, Rest 50%; Bars Out, Rest Out; Bars Out, Rest 25%; Bars 25%, Rest >50%; Bars 50%, Rest 50%; Gyms 25%	Bars Closed, Rest Close; **Bars Closed, Rest Out**; Bars Closed, Rest 25%; **Bars Closed, Rest >50%**; Bars Out, Rest 50%; Bars 25%, Rest 25%; **Gyms Closed**; Spas Closed; Spas 25%

Classification: Negative or positive trend if both the 1 and 2-week ahead estimates are of the same sign. No trend indicates that the estimates have opposite signs. **Bold** indicates both the 1-week and 2-week ahead estimates are statistically significant and have the same sign.

2.2 Baseline and Low Population Regressions: Fatality Growth Rate Forecasts 4 and 6 Weeks Out

Table 5 displays the results from restricting restaurant, bar, gym and spa operations within the baseline and low population databases.[11] The left hand side of the table displays the results for the former and the right hand side for the latter. Standard errors are clustered at the county level in these and all subsequent regressions.

Across policies, restricting bars to outside service when restaurants are permitted to open up to 25% of capacity appears to be the most useful. Among the policies with statistically significant negative growth estimates, the one for Bars Out, Rest 25% have the largest absolute value. The estimates are economically large as well. If the policy is currently in place, the coefficients imply that the fatality growth rate will be 3.2% lower in the baseline data and 3.4% lower in the low population data than if both were allowed to immediately reopen past 50% of capacity. Other policies that produce consistently negative and statistically significant estimates across forecast horizons and databases are restricting bars

---

[11] Interested readers can find the estimated parameters for a model that only includes the control variables in Appendix Table A 2. The estimated control parameters for tables where the row labeled Control as a YES entry are available from the authors.

to outside service while allowing restaurant to open to 50% of capacity and letting both bars and restaurants open to 50% of capacity. The one policy that appears to be counterproductive is letting bars open to 25% of capacity and restaurants to 50%. Based on Table 4, this policy tends to follow surprisingly large growth in rate of COVID-19 fatalities. It is possible its arrival coincides with "lockdown fatigue." In this setting, the capacity constraints may be soft enough that people feel more comfortable mixing, but tight enough that it then forces them to do so inside residences, which may be more conducive to the disease's spread. Of course, these are all hypotheses that require formal testing. Answers will likely require detailed location data. Beyond the restaurant and bar capacity restrictions already discussed, none of the other policies offer consistent results one way or the other and may have ultimately done little to reduce COVID-19's spread.

For gyms, Table 5 indicates that closing them did reduce the COVID-19 fatality growth rate. Both the 4 and 6-week forecast horizons in the baseline and low population data produce negative and statistically significant estimates. Compare this to the large positive and statistically significant estimates for Gyms Closed in Table 4. This implies that the reduced fatality rate estimates from gym closures are not simply reflecting an existing trend. Other constraints, however, appear to be either ineffective or even counterproductive. Based on the estimates, one would conclude that either gyms should be closed or allowed to reopen to over 50% of capacity.

Unlike restaurant, bar and gym closures, closing spas does not appear to be helpful at slowing fatalities due to COVID-19. All four horizon estimates are positive. Taken alone, this would imply that spa closures were counterproductive. However, spa closures were introduced when the fatality growth were abnormally high. Thus, the policies were not as harmful as the coefficient estimates imply.

Summary of Findings in Table 5. Relationships between policy variables and future new fatalities.

	Negative (positive) forecast indicates: (1) Negative (positive) sign across all 4 forecast regressions. (2) At least two estimates are statistically significant. (3) Table 4 has either no trend or a trend in the opposite direction (i.e. positive trend for negative forecast or a negative trend for a positive forecast)
Negative forecast	Bars Closed, Rest Closed; Bars Closed, Rest Out; Bars Closed, Rest 50%; Bars Out, Rest Out; Bars Out, Rest 25%; Bars Out, Rest 50%; Bars 25%, Rest 25%; Bars 50%, Rest 50%; Gyms Closed;
Positive forecast	Bars 25%, Rest 50%; Bars 50%, Rest >50%; Gyms 25%; Spas 50%

2.3   Near Neighbor Forecast Regressions

The paper's near neighbor tests are in Table 6. The left hand side of Table 6 reports the results from the Neighbor 100 mile sample the right hand side from Neighbor 200 mile sample. In addition to all of the regressors used in the earlier tables, the ones reported in Table 6 include dummies for the matching county's policy as well as for the target's.

2.3.1   Near Neighbor Pairings

For the restaurant-bar combinations, the results in Table 6 offer some support for the idea that closing bars and restaurants reduces the growth in COVID-19 fatalities. The 4 and 6-week forecasts are generally negative and significant.

For gyms and spas, the near neighbor tests also reinforce the earlier conclusions. While closing gyms helps, other policies are either ineffective or even counterproductive relative to letting them fully open. For spas, closing them completely again appears to be of little value but limiting them to 25% of capacity may help.

Summary of Findings in Table 6. Relationships between policy variables and future new fatalities.

	Negative (positive) forecast indicates: (1) Negative (positive) sign across all 4 forecast regressions. (2) At least two are statistically significant. (3) Table 4 has either no trend or a trend in the opposite direction (i.e. positive trend for negative forecast or a negative trend for a positive forecast)
Negative forecast	Bars Closed, Rest Closed; Bars Closed, Rest 50%; Bars Out, Rest 25%; Bars Out, Rest 50%; Bars 50%, Rest 50%; Gyms Closed; Spas 25%
Positive forecast	Bars 25%, Rest 50%; Gyms 25%; Spas 50%

3 Conclusion

State and county governments initially responded to the COVID-19 pandemic with policies that fully shut down an array of businesses. Since then, they have adopted a range of partial opening policies intended to help balance desires to support the health of the economy and the health of the population. These partial openings range from outdoor-only service to capacity limits that can vary from 5 to 100%. If full closures are effective, a natural question is whether it is possible to achieve similar results with less stringent policies. This paper uses data on partial openings of restaurants, bars, gyms, and spas to shed light on that question.

For gyms, the evidence indicates that shutting them completely helps, but other restrictions are ineffective or even counterproductive. It is possible that partial openings give people a false sense of security and they ultimately take fewer precautions while working out. Whatever the cause, areas that did not limit gym capacity below 50% saw slower growth in COVID-19 fatality growth than those that did; unless they closed them entirely. When it comes to barbershops, salons and other personal care services, restrictions other than capacity limits of 25% appear to be unhelpful or even counterproductive.

The fact that no county has placed tighter restrictions on restaurants than bars means that we must consider policy pairs for these establishments. We analyze the effectiveness of policies that restrict bars, given a particular restriction on restaurants. We find consistent evidence that some restaurant-bar capacity restriction pairs help reduce COVID-19 deaths. As one might expect, we find that closing them completely slows COVID-19 related fatality growth. However, less restrictive pairs that limit bars to outdoor service appear also appear help. For example, limiting bars to outside service while letting restaurants open to 25% of capacity yields the greatest estimated reduction in COVID-19 fatality growth. And limiting bars and restaurants to 50% of capacity generally works as well or better then closing them completely. Why these pairs seem particularly good at slowing deaths due to COVID-19 is a question we leave open to future research. As the literature examining mobility data progresses (e.g., that analyzed in Harris (2020)), it may eventually explain why some restaurant-bar policy combinations work better than others.

Whenever anyone studies a policy's impact, potential endogeneity issues arise. Our tests include a number of variables and filters that attempt to mitigate this problem. All of our tests allow for trends in fatalities in the six weeks prior to an order's issue, as well as the two weeks afterwards. Moreover, we conduct tests that focus on areas that are more rural and exclude those that are more populous (the latter are likely to be the focus of many policies). We also use a near neighbor pairing to try to further control for unobserved trends. Finally, we only draw conclusions when a result appears consistently, across every one of these filters. Nevertheless, no set of controls or tests can fully eliminate the chance that regression results are coincidental (i.e. due to endogeneity) rather than causal. Our tests are no different. In the end, we can predict whether fatality growth rises or falls after initiating various policies while controlling for a wide range of factors that might support alternative hypotheses.

4   Bibliography

Atkeson, Andrew, Karen Kopesky and Tao Zha, 2020. Four stylized facts about Covid-19, NBER working paper http://www.nber.org/papers/w27719.

Baker Rachel E., Wenchang Yang, Gabriel A. Vecchi, C Jessica E. Metcalf, and Bryan T. Grenfell, 2020. Susceptible supply limits the role of climate in the early SARS-CoV-2 pandemic, Science. 2020 Jul 17; 369(6501):315-319.

Baqaee, David, Emmanuel Farhi, Michael J. Mina and James Stock, 2020. Policies for a second wave. Brookings Papers on Economic Activity Summer 2020.

Barrt-Jester, Ann Maria, Angela Hart and Rachel Bluth, May 3, 2020, California's coronavirus testing still a frustrating patchwork of haves and have-nots, National Public Radio, https://www.npr.org/sections/health-shots/2020/05/03/849243723/californias-coronavirus-testing-still-a-frustrating-patchwork-of-haves-and-have.

Braunder, Jan M., Soren Mindermann, Mrinank Sharma, David Johnston, John Salvatier, Tomas Gavenciak, Anna Stephenson, Gavin Leech, George Altman, Vladimir Mikulik, Alexander John Norman, Joshua Teperowski Monrad, Tamay Besiroglu, Hong Ge, Meghan A. Hartwick, Yee Whye The, Leonid Chindelevitch, Yarin Gal, and Jan Kulveit, 2021. Inferring the effectiveness of government interventions against COVID-19. SCIENCE. 19 FEB 2021

Carlson, Colin J., Anna C. R Gomez, Shweta Bansal and Satie J. Ryan, 2020. Misconceptions about weather and seasonality must not misguide Covid-19 response, Nature Communications 11, 4312.

Chang, Serina, Emma Pierson, Pang Wei Koh, Jaline Gerardin, Beth Redbird, David Grusky and Jure Leskovec, 2020, Mobility network models of COVID-19 explain inequities and inform reopening, *Nature*, https://doi.org/10.1038/s41586-020-2923-3.

Chapman, Sandara, May 1, 2020, Indiana plan to catch up on statewide testing, WTHR, https://www.wthr.com/article/news/investigations/13-investigates/indiana-plan-catch-statewide-testing/531-6f5a24f6-5155-46cf-98d5-c273761cc0f8.

Chaudhry, Rabail, George Dranitsaris, Talha Mubashirc, Justyna Bartoszkoa and Sheila Riazia, 2020, A country level analysis measuring the impact of government actions, country preparedness and socioeconomic factors on COVID-19 mortality and related health outcomes, *EClinicalMedicine* 25, 100464.

Courtemanche, Charles, Joseph Garuccio, Anh Le, Joshua Pinkston, and Aaron Yelowitz, 2020. Strong social distancing measures in the United States reduced the COVID-19 growth rate, Health Affairs 39(7), 1237-1246.

Couzin-Frankel, Jennifer, Gretchen Vogel and Meagan Weiland, 2020. Not open and shut. *Science*17, 241-245.

Dickens, Borame L., Joel R. Koo, Jue Tao Lim, Minah Park, Sharon Quaye, Haoyang Sun, et al., 2020. Modelling lockdown and exit strategies for COVID-19 in Singapore. The Lancet Regional Health, Western Pacific 1.

Domenico, Laura Di, Giualia Pullano, Chiara E. Sabbatini, Pierre-Yves Boelle, and Vittoria Colizza, 2021. Modelling safe protocols for reopening schools during the COVID-19 pandemic in France. Nature Communications, Article 1073.

Engelberg, Isaac and Harmeet Kaur, December 3, 2020, Northern California Sheriff who Refused to Enforce Gov. Newsom's Ordered Restrictions Tests Positive for Covid-19, *The Mercury News*, https://www.mercurynews.com/2020/12/03/northern-california-sheriff-who-refused-to-enforce-gov-newsoms-ordered-restrictions-tests-positive-for-covid-19/.

Harris, Jeffrey, 2020, Geospatial analysis of the September 2020 Coronavirus outbreak at the University of Wisconsin – Madison: Did a cluster of local bars play a critical role?, NBER working paper 28132, http://www.nber.org/papers/w28132.

Holtz, David, Michael Zhaoa, Seth G. Benzellb, Cathy Y. Caoa,b, M. Amin Rahimiana, Jeremy Yanga, Jennifer Allena, Avinash Collisa, Alex Moehringa, Tara Sowrirajanc, Dipayan Ghosha, Yunhao Zhanga, Paramveer S. Dhillone, Christos Nicolaidesf, Dean Ecklesa, and Sinan Arala, 2020, Interdependence and the cost of uncoordinated responses to COVID-19, Proceedings of the National Academy of Sciences of the United States of America 117, 19837-19843, https://doi.org/10.1073/pnas.2009522117.

Jones, Thomas and Patrick Tomcheck, 2000, "Pedestrian accidents in marked and unmarked crosswalks: A quantitative study," ITE Journal, Institute of Transportation Engineers, p.42-46.

Karaivanov, Alexander, Shih En Lu, Hitoshi Shigeoka, Cong Chen and Stephanie Pamplona, 2020. Face masks, public policies and slowing the spread of Covid-19: Evidence from Canada, NBER Working Paper No. 27891.

Kissane, Erin, 11/24/2020, Daily COVID-19 Data is About to Get Weird, *The COVID Tracking Project*, https://covidtracking.com/blog/daily-covid-19-data-is-about-to-get-weird.

Lewnard, Joseph A., Lewnard,Vincent X. Liu, Michael L. Jackson, Mark A. Schmidt, Britta L. Jewell, Jean P. Flores, Chris Jentz, Graham R. Northrup, Ayesha Mahmud, Arthur L. Reingold, Maya Petersen, Nicholas P. Jewell, Scott Young, and Jim Bellows, 2020. Incidence, clinical outcomes, and transmission dynamics of severe coronavirus disease 2019 in California and Washington: prospective cohort study, BMJ; 369:m1923.

Ling, David, Chongyu Wang and Tingyu Zhou, 2020, A first look at the impact of COVID-19 on commercial real estate prices: Asset-level evidence, *The Review of Asset Pricing Studies* 10(4), 669-704.

Mervosh, Sarah and Manny Fernandez, July 6, 2020, Months into virus crisis, U.S. cities still lack testing capacity, New York Times, https://www.nytimes.com/2020/07/06/us/coronavirus-test-shortage.html.

Najmabadi, Shannon and Emma Platoff, June 10, 2020, Outside of Texas' big cities, some are still waiting nearly a month for coronavirus test results from state-run mobile units, The Texas Tribune, https://www.texastribune.org/2020/06/10/texas-coronavirus-testing-mobile-sites/.

Qian, Hua, Te Mao, Li LIU, Xiohong Zheng, Danting Luo, and Yuguo Li, 2020. Indoor transmission of SARS-CoV-2.

Risa, Alf Erling, 1994, Adverse incentives from improved technology: Traffic safety regulation in Norway, *Southern Economic Journal* 60, 844-857.

Ruelas, Richard and Andrew Oxford, May 2, 2020, These Arizona Sheriffs say they Won't Enforce Governor's Stay-at-Home Order. But they Don't Expect a Free-for-All, *AZCentral*, https://www.azcentral.com/story/news/local/arizona-health/2020/05/02/arizona-sheriffs-say-they-wont-enforce-governors-stay-home-order/3070491001/.

Schneider, Daniel, Ravi K. Grandhi, Purnima Bansal, George E. Kuntz IV, Kate E. Webster, Kelsey Logan, Kim D. Barber Foss and Gregory D. Myer, 2017, Current state of concussion prevention strategies: A systematic review and meta-analysis of prospective, controlled studies, *British Journal of Sports Medicine* 51, 1473-1482.

Schonberger, Robert B., Yair J. Listokin, Ian Ayres, Reza Yaesoubi, and Zachary R. Shelley, 2020. Cost benefit analysis of limited reopening relative to a herd immunity strategy or shelter in place for SARS-CoV-2 in the United States, Yale University Working Paper.

Spiegel, Matthew and Heather Tookes, 2021, Business restrictions and COVID fatalities, Yale School of Management working paper available at SSRN: https://ssrn.com/abstract=3725015.

Sykes, Stefan, November 17, 2020, Some New York Sheriffs say they'll Refuse to Enforce Limit on Thanksgiving Gatherings, NBC News, https://www.nbcnews.com/news/us-news/some-new-york-sheriffs-say-they-ll-refuse-enforce-gov-n1248051.

Table 1: Variable Name Key

Keyword	Description
Closed	For restaurants and bars, closed or limited to takeout. For gyms and spas, closed or limited to servicing customers outdoors.
Out	For restaurants and bars, limited to outdoor service. No indoor service permitted.
25%	Facility open, but under an indoor capacity limit of between 1% and 25% of indoor capacity.
50%	Facility open under an indoor capacity limit greater than 25% and less than or equal to 50%.
>50%	Facility open with indoor capacity over 50%, up to and including 100%.
Restaurants X	Restaurants under capacity restriction X (where X is Close, Out, 25%, 50% or >50%)
Bars X, Rest Y	Bar under capacity restriction X (where X is Close, Out, 25% or >50%) and restaurants are simultaneously under capacity restriction Y (where Y is Close, Out, 25%, 50% or >50%).
Gyms X	Gyms under capacity restriction X (where X is Close, 25%, 50% or >50%)
Spas X	Spas under capacity restriction X (where X is Close, 25%, 50% or >50%)

Table 2: Data Summary

Variable names indicate the business type and the associated capacity limit, with *50%+* indicating a limit over 50% (including 100%). *Total* is the entire database. *Baseline Data* is all available county data beginning 6 weeks after the first recorded fatality. *Low Population* is the baseline data after the 5 most populous counties in each state have been dropped. *Neighbor 100* is the baseline data using only counties that are not on the state border and for which a matching non-border county within 100 in another state exists. *Neighbor 200* is the same *as Neighbor 100* but the matching distance is extended to 200 miles.					
Variable	Total	Baseline Data	Low Population	Neighbor 100	Neighbor 200
Bars Closed, Rest Closed	24781 13.89%	2086 3.09%	1665 2.78%	687 2.80%	1137 2.94%
Bars Closed, Rest Out	4489 2.52%	1947 2.88%	1614 2.70%	189 0.77%	1032 2.66%
Bars Closed, Rest 25%	2591 1.45%	854 1.26%	728 1.22%	90 0.37%	444 1.15%
Bars Closed, Rest 50%	13183 7.39%	6116 9.06%	5300 8.86%	1876 7.64%	3669 9.47%
Bars Closed, Rest >50%	7067 3.96%	4117 6.10%	3768 6.30%	675 2.75%	1909 4.93%
Bars Out, Rest Out	4719 2.64%	1775 2.63%	1538 2.57%	700 2.85%	955 2.47%
Bars Out, Rest 25%	418 0.23%	286 0.42%	132 0.22%	78 0.32%	195 0.50%
Bars Out, Rest 50%	4284 2.40%	2081 3.08%	1892 3.16%	827 3.37%	1243 3.21%
Bars Out, Rest >50%	8 0.00%	8 0.01%	8 0.01%	4 0.02%	4 0.01%
Bars 25%, Rest 25%	1923 1.08%	1313 1.94%	1048 1.75%	414 1.69%	573 1.48%
Bars 25%, Rest 50%	2828 1.58%	1357 2.01%	1226 2.05%	423 1.72%	747 1.93%
Bars 25%, Rest >50%	598 0.34%	306 0.45%	270 0.45%	172 0.70%	250 0.65%
Bars 50%, Rest 50%	39419 22.09%	18829 27.88%	16546 27.67%	7567 30.82%	11046 28.52%
Bars 50%, Rest >50%	15091 8.46%	9133 13.52%	8628 14.43%	4441 18.09%	5717 14.76%
Bars >50%, Rest >50%	56558 31.69%	17183 25.44%	15338 25.65%	6376 25.97%	9744 25.16%

Variable	Total	Baseline Data	Low Population	Neighbor 100	Neighbor 200
Gyms Closed	31559	5823	4791	1681	3297
	17.68%	8.62%	8.01%	6.85%	8.51%
Gyms 25%	18528	8155	6713	2085	4523
	10.38%	12.08%	11.22%	8.49%	11.68%
Gyms 50%	53452	27548	24354	10990	15323
	29.95%	40.79%	40.72%	44.77%	39.57%
Gyms >50%	74928	26009	23950	9794	15585
	41.98%	38.51%	40.04%	39.89%	40.24%
Spas Closed	26780	3527	2834	962	1740
	15.01%	5.22%	4.74%	3.92%	4.49%
Spas 25%	13209	4563	3502	714	2150
	7.40%	6.76%	5.86%	2.91%	5.55%
Spas 50%	53846	28512	25493	9312	15960
	30.17%	42.22%	42.62%	37.93%	41.21%
Spas >50%	84632	30933	27979	13562	18878
	47.42%	45.80%	46.78%	55.24%	48.75%
Total Observations	178467	67535	59808	24550	38728

Table 3: Data Percentile Distribution

Panel A displays the number of weeks restrictions are left in place across dates and counties in the Baseline database. Column headers indicate percentiles. Panel B displays the distribution for the growth in the fatality rate for the four samples that we analyze. The fatality growth rate is the dependent variable in all regressions.

	5%	25%	Median	75%	95%	Mean
Panel A: Policy Duration						
Bars Closed, Rest Closed	0.714	3.429	8.286	10.286	25.857	8.578
Bars Closed, Rest Out	1.286	4.786	12.000	19.143	37.857	12.920
Bars Closed, Rest 25%	1.143	8.857	15.143	27.857	33.857	17.477
Bars Closed, Rest 50%	1.286	5.286	9.857	15.143	26.143	11.179
Bars Closed, Rest >50%	1.714	8.714	13.286	18.714	31.143	14.171
Bars Out, Rest Out	0.714	1.714	3.000	5.571	9.571	3.981
Bars Out, Rest 25%	1.857	6.000	11.714	16.714	23.286	11.782
Bars Out, Rest 50%	0.857	3.000	5.857	9.857	13.857	6.831
Bars Out, Rest >50%	0.907	1.000	1.429	2.250	3.650	1.839
Bars 25%, Rest 50%	0.714	2.143	4.714	8.714	15.143	5.749
Bars 25%, Rest >50%	0.714	2.714	13.714	18.286	25.057	12.191
Bars 50%, Rest 50%	0.893	2.000	4.143	6.821	10.000	5.135
Bars 50%, Rest >50%	1.286	4.714	11.286	19.571	29.286	12.779
Bars >50%, Rest >50%	1.714	7.000	13.714	21.286	28.286	14.357
Gyms Closed	2.143	7.571	14.286	22.286	31.286	15.373
Gyms 25%	1.286	6.143	10.714	14.857	22.857	11.075
Gyms 50%	1.143	4.286	9.714	16.286	27.286	11.232
Gyms >50%	1.714	6.714	13.143	20.143	29.143	13.829
Spas Closed	1.571	6.143	12.143	22.143	36.429	14.704
Spas 25%	2.286	8.286	10.714	14.143	25.586	12.094
Spas 50%	1.000	3.643	7.286	18.286	28.286	11.088
Spas >50%	2.286	8.000	15.286	22.714	30.286	15.530
Panel B: Fatality Growth Rate in %						
Baseline Data	0.000	0.000	2.066	9.531	35.307	7.840
Low Population	0.000	0.000	1.560	9.986	37.469	8.001
Neighbor 100	0.000	0.000	1.806	9.531	33.647	7.584
Neighbor 200	0.000	0.000	2.062	9.909	35.667	7.931

Table 4: Residual Fatality Growth near Policy Introductions

This table calculates residuals from a regression of week-ahead change in deaths ($Growth(t+1)$) during the week immediately following the introduction of policy i. Control variables are: current cumulative deaths in the county, lagged changes in deaths per capita, time controls, weather information, and demographic data are included in the regression. We also include all policies that are already in place as of period t from Table 2 other than the newly implemented policy i, where policy i is the policy listed in the first column. Mean$_{t+1}$ denotes the week $t+1$ average change fatality growth times 100. Mean$_{t+2}$ denotes the week t+2 average. *** denotes significance at the 1% level; ** denotes 5% significance; * denotes 10% significance.

	1 Week-Ahead		2 Week-Ahead	
	Mean$_{t+1}$	p-value	Mean$_{t+1}$	p-value
Bars Closed, Rest Close	1.627	0.234	0.791	0.551
Bars Closed, Rest Out	5.779***	0.001	2.783**	0.014
Bars Closed, Rest 25%	1.960	0.522	0.294	0.863
Bars Closed, Rest 50%	-0.134	0.811	2.716***	0.001
Bars Closed, Rest >50%	5.655**	0.032	5.546**	0.010
Bars Out, Rest Out	-0.361	0.565	0.031	0.966
Bars Out, Rest 25%	-0.692	0.603	0.262	0.852
Bars Out, Rest 50%	0.567	0.424	0.778	0.344
Bars 25%, Rest 25%	0.022	0.981	0.239	0.743
Bars 25%, Rest 50%	-1.727*	0.071	-1.737	0.255
Bars 25%, Rest >50%	-0.129	0.850	0.573	0.383
Bars 50%, Rest 50%	-0.548	0.153	0.056	0.896
Bars 50%, Rest >50%	-1.319**	0.014	-1.396*	0.075
Gyms Closed	2.961***	0.000	3.037***	0.001
Gyms 25%	1.492**	0.047	-1.176**	0.018
Gyms 50%	-1.306***	0.000	-0.533	0.151
Spas Closed	4.496***	0.002	2.327	0.154
Spas 25%	4.264***	0.000	1.302	0.106
Spas 50%	-0.400	0.346	-0.046	0.912

Table 5: Baseline Forecast Regression 4 and 6 Weeks Ahead

The growth rate forecast represent regression estimates for the fatality growth rate *j* periods out from the current date *t*. Each explanatory variable in the table is a dummy variable, equal to 1 if that policy is in place on date *t* and 0 otherwise. Capacity limits over 50% (including full openings) are the omitted policies. Lagged fatality growth, current and lagged cumulative fatalities per capita, demographic and weather controls are all included in the regressions, but not reported in the table. Baseline Data estimates include all counties. The Low Population estimates exclude the five most populous counties in each state. Standard errors are clustered at the county level. Significance Key: * 10%; ** 5%; *** 1%.

	Restaurant, Bar, Gym and Spa Estimates							
	Baseline Data				Low Population Counties			
VARIABLES	Mean$_{t+4}$	S.E.	Mean$_{t+6}$	S.E.	Growth$_{t+4}$	S.E.	Growth$_{t+6}$	S.E.
Bars Closed, Rest Closed	-1.921***	0.507	-2.543***	0.520	-2.200***	0.574	-2.626***	0.590
Bars Closed, Rest Out	-0.718	0.477	-1.507***	0.438	-0.569	0.568	-1.504***	0.514
Bars Closed, Rest 25%	-1.034	0.660	-0.264	0.633	-0.984	0.760	-0.247	0.718
Bars Closed, Rest 50%	-1.009***	0.358	-1.514***	0.354	-0.951**	0.406	-1.433***	0.401
Bars Closed, Rest >50%	0.065	0.370	-0.360	0.333	0.135	0.415	-0.255	0.370
Bars Out, Rest Out	-0.176	0.461	-1.127***	0.436	-0.025	0.539	-1.086**	0.508
Bars Out, Rest 25%	-4.120***	0.880	-3.200***	0.915	-5.171***	1.189	-3.446**	1.497
Bars Out, Rest 50%	-1.708***	0.429	-1.880***	0.410	-1.817***	0.483	-2.025***	0.460
Bars 25%, Rest 25%	-0.723	0.493	-1.236**	0.487	-0.824	0.592	-1.494**	0.584
Bars 25%, Rest 50%	2.716***	0.801	4.230***	0.885	3.011***	0.905	4.534***	1.004
Bars 25%, Rest >50%	1.345	1.000	1.031	0.948	1.265	1.138	0.939	1.072
Bars 50%, Rest 50%	-0.952***	0.279	-0.849***	0.268	-1.057***	0.318	-0.958***	0.304
Bars 50%, Rest >50%	0.032	0.299	0.612**	0.283	0.052	0.332	0.699**	0.313
Gyms Closed	-0.933**	0.393	-1.274***	0.385	-0.860**	0.435	-1.217***	0.425
Gyms 25%	0.427	0.357	0.461	0.346	0.814**	0.394	0.896**	0.380
Gyms 50%	0.172	0.260	-0.190	0.241	0.255	0.285	-0.124	0.263
Spas Closed	2.107***	0.373	2.495***	0.374	2.309***	0.421	2.682***	0.423
Spas 25%	-0.305	0.432	-0.491	0.403	-0.673	0.511	-0.862*	0.467
Spas 50%	0.674***	0.231	1.349***	0.219	0.669***	0.252	1.398***	0.237
Observations	67527		67526		59801		59800	
Adjusted R-squared	0.0790		0.0822		0.0759		0.0790	
Control	YES		YES		YES		YES	YES

Table 6: Near Neighbor Regressions

The regressions in this table repeat those in Table 5 with the exception that near neighbor policies are included in the list of controls. For inclusion, counties must not lie on their state's border. In addition, there needs to be a matching non-border county in another state with a population centroid within 100 miles of the target county. Among the set of possible matches the one closest in a multi-dimensional hedonic distance is se0lected based on equation (2). Standard errors are clustered at the county level. Significance Key: * 10%; ** 5%; *** 1%.

VARIABLES	100 Mile Radius				200 Mile Radius			
	Growth$_{t+4}$	S.E.	Growth$_{t+6}$	S.E.	Growth$_{t+4}$	S.E.	Growth$_{t+6}$	S.E.
Bars Closed, Rest Closed	-2.168**	0.916	-2.409***	0.931	-2.268***	0.711	-2.769***	0.710
Bars Closed, Rest Out	2.582***	0.898	2.067**	0.864	-0.841	0.708	-1.679**	0.657
Bars Closed, Rest 25%	-0.270	2.080	-2.161**	1.058	-0.745	0.974	-0.646	0.935
Bars Closed, Rest 50%	-0.352	0.583	-1.268**	0.549	-1.275***	0.454	-1.891***	0.431
Bars Closed, Rest >50%	0.605	0.750	1.295*	0.669	1.060**	0.539	0.464	0.497
Bars Out, Rest Out	1.102	0.887	-0.242	0.796	0.172	0.675	-1.138*	0.629
Bars Out, Rest 25%	-2.788***	1.075	-2.669**	1.139	-4.585***	1.080	-3.577***	1.131
Bars Out, Rest 50%	-1.671**	0.707	-1.765***	0.678	-1.650***	0.598	-2.537***	0.579
Bars 25%, Rest 25%	-0.331	0.884	-0.852	0.890	0.122	0.793	-1.152	0.752
Bars 25%, Rest 50%	4.882***	1.268	5.480***	1.275	3.123***	0.981	3.613***	0.998
Bars 25%, Rest >50%	0.717	1.619	0.544	1.531	1.511	1.183	1.043	1.125
Bars 50%, Rest 50%	-1.303***	0.504	-0.906*	0.473	-0.973**	0.397	-0.989***	0.374
Bars 50%, Rest >50%	-0.328	0.560	0.368	0.520	-0.238	0.422	0.320	0.388
Gyms Closed	-0.604	0.681	-0.742	0.667	-0.865*	0.489	-1.150**	0.487
Gyms 25%	2.297***	0.694	1.613**	0.639	0.744	0.481	0.857*	0.459
Gyms 50%	0.358	0.486	-0.454	0.441	0.218	0.352	-0.097	0.328
Spas Closed	1.384*	0.721	1.977***	0.697	1.834***	0.540	2.539***	0.522
Spas 25%	-3.121***	0.956	-2.744***	0.876	-0.644	0.631	-0.769	0.599
Spas 50%	0.405	0.424	1.410***	0.394	0.523*	0.315	1.378***	0.305
Observations	24547		24547		38725		38725	
Adjusted R-squared	0.0837		0.0879		0.0795		0.0831	
Control	YES		YES		YES		YES	
Near Neighbor Policy	YES		YES		YES		YES	

CEPR PRESS

Table 7: Summary of Coefficients from Table 4, Table 5 and Table 6

The "−" and "+" indicate negative and positive estimated coefficients for the 4 and 6 horizons (respectively) with at least two statistically significant at the 10% level.[12] The "Overall" column includes an icon if two conditions are met, (1) the Table 5 and Table 6 columns have the same icon and (2) the Table 4 column must have either indicate no trend or one of the opposite sign. A * indicates at least one of the estimates is statistically significant.

	Table 4	Table 5	Table 6	**Overall**
Bars Closed Rest Close	+	−	−	−
Bars Closed, Rest Out	+*	−		
Bars Closed, Rest 25%	+			
Bars Closed, Rest 50%		−	−	−
Bars Closed, Rest >50%	+*		+	
Bars Out, Rest Out		−		
Bars Out, Rest 25%		−	−	−
Bars Out, Rest 50%	+	−	−	−
Bars 25%, Rest 25%	+	−		
Bars 25%, Rest 50%	−*	+	+	+
Bars 25%, Rest >50%				
Bars 50%, Rest 50%		−	−	−
Bars 50%, Rest >50%	−*	+		
Gyms Closed	+*	−	−	−
Gyms 25%		+	+	+
Gyms 50%	−*			
Spas Closed	+*	+	+	
Spas 25%	+*	−		
Spas 50%	−	+	+	+

[12] The rule used for the Table 5 and Table 6 columns in this table differ from those used to classify policies in the summary tables. A policy is only included in the summary tables if the Table 4 trend is either no trend or of the opposite sign as the policy forecast. Here we only require the policy forecast to be consistent across regressions. The requirement vise-a-vis Table 4 is only imposed in the Overall column.

5   Appendix

Table A 1: Policy Interventions from Spiegel-Tookes (2021)

Policy Intervention	Description
Stay at Home	"Stay-at-home order" issued by state or county government.
State of Emergency	"State of Emergency" issued by state or county government.
Nursing Home Must Accept Positive	Nursing homes required to accept Covid-19 positive residents.
No Nursing Home Visitation	Nursing home visitors prohibited.
Schools Closed[13]	Schools closed.
Employee masks	Mandatory or recommended face coverings for employees.
Masks recommended in public	Recommended face coverings in public.
Mandatory masks in public	Mandatory face coverings anywhere. This includes policies that mandate face coverings in all public places, as well as those that require masks in a subset of public places.
Beaches and parks closed	Beaches or parks completely closed to the public. Closures must be total; no pedestrian traffic.
No elective procedures	Any elective medical procedures (medical procedures including dental and eye) prohibited.
Gatherings limited to 10	Gathering ban, where gatherings are limited to 10 people.
No gatherings over 100	Gathering ban, where the limit is less than or equal to 100 people, and greater than 10.
No gatherings, limit>100	Gathering ban, where the limit exceeds 100 people.

---

[13] The sample ends on September 1, 2020, just as schools began to reopen. Most schools in the U.S. were closed from sometime in March through the end of August. Thus, the *Schools Closed* variable captures variation in school closures at the onset of the crisis.

Policy Intervention	Description
Risk Level 1 Closed	General business closure policy in effect. Business risk levels are defined in accordance with the reopening phases set by counties. When a county adopts more than 4 phases, we group additional phases according to their proximity to one another in time. If all businesses are open, Risk Level 1, Risk Level 2, Risk Level 3, Risk Level 4 dummies all equal zero. When a general business closure policy is in effect, Risk Level 1, Risk Level 2, Risk Level 3, Risk Level 4 dummies all equal one.
Risk Level 2 Closed	Phase 1 reopening policy in effect, where all but low and medium-risk businesses remained closed. When a county is in Phase 1, the Risk Level 1 dummy equals zero and the dummies for Risk Levels 2, 3, and 4 all equal one.
Risk Level 3 Closed	Phase 2 reopening policy in effect, where higher and highest risk businesses remained closed. When a county is in Phase 2, the Risk Level 1 and 2 dummies equals zero and the dummies for Risk Levels 3 and 4 equal one.
Risk Level 4 Closed	Phase 3 reopening policy in effect, all but the highest risk businesses remain closed. When a county is in Phase 3, the Risk Level 1, 2, and 3 dummies equals zero and the dummy for Risk Levels 4 equals one.
Business re-openings reversed	Phased business reopening reversed

Table A 2: Control Variable Estimates - Baseline Data

Columns represent regression estimates for the fatality growth rate j periods out from the current date t. Each explanatory variable in the table is a dummy variable, equal to 1 if that policy is in place on date t and 0 otherwise. Capacity limits over 50% (including full openings) are omitted. Lagged fatality growth, current and lagged cumulative fatalities per capita, demographic and weather controls are all included in the regressions, but not reported in the table. Standard errors are clustered at the county level. Significance Key: * 10%; ** 5%; *** 1%.

VARIABLES	Growth$_{t+4}$	S.E.	Growth$_{t+6}$	S.E.
Constant	1.587	1.638	0.042	1.587
Deaths Per Capita	-0.057***	0.002	-0.055***	0.002
Growth$_{t-1}$	2.576***	0.634	-0.039	0.505
Growth$_{t-2}$	1.062**	0.538	-1.158**	0.465
Growth$_{t-3}$	-1.065**	0.487	-1.102**	0.461
Growth$_{t-4}$	-1.017**	0.429	-1.331***	0.408
Growth$_{t-5}$	-1.926***	0.370	-1.344***	0.343
Growth$_{t-6}$	-1.380***	0.333	-1.593***	0.290
Int$_{t-1}$	-0.019***	0.007	-0.011*	0.006
Int$_{t-2}$	-0.011**	0.006	0.002	0.005
Int$_{t-3}$	0.009**	0.004	0.008*	0.004
Int$_{t-4}$	0.004	0.004	0.006	0.003
Int$_{t-5}$	0.013***	0.003	0.006**	0.003
Int$_{t-6}$	0.005**	0.003	0.009***	0.002
Days Since First Case	-0.116***	0.013	-0.086***	0.012
Days Since Mar. 1	0.076***	0.021	0.013	0.020
Avg Temperature	0.023*	0.012	0.063***	0.011
Hot Humid Weekday	0.376***	0.084	0.156**	0.079
Hot Humid Weekend	0.499***	0.155	-0.327**	0.151
Cold Weekdays	1.204***	0.073	1.347***	0.068
Cold Weekend	0.030	0.120	0.010	0.111
Age 65+	-0.002	0.028	0.030	0.027
Age 85+	0.372**	0.183	0.252	0.172
Asian	-0.006	0.031	-0.003	0.034
Black	0.015***	0.006	0.011**	0.005
Hispanic	0.060***	0.009	0.072***	0.009
Native Americans	0.030*	0.018	0.017	0.017
Other	-0.043*	0.022	-0.058***	0.021
Per Capita Income	-0.274***	0.071	-0.265***	0.076
Population Density	7.597***	1.586	6.691***	1.483
Diabetes	0.080***	0.024	0.074***	0.023
Obesity	-0.008	0.018	-0.000	0.018
Smokers	0.102***	0.033	0.102***	0.033
Housing Density	-13.089***	3.336	-11.480***	3.167
Nursing Home Pop.	1.770***	0.275	1.602***	0.256

State of Emergency	0.353	0.328	0.195	0.331
Stay at Home	0.616*	0.342	-0.351	0.335
Nursing Accept Pos.	-0.029	0.210	0.061	0.202
No Nursing Visits	0.297*	0.173	0.411**	0.162
Employees Masks	-0.599**	0.241	-0.605***	0.227
Masks Recommended	1.707***	0.304	1.918***	0.268
Mandatory Masks	-0.629***	0.190	-0.571***	0.184
Beaches or Parks Closed	0.086	0.453	-0.067	0.456
No Elective Procedures	-0.281	0.256	-0.252	0.231
Gatherings Limited to 10	-0.536*	0.296	-0.636**	0.277
No Gatherings Over 100	0.776***	0.269	0.371	0.244
No Gatherings Limit>100	0.349	0.283	0.780***	0.269
Risk Level 1 Closed	-0.527	0.589	-0.190	0.576
Risk Level 2 Closed	0.451	0.328	-0.376	0.313
Risk Level 3 Closed	-0.999***	0.203	-0.935***	0.200
Risk Level 4 Closed	0.221	0.168	0.139	0.163
Re-openings Reversed	1.523***	0.299	0.423*	0.246
Observations	67527		67526	
Adjusted R-squared	0.0790		0.0822	

Who doesn't want to be vaccinated? Determinants of vaccine hesitancy during COVID-19[1]

Era Dabla-Norris,[2] Hibah Khan,[3] Frederico Lima[4] and Alexandre Sollaci[5]

Quick vaccine rollouts are crucial for a strong economic recovery, but vaccine hesitancy could prolong the pandemic and the need for social distancing and lockdowns. We use individual-level data from nationally representative surveys developed by YouGov and Imperial College London to empirically examine the determinants of vaccine hesitancy across 17 countries and over time. Vaccine demand depends on demographic features such as age and gender, but also on perceptions about the severity of COVID-19 and side effects of the vaccine, vaccine access, compliance with protective behaviors, overall trust in government, and how information is shared with peers. We then introduce vaccine hesitancy into an extended SIR model to assess its impact on pandemic dynamics. We find that hesitancy can increase COVID-19 infections and deaths significantly if it slows down vaccine rollouts, but has a much smaller impact if all willing adults can be immunized rapidly.

1. Introduction

Our best hope in the fight against COVID-19 and the global economic recovery rests on widespread immunization. So far, vaccine rollouts in many parts of the world have been beset by supply constraints and limited vaccine availability. Yet, even when these issues are resolved, insufficient vaccine demand could still pose a serious challenge. Vaccine hesitancy and refusal could mean that not enough people in a community are immunized above levels required for herd immunity—a threshold that remains an active area of research but that would likely require vaccinating a large majority of the population. Failure to reach herd immunity would halt progress against COVID-19 and place the economic recovery at risk. In this context, it is crucial to understand what drives vaccine hesitancy, and how it could shape pandemic dynamics.

Skepticism towards vaccines, ranging from slight hesitancy to outright refusal, is not unique to COVID-19. Vaccine hesitancy was seen as a growing challenge before the pandemic, even for well-established immunizations with proven track-records of safety and effectiveness like the measles or polio vaccines.[2] Reasons for low vaccine uptake are typically centered on concerns about their safety, potential side effects, and efficacy (Figueiredo et al., 2020), frequently fueled by misinformation or lack of trust in government and health systems (Martinez-Bravo and Stegman, 2021). Vaccination can also be a victim of its own success, as the benefits of widespread vaccination become less salient when disease incidence has been dramatically reduced (e.g., Oster, 2018).

To examine the determinants of vaccine hesitancy during COVID-19, we use individual-level data from the COVID-19 Behavior Tracker, a set of nationally representative surveys conducted by YouGov and Imperial College London across 17 countries between November 2020 and April 2021. We start our analysis by documenting stylized facts about COVID-19 vaccine hesitancy across countries and over time. Next, we investigate empirically how vaccine demand varies across demographic groups, risk attitudes, and self-reported trust in government. In the second part of the paper, building on work by Radzikowski and Dizioli (2021), we extend a canonical susceptible-infectious-recovered (SIR) epidemic model to examine the implications of vaccine hesitancy for pandemic dynamics.

Averaging across our entire sample of over 114,000 observations, only about 61 percent of respondents agree or strongly agree that they would take the COVID-19 vaccine if available to them. However, there is significant variation in vaccine hesitancy across and within countries, and over time. This suggests that even if certain countries or regions achieve herd immunity through

---

[2] The WHO defines vaccine hesitancy as "*the reluctance or refusal to vaccinate despite the availability of vaccines.*" This trend has accelerated in recent years, in tandem with a general decline in trust in institutions, so much so that the WHO named vaccine hesitancy as one of the top ten threats to global health in 2019, alongside climate change. See https://www.who.int/news-room/feature-stories/ten-threats-to-global-health-in-2019 (Jan 18, 2019).

vaccination, neighboring countries or regions may not, which could prolong the pandemic. Vaccine hesitancy also seems to be declining over time, especially in countries that are further along their vaccination programs.

Next, we focus on systematic differences in vaccine intent across demographic groups. Older people, who have the largest morbidity and mortality risks linked to COVID-19, report a much higher willingness to be vaccinated than younger cohorts. Women are also less likely to want to be vaccinated than men. While observed across all age cohorts, the gender gap is largest among working-age women, and could be driven by gender-specific concerns about side effects, differences in access to information, trust in healthcare systems, or risk tolerance (e.g., Croson and Gneezy, 2009; Dabla-Norris et al., 2021).

As in the case of previous vaccines, concerns about the safety and efficacy of the COVID-19 vaccine have the largest direct impacts on vaccination intent. Respondents who strongly believe that the government will provide them with an effective vaccine are almost 50 percentage points more likely to take the vaccine than those who do not. Similarly, respondents who express strong concerns about side effects are 30 percentage points less likely to take the vaccine than those who do not. This suggests that public health policies and communication targeted at informing the public about vaccine safety and effectiveness are key to containing vaccine hesitancy.

The decision to be vaccinated against COVID-19 is also shaped by interactions with peers. People who are more exposed to warnings against the vaccine from family and friends are also less willing to take it, and more likely to share negative information about vaccines with their peers. Quantitatively, a one percentage point increase in the number of friends and family that want the vaccine is associated with a 0.5 percentage point increase in own probability of vaccination, suggesting that vaccine hesitant respondents tend to cluster across the same social networks.

Lastly, we assess the effects of vaccine hesitancy on the number of COVID-19 cases and deaths in an extended SIR model. The model accounts for several relevant features of the COVID-19 pandemic, including endogenous infection rates, asymptomatic transmission, random testing and vaccines. We introduce vaccine hesitancy into the model by changing rollouts in two ways: (i) imposing a cap on the share of population that gets vaccinated; and (ii) reducing the number of vaccinations administered each day, as demand for vaccines fails to meet supply (including due to vaccine shopping) and doses go unused. When both effects are at play, vaccine hesitancy can have a large impact on excess deaths from COVID-19. For example, a counterfactual exercise which increases vaccine hesitancy in the United Kingdom (the lowest in our sample) to the levels observed in France (the highest in our sample) could have worsened the death toll in the United Kingdom by approximately 18,000 deaths between February and July 2021.

In contrast, if governments are able to maintain a rapid pace of vaccine rollouts so that the only effect of vaccine hesitancy is to cap the number of people vaccinated in line with the levels of

vaccine hesitancy measured in our survey, our model predicts that number of excess deaths would be almost 20 times lower. This is because the consequences of a slower rollout compound over time, resulting in much larger effects than what would be expected if the population were quickly vaccinated. The policy benefits of increasing vaccination speed are therefore significant, as it decreases both cumulative deaths, time to herd immunity, and risk of new variants. At the same time, it is also important to take action to reduce hesitancy, especially if the share of population that is willing to be vaccinated is lower than the levels needed for herd immunity.

1.2 Relationship to the Literature

Our paper is part of a growing literature that examines the drivers of behavioral responses to COVID-19, including whether or not to get vaccinated.[3] Using an online survey of 6 European countries conducted in April of 2020, Neumann-Böhme et al. (2020) and Bughin et al. (2021) argue that vaccination preferences are shaped by individual perceptions of benefits and risks, which in turn depend on information from peers and trusted institutions. Lazarus et al. (2021) also find large cross-country heterogeneity in vaccine hesitancy in surveys completed in June of 2020, with higher acceptance rates in countries with stronger trust in government. Much of this research focuses on just a few advanced economies and surveys conducted while COVID-19 vaccines were still being developed. We contribute by analyzing a richer dataset covering 17 countries that begins around the time the first vaccine efficacy results were announced, allowing us to explore how hesitancy has changed as mass vaccination programs were rolled out.

An extensive pre-pandemic literature links lower vaccine take-up with vaccine mistrust, concerns about potential side effects, and inconsistent risk messages from experts and elected officials (Das and Das, 2003; Martinez-Bravo and Stegmann, 2021).[4] These factors are also important for COVID-19 vaccines, with recent work showing that misinformation and conspiracy theories—especially when crafted in pseudo-scientific language—significantly impacts COVID-19 vaccine acceptance (Thunström et al., 2020; Loomba et al., 2021).[5] Hesitancy also appears to be

---

[3] Another strand of the literature looks at compliance with social distancing recommendations, and finds that it is associated with a range of socio-demographic characteristics and personal attitudes (Dabla-Norris et al, 2021, Galasso et al., 2020), as well as other factors such as work flexibility (Papageorge et al., 2020), political beliefs (Allcott et al., 2020; Barrios and Hochberg, 2020), risk preferences (Fan et al., 2020), media choices (Simonov et al., 2020), and civic capital (Barrios et al., 2021).

[4] Omer et al. (2009) and Sadaf et al. (2013) provide useful reviews of the medical literature on this topic, while Dupas and Miguel (2017) offer a useful review of the determinants of broader health care demand, especially in developing countries.

[5] Carrieri et al. (2019) show that misinformation about MMR vaccines led to a reduction in child immunization rates in Italy, especially among households with more online access, while Du et al. (2020) argue that news about manufacturing malpractices was linked to increased hesitancy about vaccination in China. In fact, events that erode trust can impact healthcare demand beyond vaccinations. For instance, Alsan and Wanamaker (2018) show that the disclosure of the Tuskegee study is linked to increases in medical mistrust and mortality and decreases in both outpatient and inpatient physician interactions for older black men in the US.

connected with social media use, and distrust of traditional and authoritative media sources (Murphy et al., 2021). Our results confirm the importance of trust and peers in driving vaccine acceptance, even after controlling for demographics and perceived risk of COVID-19. We also show that policies that increase trust, in particular by assuaging concerns about potential side effects, can have a dramatic impact on vaccine hesitancy, and should be prioritized.

Finally, our paper is related to studies of vaccination policy using epidemiological models.[6] There is a large pre-pandemic body of work examining the impact of vaccination on the dynamics of infectious diseases—Chen and Toxvaerd (2014) and Rowthorn and Toxvaerd (2020) provide useful reviews. We extend the canonical SIR model to include vaccination policy and hesitancy, building on work by Radzikowski and Dizioli (2021), and use this model to study the potential impact of vaccine hesitancy on COVID-19 dynamics.

The rest of the paper proceeds as follows. Section 2 presents the data and discusses stylized facts about COVID-19 vaccine hesitancy. Section 3 describes the empirical approach and Section 4 examines the different drivers of vaccine hesitancy in our data. Section 5 then studies how hesitancy can impact vaccination programs and draws policy implications. The last section concludes.

2. Data and stylized facts

2.1 Data sources

Our main data source is the COVID-19 behavior tracker, a publicly available survey developed by YouGov and the Institute of Global Health Innovation (IGHI) at Imperial College London. Since the start of the pandemic, this tracker has interviewed tens of thousands of people per week across several countries to gather global insights on people's behaviors in response to COVID-19.

The tracker covers a wide range of questions on COVID-19 symptoms, testing, attitudes, and compliance with social distancing recommendations (see Dabla-Norris et al., 2021). Each survey wave is designed to be broadly representative of the general public in each country, and typically has around 1,000 respondents.

For this paper, we focus only on survey waves with information on attitudes towards COVID-19 vaccines. These waves were mostly conducted bi-weekly between November 2020 and April 2021, and span over 114,000 individual observations. Data on vaccine attitudes is available for 17 countries, although with a slight change in country coverage over time, as Finland was replaced

---

[6] A recent literature extends the canonical SIR model to study the interaction between pandemics and economic behavior and social distancing – see, for example, Acemoglu et al. (2020) and Eichenbaum et al. (2021).

by Israel in early 2021, and the United States replaced the Netherlands in March 2021.[7] Details on data availability by country and over time are presented in Figure B.1 in the appendix.

Each respondent is asked a number of questions on COVID-19 vaccines. These include whether respondents would take a COVID-19 vaccine if available (or if they've already been vaccinated), how they perceive vaccine effectiveness and potential side effects, and what obstacles they might face to get their shot. The surveys also ask questions about vaccination attitudes among friends and family, and about sharing and receiving information about vaccines from friends and family and online. Individual responses are self-reported, and coded categorically, either has a binary choice ("Yes" or "No"), or on a sliding scale of agreement or importance. Table A.5 in the appendix provides additional details on some of the questions included in the survey.

Our dataset also includes a wealth of controls, including information on location (state or region), gender, age, health, employment status and occupation, household size, and the number of children in the household. It also includes measures of individual attitudes towards COVID-19, as well as confidence in their government's handling of the COVID-19 crisis and the ability of national health systems to respond to the crisis. Finally, we take data on daily vaccinations from the COVID-19 databases compiled by Our World in Data.

2.2 Vaccine hesitancy across countries and over time

We measure vaccine hesitancy by looking at whether respondents agree, disagree or are unsure about the statement "*I will take the COVID-19 vaccine if it becomes available to me.*" Across our full sample, only 61 percent agree that they will take the COVID-19 vaccine, while 22 percent are unsure, and 17 percent disagree. However, there is significant variation across countries and over time. Figure 1 plots the share of the population that is willing to take the vaccine by country (averaged across all periods in the sample), which ranges from less 40 percent in France to about 77 percent in the United Kingdom. This suggests that simply shoring up the number of available vaccines without simultaneously increasing vaccine demand could be insufficient to boost immunization rates to the levels required to achieve herd immunity in many countries.

---

[7] Our sample covers Australia, Canada, Denmark, Finland, France, Germany, Israel, Italy, Japan, Netherlands, Norway, Singapore, Republic of Korea, Spain, Sweden, United Kingdom, and the United States. The data can be downloaded from https://github.com/YouGov-Data/covid-19-tracker/

Figure 1: Variation in vaccine hesitancy across countries
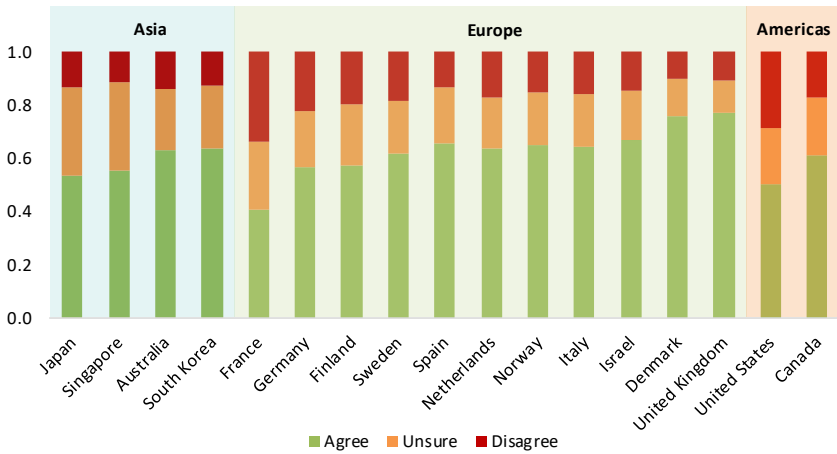
Note: This figure shows average responses to the question of whether respondents agree, disagree or are unsure about the statement "*I will take the COVID-19 vaccine if it becomes available to me.*" We average responses across all waves available for each country.

There is also considerable regional variation in vaccine hesitancy in many countries. To illustrate this point, Figure 2 plots average vaccine hesitancy across German states and regions in the United Kingdom (the nine regions in England, plus Scotland, Wales and Northern Ireland). There are clear and significant regional differences in willingness to take the vaccine in Germany and in the United Kingdom, although the variation is much smaller in the latter. This suggests that even countrywide acceptance of COVID-19 vaccines may not be sufficient when associated with high vaccine rejection rates at the local level. Indeed, if vaccine refusers cluster geographically or share the same social networks, achieving herd immunity is more challenging. This is because the clustering of people that are not immunized can disproportionately increase the percentage of vaccination coverage required to achieve herd immunity in neighboring regions or networks.

It is also important to bear in mind that vaccine hesitancy is not exclusive to the COVID-19 vaccines. In fact, Figure 3 suggests that the level of hesitancy to the COVID-19 vaccine in most countries tracks hesitancy for all other vaccinations. The vertical axis in the figure shows the average share of the population that strongly agrees that vaccines are safe, that they are important and that they are effective. These shares were computed based on data collected by Figueiredo et al. (2020) between 2015 and 2019 and predate the current pandemic. The horizontal axis measures the share of the population who strongly agrees that they will take the COVID-19 vaccine if available. General vaccine hesitancy is strongly and positively correlated with COVID-19 vaccine hesitancy across countries.

Figure 2: Variation in vaccine hesitancy within countries: UK and Germany

UK   Germany
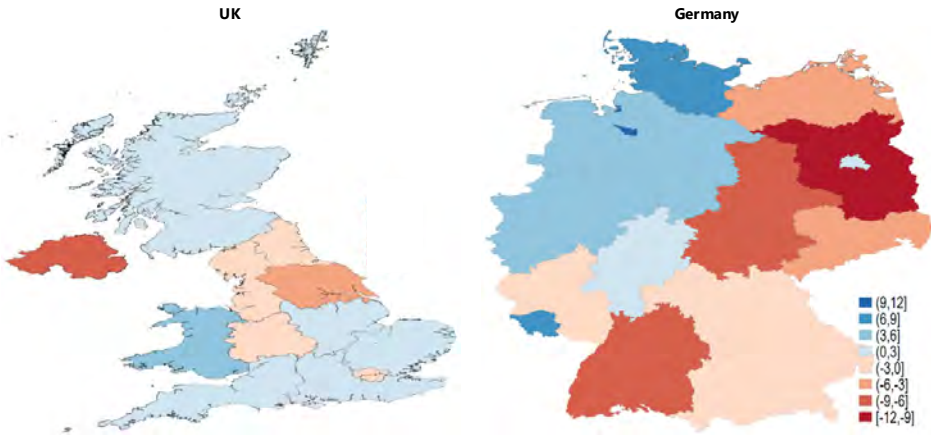
Note: This figure computes share of the population that agrees with the statement *"I will take the COVID-19 vaccine if it becomes available to me"* across regions and plots the deviation from the country mean in percentage points. For example, regions in dark blue are 9-12 p.p. more likely to take the vaccine than the average citizen in the country. We average responses across all waves available for the UK and Germany.

Figure 3: General vaccine acceptance and COVID-19 vaccine intent
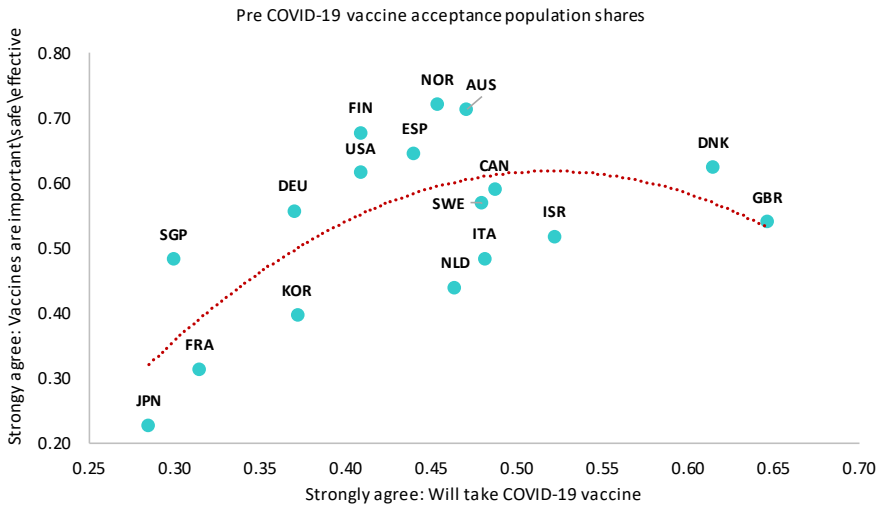
Note: This figure plots the country share of respondents who strongly agree with the statement *"I will take the COVID-19 vaccine if it becomes available to me"* against the share of people who strongly agree that vaccines are important, safe and effective. This second indicator is taken from Figueiredo et al. (2020).

However, Figure 3 also shows that some countries seem to have been able to reduce hesitancy for the COVID-19 vaccine past the level that would be expected given hesitancy for other

vaccines. This is encouraging and suggests that reducing vaccine hesitancy is possible. In fact, and despite the high levels of COVID-19 vaccine hesitancy observed in many countries, our data suggests that skepticism about COVID-19 vaccines has broadly declined over time, with the start of the decline coinciding with vaccine approval and rollouts. The speed of decline is heterogeneous across countries, but the trends shown in Figure 4 suggest that the countries that started vaccination early (in Europe and North America) have experienced a faster decline in hesitancy over time. This suggests that vaccine hesitancy can be lowered and demands further investigation to uncover what actions should be taken to achieve that goal. We turn to this question in the next section.

Figure 4: Trends in COVID-19 vaccine demand over time
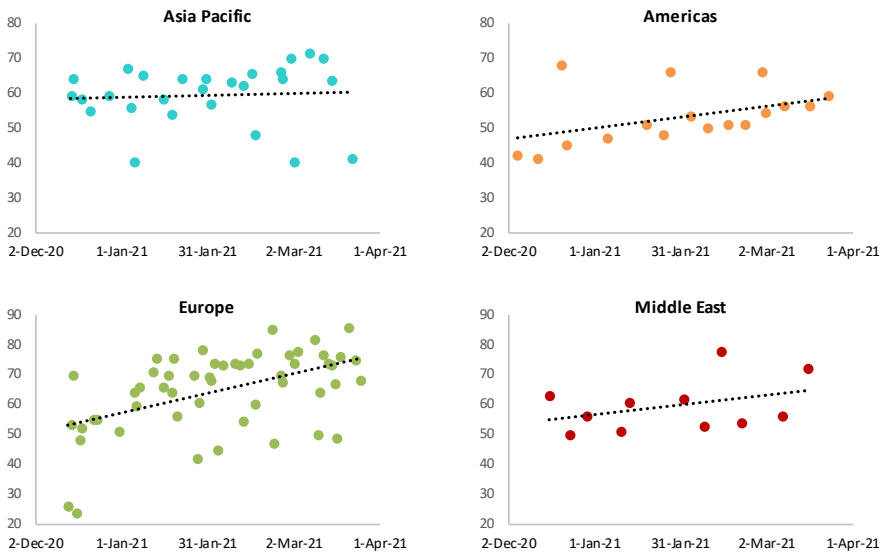
Note: This figure shows the average share of the population that will take the vaccine (or have already done so), plotted against calendar weeks. The regions are defined as follows: Asia Pacific (Australia, China, Hong Kong, India, Indonesia, Malaysia, Philippines, Singapore, Taiwan, Thailand, Vietnam), Americas (Canada, Mexico, USA), Europe (Denmark, Finland, France, Germany, Italy, Norway, Poland, Spain, Sweden, UK) and Middle East (Saudi Arabia, UAE). Data is available at the YouGov COVID-19 Public Monitor: Vaccine Willingness. This is a slightly different dataset, which focuses exclusively on vaccine hesitancy/willingness and thus has a higher coverage than the microdata described in section 2.1.

3. Empirical Model

The baseline empirical model examines the impact of perceptions regarding COVID-19, the vaccine itself, and the response of government and health care systems to the pandemic on vaccination intent. We adopt the following linear probability model:

$$\mathbb{I}\big(\text{take vaccine}_{i(j,t)}\big) = D'_{i(j,t)}\beta_D + A'_{i(j,t)}\beta_A + V'_{i(j,t)}\beta_V + S'_{i(j,t)}\beta_S + G'_{i(j,t)}\beta_G + \delta_{j,t} + \varepsilon_{i(j,t)} \quad (1)$$

where $\mathbb{I}\left(\text{take vaccine}_{i(j,t)}\right)$ is an indicator that assumes a value of 1 if individual i in location j (states or regions within a country) and date t agrees to take the COVID-19 vaccine if available and 0 if she does not.[8] In our main specification (1), we focus on respondents that agree or disagree that they would take the vaccine if available, and exclude respondents that are unsure. In the robustness section, we estimate an alternative model that compares unsure respondents with those that agree to take the vaccine. Our baseline model also favors a linear probability specification, due to the ease in interpreting coefficients and because of the large number of fixed-effects (location-by-time) that are included. The robustness section also discusses an alternative logit specification of the same model. Our findings are consistent across all specifications.

The vector of controls D_i contains demographic variables, consisting of age cohort, gender, household size, number of children, and occupation.[9] We also include indicator variables capturing pre-existing health conditions that may increase COVID-19 risk, and whether individuals have experienced COVID-19-related symptoms.

Next, A_i measures self-reported ease of access to a vaccination site. Some of the reasons indicated by respondents explaining why access might be difficult include: vaccination sites are too far from where they live, or are open in inconvenient times; they are not able to go to a site by themselves; the waiting times at a vaccination are too long; and they fear they might be turned away from a site without receiving the vaccine. Our measure aggregates all those issues into a single dummy variable that indicates whether or not individuals think it would be hard to get a COVID-19 vaccine.

V_i includes the responses to questions about COVID-19 and perceptions about the safety and efficacy of vaccines. These questions ask participants to indicate how worried they are about COVID-19 and about the potential side-effects of the vaccine. They also indicate how confident they are that their government will provide them with an effective vaccine.

Compliance with social distancing recommendations is captured by S_i. For simplicity, we aggregate all questions in that regard into two categories, including compliance with safe behaviors and mask wearing. Safe behaviors include handwashing frequency, and whether individuals follow health authorities' advice and practice social distancing. Mask wearing summarizes how diligent people are about wearing masks outside their home. Since each category includes multiple questions, we summarize the data by assigning a sliding scale to each

---

[8] We assume that that i agrees to be vaccinated if she answered *"4- Agree"* or *"5 – Strongly Agree"* to the question *"If a Covid-19 vaccine becomes available to me in 2021, I definitely intend to get it."* Alternatively, if she answered *"1 – strongly disagree"* or *"2 – disagree"* to that question, the indicator takes the value of zero.

[9] We break age into separate groups (18—24, 25—34, 35—44, 45—54, 55—64, 65—75, and 75+) to capture potential non-linear effects.

answer (with 1 representing "strongly disagrees" and 5 meaning "strongly agrees") and computing the share of the possible total that each individual obtained. This provides an index between 0 and 1 that describes how much each person has adhered to the safety recommendations.

Finally, G_i includes information on the extent to which individual i trusts their government. It consists of two questions: "*how well or badly do you think the government is handling COVID-19*" and "*how much confidence do you have in the healthcare system to respond to COVID-19*". To abstract from time and spatial variation in COVID-19 transmission dynamics and related policy interventions (e.g., stringency of government lockdowns), all regressions condition on location-by-week fixed effects, $\delta_{j,t}$. The results of the estimation are discussed below (see also Table A.1).

4. Results

For ease of exposition, we classify our variables into two groups: *demographic characteristics and vaccine access*, and *beliefs concerning COVID-19, vaccines and trust in government*, which can potentially be affected by policy and/or new information. The results below all use regression model described in equation (1) and include the full set of controls. All standard errors are clustered at the country level.

4.1 Demographic characteristics and Vaccine access

We first examine how vaccine hesitancy varies with age and gender. We estimate the average age profiles in the probability of vaccination among men and women, across all countries in our sample. These profiles are plotted in the left panel of Figure 5.

Age is an important driver of vaccine intent and older people are much more willing to get vaccinated against COVID-19. This is not entirely surprising, since COVID-19 mortality and morbidity increases exponentially with age (Levin et al., 2020; O'Driscoll et al., 2020), and the initial public health messaging on vaccines focused heavily on raising awareness and building trust among older individuals. We would therefore expect vaccine hesitancy to decrease with age, and especially so for the oldest cohorts.

In addition, we find that women are less likely to want to be vaccinated compared to men. This gender gap is observed across all age cohorts, and is consistent with previous findings on hesitancy towards COVID-19 vaccines (e.g., Paul et al., 2021; Neumann-Bohme et al., 2020). The gap largely reflects a higher proportion of women being unsure about their vaccine intent, with smaller gender differences among those who outright reject vaccination. This could have multiple drivers, including specific concerns about side effects among women, gender differences in access to information, trust in healthcare systems or risk aversion. However, it is unlikely to be explained by women perceiving a lower risk of COVID-19, since women in the same survey also

report higher rates of compliance with social distancing recommendations (Dabla-Norris et al., 2021).

Figure 5: Impact of age, gender, and occupation on the probability of vaccination

Note: The left panel shows the impact of age and gender on the probability of vaccination, relative to the cohort of 18 to 24 year old females. The right panel shows the effect on the likelihood of vaccination by employment sector, in percentage points. Coefficients are estimated based on equation (1), including all controls and location and week fixed effects. The 95 percent confidence intervals are computed using robust standard errors clustered by country.

We also examine differences in compliance across employment sectors. We start by restricting the analysis to full and part-time workers, separating between those working from home from those who are not able to telework. Among the latter, we also have information on the sector they work in. We continue to use the specification described in equation (1), and control for age, gender, household composition, and health status. As shown in the right panel of Figure 5, the effect of occupation on the probability of vaccination can be quite large. Controlling for all other characteristics, healthcare workers are the most likely occupation to want to take the vaccine, followed by public servants (essential workers) and the retired. This is potentially good news given that they are the most exposed and the first to be vaccinated in most countries.

Having easy access to a vaccination site also increases the chances that a person gets the vaccine by 4 to 12 percentage points, depending on the model specification. About 55 percent of respondents in our sample report some obstacle to get to a vaccination site. The most common concerns are that the waiting times at vaccination sites are too long (11 percent of respondents), that the sites are too far (6 percent), and that they are open during inconvenient times (4 percent). About 6 percent of respondents also fear that they might be turned away from a vaccination site without receiving a vaccine.

4.2 Beliefs Concerning COVID-19, Vaccines, and Trust in the Government

Next, we turn to the relation between vaccine intent and attitudes towards COVID-19. Anecdotally, concerns about catching COVID-19 play a significant role in determining the probability of vaccination, suggesting that actual or perceived risk of COVID-19 is a major driver of vaccine intent. This is confirmed by the results in the left panel of Figure 6. The probability that a person who strongly agrees with the statement "*I am worried about getting COVID-19*" takes the vaccine is more than 20 percentage points higher than a person who strongly disagrees with it. Similarly, individuals that wear masks and comply with social distancing guidance are also more likely to want to take the vaccine. Each of these variables likely reflects (and partly controls for) risk attitudes and individual beliefs about the severity of the disease.

Figure 6: Probability of vaccination and perceptions of COVID-19
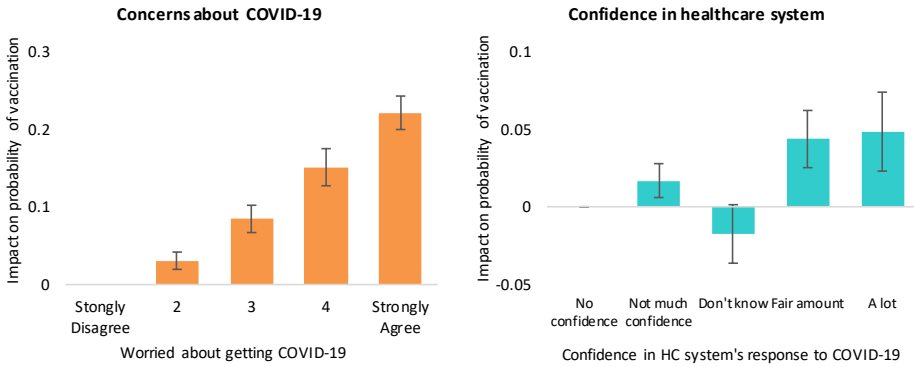
Note: The left panel shows impact of respondents' concerns about COVID-19 on their probability of vaccination. The right panel plots the same effects, this time based on respondents' trust in their nation's healthcare system. Each panel plots the effect on vaccination relative to the lowest category of the survey. Coefficients are estimated based on equation (1), including all controls and location and week fixed effects. The 95 percent confidence intervals are computed using robust standard errors clustered by country.

We also examine how vaccine hesitancy varies with respondents' confidence in the ability of the country's health system ability to respond to COVID-19 crisis. We continue to use the specification in equation (1). Therefore, our results are based on cross-individual variation in self-reported trust, relative to a location-week specific average. This helps to attenuate concerns about omitted variable bias, including from potential differences in the quality of trust questions across countries. The right panel in Figure 6 shows that trust in the capacity of the healthcare system to respond to COVID-19 increases the probability of vaccination by up to 5 percentage points.

Finally, we plot in Figure 7 the impact of the two variables with the largest effects on the probability of vaccination against COVID-19: concerns about the vaccine side effects and whether people believe that their government will provide them with an effective vaccine. These variables have been previously found to impact demand for other vaccines, so it is not surprising that they explain a significant share of the variation in COVID-19 vaccine intent across individuals (Das and

Das, 2003; Martinez-Bravo and Stegmann, 2021). However, the magnitude of the impacts on vaccine hesitancy during COVID-19 is notable. Widespread concerns about side effects reduces vaccination intent by about 30 percentage points, conditional on controls, while strong trust that the government will provide an effective vaccine increases vaccine demand by almost 50 percentage points relative to those with no trust at all. Importantly, both variables can potentially be shaped by public health policies and communication to inform the public about existing evidence on vaccines.

Figure 7: Probability of vaccination, side effects and effectiveness
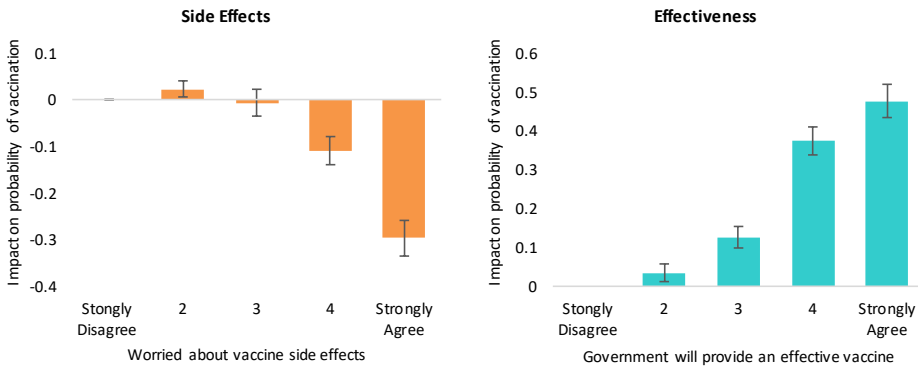
Note: The left panel shows the impact of respondent's concerns about vaccine side effects on their probability of vaccination; the right panel shows the same effects based on the confidence that their government will provide an effective vaccine. Each panel plots the effect on vaccination relative to the lowest category of the survey. Coefficients are estimated based on equation (1), including all controls and location and week fixed effects. The 95 percent confidence intervals are computed using robust standard errors clustered by country.

4.3 Peer Effects and Information

Another factor that can potentially influence an individual's decision to be vaccinated against COVID-19 is whether their peers are doing so as well. Evidence suggests that there are significant peer effects on vaccination through various mechanisms such as information sharing and imitation (Rao 2007; Bodine-Baron et al., 2013; Sato and Takasaki 2019). One particularly important mechanism in the context of the COVID-19 vaccines is the quantity and quality of information that people have access to. People that are more exposed to warnings against the vaccine from their peers could also be less willing take the shot. Conversely, someone who has reservations against the vaccine could also be more likely to look for negative information about in and share it with their peers. The results plotted in Figure 8 suggest that both these channels seem to play a role on vaccine hesitancy.

Vaccine rollouts could also influence uptake. In countries with faster rollouts, the likelihood that people know someone who has received the vaccine is higher, and network effects could play a role in shaping intent. In addition, faster vaccine rollouts provide more data on vaccine side

CEPR PRESS

effects. If the number of people experiencing adverse side effects is consistently small, concerns about uptake could taper off.

Figure 8: Misinformation and vaccine hesitancy
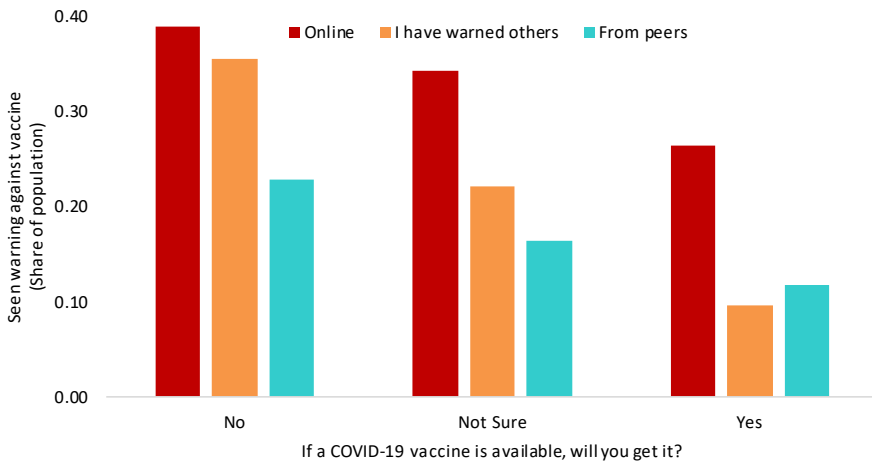
Note: This figure plots the average share of respondents who have seen warnings against the COVID-19 from different sources (online or peers), or that have shared similar warnings themselves. The shares are plotted separately depending on whether respondents agree, disagree or are unsure about taking a COVID-19 vaccine themselves if available.

To measure the effect of the vaccination rollout, we include the log of vaccinations per 100 people in individual i's country at date t into equation (1). We also change the fixed effects from location-by-week to location *and* week $(\delta_j + \delta_t)$ to capture the impact rollouts on vaccine intent. Measuring the impact from peer effects on the probability of vaccination is more challenging, as the network of one's peers is most likely endogenous. One naïve approach is to include the percentage of each respondent's close friends and family that will take the vaccine as a measure of peer effects (the endogeneity issue comes from the fact that individual i wanting to be vaccinated makes it more likely that their friends and family are like-minded).

Table A.2 in the appendix shows the results of this regressions. Columns (1) – (4) mimic Table A.1, with the addition of the vaccination rollout as a control. Note that the rollout has a positive and significant effect on the probability of vaccination in columns (1) and (2) which only include demographic variables. Once concerns about side effects and the effectiveness of the vaccines are controlled for (columns 3 and 4), the effect of the rollout becomes considerably smaller and is no longer statistically different from zero. The effects of the other variables are largely in line with Table A.1.

Columns (5) of Table A.2 includes the share of friends and family than plan to take the vaccine as a measure of peer effects. In this specification, peer effects are strongly correlated with the probability of vaccination: a one percentage point increase in the share of friends/family that will take the vaccine is associated with a 0.5 percentage point increase in the probability of vaccination. However, this coefficient could be biased upwards for the reasons mentioned above.

Comparing columns (4) and (5) in Table A.2 we find that peer effects as an explanatory variable attenuates the size of other coefficients, but the effects retain their directions. While this result is hard to interpret, given the endogeneity of peer effects noted above, it suggests that the effects that were found in the previous section are qualitatively robust to network and peer effects.

4.4 Robustness

Uncertainty about vaccines. Individuals that do not want a vaccine and those who are unsure about getting one if available could have different motivations. As a result, we estimate two separate models. The baseline (discussed above) compares people who will take the vaccine with those that will not. An alternative model compares those who will take the vaccine with those that are unsure about it. Both models are described by equation (1), with the only difference being how we define the dependent (left-hand side) variable.

Qualitatively, the results in both models are very similar. Quantitatively, the effects estimated in the alternative unsure model tend to be smaller. This suggests that the same concerns determine vaccine hesitancy for both the unsure and the people who refuse vaccines. The difference between them is the extent to which these concerns outweigh the potential benefits of the vaccine. The coefficients for the alternative model are reported in Table A.3.

Latent Variable Logit Model. In our main specification, we favor the linear probability model due to the ease of interpretating its coefficients. One drawback of this model is that the predicted probabilities of vaccination can be smaller than zero or greater than one (although this happens for only 4 percent of observations). As an alternative, we adopt a "latent variable" interpretation of the data and construct a logit model for the probability of vaccination.

Suppose that individual i computes the net benefit of taking a vaccine, y_i^*, based on the following equation

$$y_{i(j,t)}^* = D_{i(j,t)}' \beta_D + A_{i(j,t)}' \beta_A + V_{i(j,t)}' \beta_V + S_{i(j,t)}' \beta_S + G_{i(j,t)}' \beta_G + \delta_j + \delta_t + v_{i(j,t)}$$

where the matrices on the RHS represent the same data as in equation (1) and v_i has a standard logistic distribution.[10] In this case, individual i will take the vaccine if $y_i^* \geq 0$ and will not take the vaccine if $y_i^* < 0$. We do not directly observe y_i^*, but $\mathbb{I}(y_i^* \geq 0)$ is known and can be used to estimate the coefficients in the equation above. Table A.4 in the appendix presents the odds ratios associated with the variables in our model. Once again, the results are consistent with our baseline.

4.5 Policies to Decrease Vaccine Hesitancy

We start this section by pointing out that some of the variables included in our empirical model could be endogenous. Our estimates, therefore, might not reflect causal effects and should be interpreted with this caveat in mind. Nevertheless, the regression coefficients are consistent across various model specifications and the inclusion of numerous control variables, including location-week fixed effects to control for country- and region-specific changes in the propensity of vaccination week-by-week. As such, we believe they can still be informative some of the underlying reasons for vaccine hesitancy.

As mentioned above, one of the variables with the largest impact on the propensity to take the COVID-19 vaccine is concern about the potential side effects of the vaccine. It is not uncommon to hear concerns about that vaccines were "rushed" and that more testing is needed. In addition, a large share of the available information that cautions against the use of the vaccine (especially online) is either at odds with scientific evidence or overemphasizes potential side effects.[11] This type of information can lower the willingness to get the vaccine. It is important, therefore, that health authorities accurately and repeatedly inform the public about the evidence on the safety of vaccines, and address some of their salient concerns. This can have positive peer effects as well. For instance, if a person's vaccination decision is positively influenced by his or her peers' vaccination behavior, interventions to promote vaccine take-up among selected individuals not only directly encourage their own take-up but also indirectly encourage take-up among peers.

The second variable with a strong relationship to the propensity of vaccination is whether or not people think that their government will provide them with an effective vaccine. This question

---

[10] We also adopt "location plus time" fixed effects, instead of "location by time". This is done to reduce the number of incidental parameters (dummies) that need to be estimated in the model, which could bias our coefficients.

[11] Note that this is not an issue restricted to on-line message boards or chat groups. In first half of March 2021, more than a dozen European countries temporarily halted the use of one COVID-19 vaccine because of a fear that it might cause potentially fatal blood clots in rare cases. This decision was made despite statements by the World Health Organization and the European Medicines Agency that there was no reason to suspect that the vaccine was unsafe. Source: The Economist, March 2021. "EU countries pause AstraZeneca's covid-19 jab over safety fears." url: https://www.economist.com/science-and-technology/2021/03/15/eu-countries-pause-astrazenecas-covid-19-jab-over-safety-fears, accessed on March 23, 2021.

speaks to at least three issues: trust that the government will be able to secure and distribute a vaccine; that the vaccine will be provided to them (and at reasonable cost); and that the vaccine will be effective in tackling infection and transmission of the COVID-19 virus (Figure 9).

Figure 9: Trust in government and vaccine effectiveness

Country authorities will provide me with an effective vaccine

Once again, these correlations suggest that one effective way to reduce vaccine hesitancy is for health authorities to: (i) credibly and regularly inform the public about the vaccine's efficacy as new information becomes available; and (ii) keep updated information about the vaccination rollout (when and where can each person expect to get the shot). Experimental work has also shown the effectiveness of cues and nudges (Milkman et al., 2011) or increased accessibility (Brewer et al. 2017), particularly among those planning to be vaccinated. For those unwilling to be vaccinated, some recent evidence suggests a role for communicating information by diverse individuals, including in terms of expertise (Alsan and Eichmeyer, 2021). By and large, these strategies have already been adopted in some form or another in most countries that have started to vaccinate their population.[12] Our results and anecdotal evidence given by declining hesitancy (Figure 4) suggest that these approaches could be fruitful.

---

[12] This includes the public vaccination of politicians and celebrities, appearances by health experts on various broadcasts (from news casts to YouTube videos) to clarify some of the main misunderstandings, and easy-to-access and up-to-date web pages in the websites of major hospitals and healthcare organizations containing information on COVID-19 and the vaccines.

5. Vaccine Hesitancy and Pandemic Dynamics

Vaccination is not a goal on itself, but a means to protect people against disease. As such, it is important to gauge how vaccine hesitancy can affect vaccine rollouts, and therefore the speed at which a country's population gains immunity and the consequences in terms of the number of preventable deaths and infections.

5.1 An Extended SIR Model

To examine the impact of behavioral choices of individuals and how they interact with the vaccine rollout, we make use of the extended SIR model developed by Radzikowski and Dizioli (2021). The basic framework divides the population (N) into susceptible (S), infectious (I), recovered (R), quarantined (Q), and dead (D). The "quarantined" category is adopted to accommodate asymptomatic transmission and allow for random testing in the model. The difference between people who are infectious and people who are quarantined is that those in quarantine have either developed symptoms or tested positive for COVID-19. As a result, they do not have contact with (and thus do not infect) others.

At each point in time, the population can be partitioned into the five categories mentioned above, so that $N = S_t + I_t + R_t + Q_t + D_t$. When a share q of the population is tested each day, the laws of motion are:

$$\frac{dS_t}{dt} = \underbrace{-\beta S_t I_t / N}_{new\ infections}\ ;$$

$$\frac{dI_t}{dt} = \beta S_t I_t / N - \underbrace{\frac{\gamma}{1-q} I_t}_{discover\ infection}\ ;$$

$$\frac{dQ_t}{dt} = \underbrace{-\theta Q_t}_{resolving\ infections} + \frac{\gamma}{1-q} I_t;$$

$$\frac{dR_t}{dt} = \underbrace{(1-\delta)\theta Q_t}_{recovered}\ ;$$

$$\frac{dD_t}{dt} = \underbrace{\delta \theta R_t}_{dead}.$$

The parameter β measures the rate of infection, γ is the rate at which symptoms develop, θ is the time it takes to recover from an infection, and δ is the probability of death. The expression $\gamma/(1-q)$ is the rate at which an infected individual discovers that he/she is infected. The reasoning is as follows: once a person is infected, symptoms develop at a Poisson rate γ, which means that the average number of days until the first symptoms appear is $1/\gamma$. When a fraction q of the population is tested each day, the average number of days until a person either develops

symptoms or receives a positive result is $q \times 0 + (1 - q) \times (1/\gamma)$. Inverting this expression gives the rate at which infections are discovered.

Following Radzikowski and Dizioli (2021), this framework is extended in several ways, which we discuss below.

Endogenous rate of infection. The rate of infection can change based on behavioral patterns:

$$\beta_t = n_t \times inf_t \times \omega_t,$$

where n_t is the average number of contacts with other people per day (affected by lockdowns and social distancing), inf_t is the probability of infection (which can be reduced by wearing masks, washing hands, etc.) and ω_t that adjusts the scale and controls for seasonal differences in the rate of infections. The number of contacts per day is defined as:

$$ln(n_t) = ln(n_0) - m_t$$

where n_0 is the initial number of contacts and m_t is relative mobility in day t, measured in relation to the pre-pandemic baseline. Similarly,

$$inf_t = \beta_0\, e^{-\lambda t} + \beta_s\big(1 - e^{-\lambda t}\big)$$

where β_0 is the initial rate of infection, β_s is the rate of infection when safety precautions are taken, and λ measures the time to make that transition.

Population subgroups and mutant strains. The model also allows for the population to be subdivided into vulnerable and non-vulnerable (young), who differ in their probabilities of death if infected. This is done by differentiating between the vulnerable and the young in each of the five categories of the population. Importantly, the categories remain related. For example, a susceptible vulnerable person can still be infected by a young person. The differentiation between the vulnerable and young is relevant to capture the dynamics of the pandemics, especially policies that prioritize one of these groups are implemented (e.g., vaccines).

Equally important to capture the dynamics of COVID-19 is the introduction of mutant strains, which are potentially more infectious to humans than the wild virus. Including this into the model requires accounting for infections from both types of the virus, which is reflected by including extra terms for the mutant variant in the law of motion of each category in the population. For example, the law of motion for the susceptible becomes:

$$\frac{dS_t^v}{dt} = -\beta_t\, S_t^v\big(I_t^v + I_t^y\big)/N - \beta_t^n S_t^v\big(I_t^{v,n} + I_t^{y,n}\big)/N$$

where I_t^v is the number of infectious vulnerable people that are infected with the wild virus and $I_t^{v,n}$ is the number of infectious vulnerable people that are infected with the new (mutant) strain

of the virus (similarly, these categories for the young are represented by I_t^y and $I_t^{y,n}$, respectively). β_t^n is the rate of infection for the new strain, defined as

$$\beta_t^n = \mu \times \beta_s \times n_t \times \omega_t,$$

where β_s, n_t and ω_t are as defined above, and μ is a constant that captures how much more infectious the new strain is, relative to the wild one.

Vaccines. Vaccination can benefit a susceptible person by generating an immune response that prevents illness/death if they are infected with the virus. It may also reduce the probability that this person, if infected, transmits the virus to other susceptible persons. In the model, vaccinated people do not die if they get infected, and the rate of transmission is $vi = 3$ times lower when a susceptible person comes in contact with a vaccinated and infected one.[13] To accommodate these features, new categories are created for the vaccinated and not infected, vaccinated and infected with the wild strain, and vaccinated and infected with the mutant strain. Each of these categories is also subdivided into vulnerable and young, and the laws of motion are modified accordingly.

Calibration. The parameters in the model are calibrated to capture the evolution of the pandemic in the United Kingdom. The path of vaccinations, number of deaths, and other daily data on the COVID-19 pandemic can be found on the datasets compiled by Our World in Data. In addition, the pre-pandemic number of daily contacts is calibrated using results from the American Community Survey, and relative mobility is available through Google's Community Mobility Reports. Figure B.2 in the appendix shows the predicted path for mobility and number of contacts based on estimates by Radzikowski and Dizioli (2021).

The top panel of table 1 shows the values of the parameters mentioned above. We also include the probability of death (once infected) of the vulnerable (δ^v) and the young (δ^y), as well as their population shares (s^v and s^y). Those parameters are calibrated by defining the vulnerable population as those aged 65 and over. The bottom panel of table 1 shows the share of the population that is tested in each of the 500 days between February 29th, 2020 and July 12th, 2021.

---

[13] While this number is still highly uncertain, there is increasing evidence that vaccines can substantially cut the rates of transmission for infected individuals. A recent study by Levine-Tiefenbrun et al. (2021) shows that vaccines can significantly reduce the viral load if an infection occurs, which has been shown to affect the probability of transmission. A *Daily Briefing* by the Advisory Board published on March 4th also discusses the issue (available at: https://www.advisory.com/en/daily-briefing/2021/03/04/vaccine-transmission).

Table 1: Parameter Values

Parameter	β_0	β_s	λ	γ	θ	μ	vi	δ^v	δ^y	s^v	s^y
Value	0.1658	0.0835	0.028	0.17	0.12	1.47	3	0.02473	0.00069	0.18	0.82

Days	1-14	15-27	28-40	41-60	61-101	102-500
q	0	0.05	0.25	0.28	0.3	0.3

5.2 Adding Vaccine Hesitancy

Vaccine hesitancy can affect vaccination efforts in two ways. First, it introduces a cap in terms of the number of people who will get the vaccine, which can be below the required threshold needed to achieve herd immunity. Second, it can reduce the average number of people getting vaccinated per day, for example when a low uptake of the vaccine leaves doses unused or discarded.[14] We introduce each of those outcomes in the model by changing the number of vaccines delivered per day, relative to the baseline vaccination path depicted in Figure 10.[15]

5.3 The Toll of Vaccine Hesitancy

Our counterfactual exercises compare the predicted number of COVID-19 related deaths in the baseline model with the same measure when we increase the levels of vaccine hesitancy to a country in the median of the distribution of hesitancy levels, and to a country at the top of the distribution. Our baseline roughly reflects the average hesitancy observed in the United Kingdom, the median of the distribution is represented by the hesitancy levels observed in the Republic of Korea, and the maximum hesitancy level we consider is that of France. In all cases, we average the data across the entire sample period to arrive at the level of hesitancy in each country.

---

[14] There have been several reports of this across Europe and the U.S. In European countries, low willingness to take the AstraZeneca vaccine has left millions of doses unused by the end of February – that is, before the temporary prohibition of the vaccine due to potential side effects, including blood clots (see https://www.ft.com/content/767fdd85-5329-479d-b565-4ec85d28b492). In the U.S., some healthcare providers have struggled to administer doses due to lack of demand, especially in underserved communities (for one example, see https://www.nbcbayarea.com/investigations/nearly-100k-vaccine-doses-unused-in-santa-clara-county-smaller-providers-struggling/2454267/).

[15] The path depicted in Figure 10 is the projected vaccination path for the United Kingdom produced by Radzikowski and Dizioli (2021) based on data obtained from Airfinity, a predictive science company. Note that this does not necessarily reflect the realized number of daily vaccinations in the country going forward, and is used only as a common baseline to compare the effects of different levels of vaccine hesitancy.

We analyze two different scenarios, described below, that differ based on authorities' ability to keep the pace of vaccination rollouts. The vaccine hesitancy levels are described in Table 2.

Figure 10: Baseline Vaccination Path
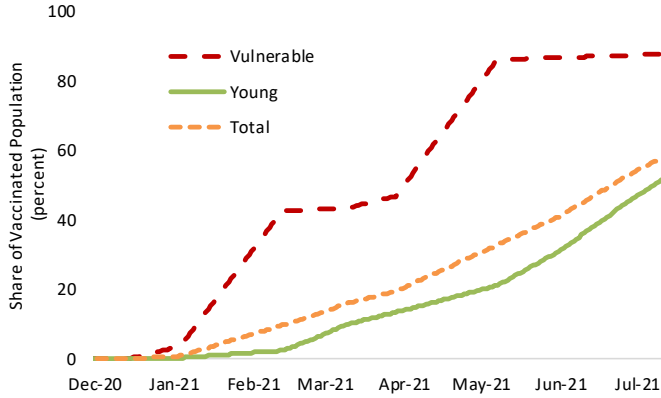
Source: Projected vaccination path from Radzikowski and Dizioli (2021) based on UK data

Table 2: Vaccine Hesitancy Thresholds

	Baseline		Median		Max	
Share of the population that	Vulnerable	Young	Vulnerable	Young	Vulnerable	Young
wants to be vaccinated	89.35%	74.07%	75.00%	62.07%	57.48%	34.71%

Scenario 1. In the first scenario, we assume that both effects discussed above play a role in the vaccination effort. The vaccine hesitancy cap is introduced by changing the number of daily vaccinations for a group to zero once the cumulative share of the vaccinated population in that group reaches the threshold in Table 2.[16] In addition, the pace of the vaccine rollout is also affected by, and in proportion to, the level of hesitancy in the country. For example, suppose that the baseline has v_t^v vaccines administered each day to the vulnerable, v_t^y vaccines administered to the young, and we change vaccine hesitancy to its maximum level in Table 2. The counterfactual number of vaccines administered per day would be $v_t^v \times \frac{57.48}{89.35}$ for the vulnerable population and $v_t^y \times \frac{34.71}{74.07}$ for the young.

---

[16] This does not affect the vaccination path for other groups (i.e., we do not allow for the reallocation of vaccines to another group that hasn't reached its cap yet). The cap is only binding in scenario 2; in scenario 1, due to the slower vaccination pace, the thresholds are not reached in the time frame we consider.

Figure 11: Deaths related to COVID-19 in scenario 1

Daily number of deaths related to COVID-19

Total number of deaths related to COVID-19
(Feb-20 to Jul-21)
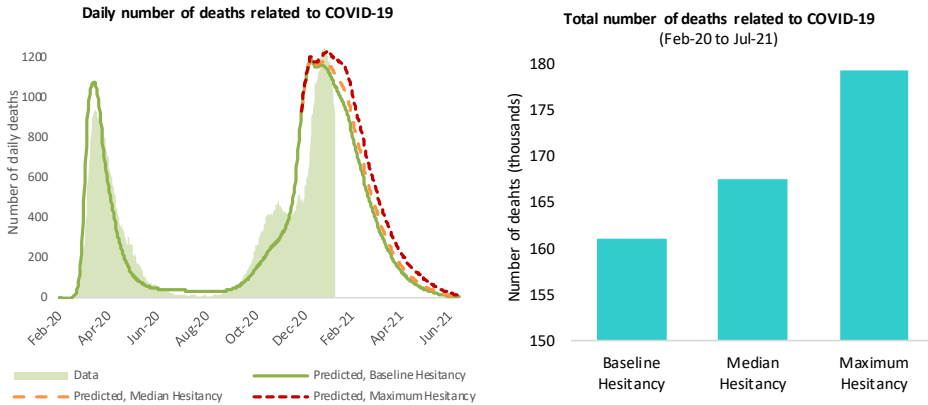
Figure 11 (left panel) shows the number of daily deaths related to COVID-19 in the United Kingdom (baseline), both in the data and predicted by the model. Note that the model does a fairly good job at matching the data picking up both waves of the disease in the UK (in large part, this reflects the mobility and contact data; see Figure B.2). Starting in end-December 2020, when the vaccine rollout began, we also include the predicted number of daily deaths under different levels of vaccine hesitancy.

In the right panel of the figure, we aggregate the predicted number of deaths between February 2020 and mid-July 2021. From these numbers, it is clear that vaccine hesitancy could have devastating effects for the country. Increasing the level of vaccine hesitancy from the baseline to the maximum level would cost roughly 18,000 lives, which is 11 percent more than the total death toll that is expected in the current baseline. We see a similar picture when looking at the number of infected cases (see appendix C).

Scenario 2. In the second scenario, we maintain the hesitancy cap but assume that governments are able to maintain the same pace of vaccination, regardless of the level of vaccine hesitancy. The outcomes in this case, shown in Figure 12, are very different. The trajectory of all three curves (left panel) are the same until the beginning of April, when the 57.48 percent vaccination cap is reached for the vulnerable population. The second cap (75 percent) for the vulnerable is reached by the end of that same month, and the 34.71 percent cap for the young is reached in June.

The key difference in this case is that by the time the curves diverge, the number of cases is considerably lower (see appendix C) and a much larger share of the population is already immune. Vaccine hesitancy is still costly in terms of more cases and deaths, but as shown in the right panel of Figure 12, the number of excess deaths relative to the baseline is much lower.

When comparing hesitancy levels between the baseline and the maximum hesitancy levels, there are slightly less than 1,000 excess deaths, which is almost 20 times lower than scenario 1.

Figure 12: Deaths related to COVID-19 in Scenario 2

Discussion. There are two important factors that play a role in the predictions above. First, note that the impact of vaccine hesitancy on the number of deaths is non-linear. In scenario 2, this happens because the vaccination cap is reached earlier when vaccine hesitancy is higher, which means that fewer people are vaccinated by then, and the chances of infection are higher. In scenario 1, this effect is compounded by a slower rollout, which causes the curves to diverge earlier and further increases the chances of viral infection at any point in time.

Second, the number of new infections crucially depends on assumptions about mobility and the extent to which people adhere to safety protocols, and we adopt a conservative stance in both cases. The rate of infections remains at its lowest point (β_s) once it is reached, which means that adherence to mask-wearing and other safety protocols is maintained even as most people get vaccinated. In addition, the number of daily contacts slowly increases, but never reaches its pre-pandemic levels in the time frame considered. It is also worth mentioning that a higher level of hesitancy leads to a higher peak in daily cases (in scenario 1). While not included in our model, this can lead to a higher rate of mortality due to overcrowded ICUs. A relaxation of any of those assumptions in the model could lead to more deaths due to vaccine hesitancy in both scenarios.

5.4 Policy Implications

The results presented in this section point to one clear policy implication: vaccine hesitancy can have a significant impact on pandemic dynamics if allowed to slow down the pace of vaccination. As a result, health authorities should focus on speeding up vaccine rollout until every willing adult has received their shot.

Achieving this goal might involve moving quickly through priority groups if demand for vaccines by the vulnerable population fails to meet supply. It could also entail having flexible days and times for appointments to receive a vaccine, multiple vaccination sites to avoid difficulties in commuting to a site, allowing individuals to choose their vaccine brand within the possibilities of each location (to avoid brand hesitancy), and allowing the population to pre-register or signal their intent to receive a vaccine so that authorities can assess demand with some presence. Because hesitancy and impediments to vaccinations can vary between regions/locations, it could also be advantageous to allow different vaccination sites to tailor the mix of policies to better meet their needs.

Another conclusion from the model presented above is that the path for mobility and other individual precautions taken to avoid infection (wearing masks and adhering to social distancing) could also play a key role in determining the toll of the pandemic. In our simulations, we assume that people remain vigilant, so that transmission is never as easy as it was at the beginning of the pandemic. Relaxing those measures could increase the number of deaths even if authorities are able to keep the pace of vaccination. It is also worth noticing that when the rate of infection is low, the level of vaccination required to achieve herd immunity is lower.

Finally, decreasing vaccine hesitancy also has myriad other benefits that are not directly related to the number of COVID-19 related deaths. A faster vaccine rollout will allow communities to quickly reach herd immunity and safely reopen the economy, allowing for the recovery of economic activity and growth. In addition, having a larger share of the population vaccinated can decrease the chances that new variants of the virus evolve.

6. Conclusion

While there has been steady progress towards resolving global vaccine supply constraints in recent months, vaccine hesitancy still poses serious challenges to achieving herd immunity. In this paper, we use individual-level survey data for 17 countries between November 2020 and April 2021 to understand the drivers and implications of COVID-19 vaccine demand, and what aspects should be prioritized when designing policies to tackle hesitancy.

Across our entire sample, nearly 40 percent of respondents are either unsure or unwilling to take the vaccine. There are also considerable differences in vaccine hesitancy across and within countries, which are correlated with pre-pandemic measures of general vaccine hesitancy. However, a positive development is that COVID-19 vaccine demand has been rising since the first effective vaccine was announced in November 2020, suggesting that hesitancy can be mitigated to some extent.

Turning to individual-level data, we find that there are systematic differences in vaccine demand across demographic groups, with higher demand among older cohorts, and also among men—a

gender gap that is consistent across ages and robust to controlling for household composition. Consistent with the pre-pandemic literature on vaccine hesitancy, these demographic differences are, in turn, partly attributable to differences in attitudes and beliefs. Important drivers of vaccine demand include individual concerns about the (potential) severity of COVID-19, self-reported compliance with protective behaviors, and overall trust in government.

Finally, we also show that vaccine demand is linked to how information is shared with peers, such as friends and family. People who are more exposed to warnings against the vaccine from peers are also less likely to want to take the shot, and, conversely, hesitant respondents are also more likely to share negative information about COVID-19 vaccines with their peers. This correlation suggests that managing information about COVID-19 vaccines, including through public health policies and communication targeted at informing the public about vaccine safety and effectiveness, are key to containing vaccine hesitancy.

Building on these empirical results, in the second part of the paper we extend a canonical SIR model to examine the implications of vaccine hesitancy for pandemic dynamics. We consider two channels through which hesitancy can affect vaccine rollouts, namely a reduction in the overall share of the population that gets vaccinated and a slowing down of mass vaccination programs. When both effects are operational, hesitancy can have a dramatic impact on COVID-19 dynamics. This suggests that the policy benefits of tackling vaccine hesitancy and increasing vaccination speed could be very large.

Our paper points towards several interesting avenues for future research. First, our results show that vaccine hesitancy can cluster geographically and across specific social networks, partly due to peer effects. This suggests ways to target policies to promote vaccine demand, but also the need to monitor potential local flare-ups of COVID-19 cases even after herd immunity is achieved at the country level. Second, it would be useful to understand if mandatory vaccination strategies are optimal and how they could be implemented effectively (e.g., linked to employment, school), given that COVID-19 vaccine hesitancy is also strongly correlated with broader distrust of institutions and noncompliance with social distancing recommendations. Third, further research on higher hesitancy among working-age women, and also households with children is important given that many of these individuals will be closely involved in COVID-19 vaccination decisions for their children once those vaccines become available. Finally, it would be interesting to embed our modelling results in a broader economic model, to understand the potential economic impact of vaccine hesitancy, both over the short and medium-term.

CEPR PRESS

References

Acemoglu, D., Makhdoumi, A., Malekian, A. and Ozdaglar, A., (2020). Testing, voluntary social distancing and the spread of an infection, (Working Paper No. 27483). *National Bureau of Economic Research*. http://www.nber.org/papers/w27483

Allcott, H., Boxell, L., Conway, J. C., Gentzkow, M., Thaler, M., and Yang, D. Y., (2020). Polarization and public health: partisan differences in social distancing during the coronavirus pandemic, (Working Paper No. 26946). *National Bureau of Economic Research*. https://doi.org/10.3386/w26946

Alsan, M., and Wanamaker, M., (2018). Tuskegee and the health of black men. *The Quarterly Journal of Economics*, Vol. 133, Issue 1, 407-455.

Alsan, M., and Eichmeyer, S., (2021). Persuasion in medicine: messaging to increase vaccine demand, (Working Paper No. 28593). *National Bureau of Economic Research*. https://www.nber.org/papers/w28593

Barrios, J. M., Benmelech, E., Hochberg, Y. V., Sapienza, P. and Zingales, L., (2021). Civic capital and social distancing during the Covid-19 pandemic. *Journal of Public Economics,* 193. 104310. 10.1016/j.jpubeco.2020.104310.

Barrios, J. M., and Hochberg, Y., (2020). Risk perception through the lens of politics in the time of the covid-19 pandemic, (Working Paper No. 27008). *National Bureau of Economic Research*. https://doi.org/10.3386/w27008

Bodine-Baron, E., Nowak, S., Varadavas, R. and Sood, N., (2013). Conforming and non-conforming peer effects in vaccination decisions (Working Paper No. 19528). *National Bureau of Economic Research*. https://www.nber.org/papers/w19528

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

Bughin, J., Cincera, M., Peters, K., Reykowska, D., Zyszkiewicz, M. and Ohmed, R., (2021). Make it or break it: vaccination intent at the time of Covid-19. *Centre for Economic Policy Research*, Covid Economics 65, 20 January 2021: 188-222.

Carrieri, V., Madio, L., Principe, F., (2019). Vaccine hesitancy and (fake) news: quasi-experimental evidence from Italy. *Health Economics*, 2019; 28: 1377– 1382. https://doi.org/10.1002/hec.3937

Chen, F. and Toxvaerd, F. (2014). The economics of vaccination. *Journal of Theoretical Biology*, 363, 105-117.

Croson, R., and Gneezy, U., (2009). Gender differences in preferences. *Journal of Economic Literature*, 47 (2): 448-74. DOI: 10.1257/jel.47.2.448

Dabla-Norris, E., Khan, H. and Lima, F., (2021). What determines social distancing? Evidence from advanced and emerging market economies. *IMF Working Paper*, forthcoming

Das, J., and Das, S., (2003). Trust, leaning and vaccination: A case study of a North Indian village. *Social Science & Medicine,* 57(1):97-112. doi: 10.1016/s0277-9536(02)00302-7. PMID: 12753819.

Du, F., Chantler, T., Francis, M. R., Sun, F. Y., Zhang, X., Han, K., Rodewald, L., Yu, H., Tu, S., Larson, H. and Hou, Z., (2020). The determinants of vaccine hesitancy in China: a cross-sectional study following the Changchun Changsheng vaccine incident. *Vaccine*, 2020 Nov 3;38(47):7464-7471. doi: 10.1016/j.vaccine.2020.09.075. Epub 2020 Oct 9. PMID: 33041097.

Dupas, P. and Miguel, E., (2017). Chapter 1 - Impacts and determinants of health levels in low-income countries. *Handbook of Economic Field Experiments*, Editors: Abhijit V. Banerjee, Esther Duflo, North-Holland, Vol. 2, 3-93.

Eichenbaum, M., Rebelo, S. and Trabandt, M., (2021). The macroeconomics of epidemics. Working Paper.

Fan, Y., Orhun, A. Y., and Turjeman, D., (2020). Heterogeneous actions, beliefs, constraints and risk tolerance during the COVID-19 pandemic (Working Paper No. 27211). *National Bureau of Economic Research*. doi: 10.3386/w27211

Figueiredo, A., Simas, C., Karafillakis, E., Paterson, P. and Larson, H. J., (2020). Mapping global trends in vaccine confidence and investigating barriers to vaccine uptake: a large-scale retrospective temporal modelling study. *The Lancet,* Volume 396, Issue 10255, Pages 898-908, https://doi.org/10.1016/S0140-6736(20)31558-0

Galasso, V., Pons, V., Profeta, P., Becher, M., Brouard, S. and Foucault, M., (2020). Gender differences in COVID-19 attitudes and behavior: panel evidence from eight countries. *Proceedings of the National Academy of Sciences of the United States of America*, 202012520; DOI: 10.1073/pnas.2012520117

Lazarus, J.V., Ratzan, S.C., Palayew, A. et al., (2021). A global survey of potential acceptance of a COVID-19 vaccine. *Nature Medicine* 27, 225–228. https://doi.org/10.1038/s41591-020-1124-9.

Levin, A.T., Hanage, W.P., Owusu-Boaitey, N., et al. (2020). Assessing the age specificity of infection fatality rates for COVID-19: systematic review, meta-analysis, and public policy implications. *European Journal of Epidemiology,* volume 35, 1123–1138. https://doi.org/10.1007/s10654-020-00698-1

Levine-Tiefenbrun, M., Yelin, P., Katz, R., Herzel, E., Golan, Z., Schreiber, L., Wolf, T., Nadler, V., Ben-Tov, A., Kuint, J., Gazit, S., Patalon, T., Chodick, G., and Kishony, R. (2021). Decreased SARS-CoV-2 viral load following vaccination. medRxiv, pre-print. https://doi.org/10.1101/2021.02.06.21251283

Loomba, S., Figueiredo, A., Piatek, S.J. et al., (2021). Measuring the impact of COVID-19 vaccine misinformation on vaccination intent in the UK and USA. *Nature Human Behaviour* 5, 337–348. https://doi.org/10.1038/s41562-021-01056-1

Martinez-Bravo, M., and Stegmann, A., (2021). In vaccines we trust? the effects of the CIA's vaccine ruse on immunization in Pakistan. *Center for Economic and Policy Research,* Discussion Paper No. DP15847, Available at SSRN: https://ssrn.com/abstract=3795231

Milkman, K., Beshears, J., Choi, J., Laibson, D., and Madrian, B., (2011). Using implementation intentions prompts to enhance influenza vaccination rates. *Proceedings of the National Academy of Sciences of the United States of America*, 108, 10415–10420.

Murphy, J., Vallières, F., Bentall, R.P. et al., (2021). Psychological characteristics associated with COVID-19 vaccine hesitancy and resistance in Ireland and the United Kingdom. *Nature Communications,* volume 12, 29. https://doi.org/10.1038/s41467-020-20226-9

Rao, N., Mobius, M. M. and Rosenblat, T., (2007). Social networks and vaccination decisions. *Federal Reserve Bank of Boston Research,* Working Paper No. 07-12, Available at SSRN: http://dx.doi.org/10.2139/ssrn.1073143

Neumann-Böhme, S., Varghese, N. E., Sabat, I., Barros, P. P., Brouwer, W., Van-Exel, J., Schreyögg, J. and Stargardt, T., (2020). Once we have it, will we use it? A European survey on willingness to be vaccinated against COVID-19. *The European Journal of Health Economics*, 2020 Sep;21(7):977-982. doi: 10.1007/s10198-020-01208-6.

O'Driscoll, M., Ribeiro Dos Santos, G., Wang, L. et al., (2021). Age-specific mortality and immunity patterns of SARS-CoV-2. *Nature,* 590, 140–145. https://doi.org/10.1038/s41586-020-2918-0

Omer, S. B., Salmon, D. A., Orenstein, W. A., DeHart, M. P. and Halsey, N., (2009). Vaccine refusal, mandatory immunization, and the risks of vaccine-preventable diseases. *New England Journal of Medicine*; 360:1981-1988 DOI: 10.1056/NEJMsa0806477

Oster, E., (2018). Does disease cause vaccination? disease outbreaks and vaccination response. *Journal of Health Economics*, Volume 57, Pages 90-101

Papageorge, N. W., Zahn, M. V., Belot, M., Broek-Altenburg E. V. D., Choi, S., Jamison, J. C., and Tripodi, E., (2020). Socio-demographic factors associated with self-protecting behavior during the Covid-19 pandemic, (Working Paper No. 27378). *National Bureau of Economic Research*. doi 10.3386/w27378

Paul, E., Steptoe, A. and Fancourt, D., (2021). Attitudes towards vaccines and intention to vaccinate against COVID-19: Implications for public health communications. *The Lancet Regional Health Europe*, Volume 1, 2021, 100012, ISSN 2666-7762, https://doi.org/10.1016/j.lanepe.2020.100012.

Radzikowski, A. and Dizioli, A., (2021). IMF Research Department's Pandemic Tool. *International Monetary Fund,* forthcoming.

Rowthorn, R. and Toxvaerd, F., (2020). On the management of population immunity. *Bennett Institute,* Working Paper.

Sadaf, A., Richards, J. L., Glanz, J., Salmon. D. A. and Omer, S. B., (2013). A systematic review of interventions for reducing parental vaccine refusal and vaccine hesitancy. *Vaccine*. 13:31(40):4293-304. doi: 10.1016/j.vaccine.2013.07.013. Epub 2013 Jul 13. PMID: 23859839.

Sato, R. and Takasaki, Y., (2019). Peer effects on vaccination behavior: experimental evidence from rural Nigeria. *Economic Development and Cultural Change* 68, 93–129 (2019).

Simonov, A., Sacher, S., Dubé, J. P and Biswas, S., (2020). The persuasive effect of Fox News: non-compliance with social distancing during the COVID-19 pandemic. *Social Science Research Network Electronic Journal* DOI: 10.2139/ssrn.3604214

Thunström, L., Ashworth, M., Finnoff, D. and Newbold, S. C., (2020). Hesitancy towards a COVID-19 vaccine and prospects for herd immunity. *Centre for Economic Policy Research,* Covid Economics 35, 7 July 2020: 1-50.

CEPR PRESS

A Tables

Table A.1 Agree to be Vaccinated vs Disagree to be Vaccinated – Baseline Model

	(1)	(2)	(3)	(4)
Number Children	-0.013**	-0.015***	-0.006**	-0.006*
	(0.004)	(0.004)	(0.003)	(0.003)
Easy Access to Vaccine	0.115***	0.113***	0.044***	0.045***
	(0.017)	(0.016)	(0.010)	(0.012)
Safe Behavior		0.528***	0.162***	0.169***
		(0.083)	(0.034)	(0.033)
Wears Masks		0.098***	0.045**	0.047***
		(0.026)	(0.016)	(0.015)
Worried About COVID-19				
2			0.044***	0.031***
			(0.006)	(0.005)
3			0.090***	0.084***
			(0.010)	(0.008)
4			0.159***	0.151***
			(0.013)	(0.011)
5 (Str. Agree)			0.225***	0.221***
			(0.015)	(0.010)
Worried About Side Effects				
2			0.030***	0.023**
			(0.007)	(0.008)
3			-0.003	-0.007
			(0.012)	(0.013)
4			-0.098***	-0.109***
			(0.015)	(0.014)
5 (Str. Agree)			-0.284***	-0.297***
			(0.018)	(0.018)
Gov. Provide Effective Vaccine				
2			0.045***	0.035***
			(0.009)	(0.011)
3			0.137***	0.126***
			(0.012)	(0.013)
4			0.393***	0.374***
			(0.014)	(0.017)
5 (Str. Agree)			0.496***	0.476***
			(0.017)	(0.020)
Confidence in HC System				
2 (Not Much)				0.017***
				(0.005)
3 (Not Sure)				-0.017*
				(0.009)
4 (Fair)				0.044***
				(0.009)
5 (A Lot)				0.048***
				(0.012)
Gov Handled Pandemic Well				
2				-0.009
				(0.009)
3				-0.074***
				(0.012)
4				-0.024*
				(0.011)
5 (Str. Agree)				-0.020
				(0.015)
Observations	83805	83805	83805	60521
Within R2	0.0300	0.0452	0.3141	0.3201

All Specifications control for age group by gender, pre-existing health conditions, whether individuals had COVID-19 symptoms, and household size. We also include occupation and location-by-week fixed effects. Standard errors are clustered at the country level and shown in parenthesis. *, ** and *** indicate that coefficients are statistically different from 0 at the 10%, 5%, and 1% levels, respectively.

Table A.2 Agree to be Vaccinated vs Disagree to be Vaccinated – Peer Effects

	(1)	(2)	(3)	(4)	(5)
Number Children	-0.017***	-0.018***	-0.007***	-0.007***	-0.004
	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)
Easy Access to Vaccine	0.118***	0.117***	0.043***	0.040***	0.028***
	(0.018)	(0.018)	(0.010)	(0.011)	(0.009)
Log(Vaccines/100 people)	0.015**	0.022**	0.008	0.012	0.004
	(0.006)	(0.009)	(0.006)	(0.008)	(0.008)
Safe Behavior		0.237***	0.036	0.156***	0.092**
		(0.068)	(0.029)	(0.047)	(0.042)
Wears Masks		0.016	0.021	0.053***	0.017
		(0.034)	(0.012)	(0.016)	(0.013)
Worried About COVID-19					
2			0.047***	0.035***	0.026***
			(0.005)	(0.005)	(0.006)
3			0.094***	0.087***	0.073***
			(0.011)	(0.010)	(0.011)
4			0.160***	0.151***	0.126***
			(0.017)	(0.014)	(0.012)
5 (Str. Agree)			0.230***	0.224***	0.171***
			(0.018)	(0.012)	(0.012)
Worried About Side Effects					
2			0.031***	0.027***	0.015*
			(0.007)	(0.009)	(0.008)
3			-0.007	-0.010	0.005
			(0.011)	(0.012)	(0.008)
4			-0.108***	-0.118***	-0.083***
			(0.016)	(0.017)	(0.014)
5 (Str. Agree)			-0.293***	-0.304***	-0.204***
			(0.019)	(0.021)	(0.017)
Gov. Provide Effective Vaccine					
2			0.049***	0.031**	0.005
			(0.013)	(0.013)	(0.017)
3			0.147***	0.132***	0.130***
			(0.017)	(0.017)	(0.023)
4			0.395***	0.371***	0.304***
			(0.020)	(0.021)	(0.028)
5 (Str. Agree)			0.495***	0.464***	0.362***
			(0.019)	(0.023)	(0.030)
Confidence in HC System					
2 (Not Much)				0.014*	-0.007
				(0.007)	(0.007)
3 (Not Sure)				-0.014	-0.021
				(0.011)	(0.014)
4 (Fair)				0.037***	0.002
				(0.009)	(0.007)
5 (A Lot)				0.036**	-0.012
				(0.013)	(0.012)
Gov Handled Pandemic Well					
2				-0.012	-0.014
				(0.011)	(0.009)
3				-0.081***	-0.049***
				(0.012)	(0.013)
4				-0.021*	-0.024***
				(0.012)	(0.008)
5 (Str. Agree)				-0.009	-0.005
				(0.016)	(0.015)
% Peers Will Vac					0.005***
					(0.000)
Observations	53396	53396	53396	38335	29826
Within R2	0.0337	0.0386	0.3277	0.3372	0.3800

All Specifications control for age group by gender, pre-existing health conditions, whether individuals had COVID-19 symptoms, and household size. We also include occupation, week, and location fixed effects. Standard errors are clustered at the country level and shown in parenthesis. *, ** and *** indicate that coefficients are statistically different from 0 at the 10%, 5%, and 1% levels, respectively.

Table A.3 Agree to be Vaccinated vs Unsure About Vaccination

	(1)	(2)	(3)	(4)
Number Children	-0.007*	-0.009**	-0.006*	-0.007**
	(0.004)	(0.004)	(0.003)	(0.003)
Easy Access to Vaccine	0.069***	0.068***	0.027***	0.022**
	(0.011)	(0.011)	(0.008)	(0.008)
Safe Behavior		0.342***	0.206***	0.217***
		(0.059)	(0.039)	(0.037)
Wears Masks		0.074***	0.055***	0.056***
		(0.019)	(0.013)	(0.014)
Worried About COVID-19				
2			-0.007	-0.019**
			(0.008)	(0.008)
3			-0.033***	-0.040***
			(0.009)	(0.010)
4			0.054***	0.046***
			(0.006)	(0.006)
5 (Str. Agree)			0.101***	0.094***
			(0.007)	(0.007)
Worried About Side Effects				
2			0.008*	0.009*
			(0.004)	(0.005)
3			-0.063***	-0.061***
			(0.010)	(0.010)
4			-0.066***	-0.068***
			(0.011)	(0.011)
5 (Str. Agree)			-0.168***	-0.171***
			(0.014)	(0.012)
Gov. Provide Effective Vaccine				
2			-0.056***	-0.055***
			(0.010)	(0.010)
3			-0.143***	-0.137***
			(0.025)	(0.021)
4			0.107***	0.101***
			(0.016)	(0.013)
5 (Str. Agree)			0.163***	0.153***
			(0.016)	(0.014)
Confidence in HC System				
2 (Not Much)				0.023*
				(0.011)
3 (Not Sure)				-0.000
				(0.016)
4 (Fair)				0.058***
				(0.010)
5 (A Lot)				0.071***
				(0.012)
Gov Handled Pandemic Well				
2				-0.012*
				(0.006)
3				-0.071***
				(0.012)
4				-0.019**
				(0.007)
5 (Str. Agree)				-0.018*
				(0.010)
Observations	71701	71701	71701	51615
Within R2	0.0191	0.0272	0.1772	0.1814

All Specifications control for age group by gender, pre-existing health conditions, whether individuals had COVID-19 symptoms, and household size. We also include occupation and location-by-week fixed effects. Standard errors are clustered at the country level and shown in parenthesis. *, ** and *** indicate that coefficients are statistically different from 0 at the 10%, 5%, and 1% levels, respectively.

Table A.4. Agree to be Vaccinated vs Disagree to be Vaccinated – Logit Specification (Odds Ratio)

	(1)	(2)	(3)	(4)
Number Children	0.946***	0.938***	0.963**	0.958**
	(0.018)	(0.018)	(0.016)	(0.017)
Easy Access to Vaccine	1.695***	1.703***	1.349***	1.349***
	(0.138)	(0.138)	(0.093)	(0.100)
Safe Behavior		3.090***	1.551***	3.090***
		(0.876)	(0.250)	(0.636)
Wears Masks		1.017	1.040	1.353***
		(0.140)	(0.094)	(0.133)
Worried About COVID-19				
2			1.444***	1.311***
			(0.071)	(0.054)
3			1.964***	1.844***
			(0.153)	(0.121)
4			3.093***	2.812***
			(0.313)	(0.253)
5 (Str. Agree)			5.001***	4.652***
			(0.570)	(0.430)
Worried About Side Effects				
2			1.129**	1.051
			(0.060)	(0.065)
3			0.852*	0.818**
			(0.080)	(0.081)
4			0.480***	0.442***
			(0.051)	(0.048)
5 (Str. Agree)			0.162***	0.146***
			(0.022)	(0.020)
Gov. Provide Effective Vaccine				
2			1.570***	1.526***
			(0.086)	(0.080)
3			2.528***	2.470***
			(0.157)	(0.139)
4			8.477***	7.888***
			(0.592)	(0.594)
5 (Str. Agree)			19.058***	17.361***
			(1.680)	(1.491)
Confidence in HC System				
2 (Not Much)				1.159***
				(0.043)
3 (Not Sure)				0.949
				(0.061)
4 (Fair)				1.342***
				(0.071)
5 (A Lot)				1.420***
				(0.111)
Gov Handled Pandemic Well				
2				0.984
				(0.059)
3				0.663***
				(0.054)
4				0.906
				(0.067)
5 (Str. Agree)				0.901
				(0.092)
Observations	83799	83799	83799	60519
Pseudo R2	0.0656	0.0693	0.3035	0.3098

Note: table displays odds ratios.
All Specifications control for age group by gender, pre-existing health conditions, whether individuals had COVID-19 symptoms and household size. We also include occupation, week, and location. Standard errors are clustered at the country level and shown in parenthesis. *, ** and *** indicate that logit coefficients are statistically different from 0 at the 10%, 5%, and 1% levels, respectively.

Table A.5: Select survey questions

Safe behaviors:
Thinking about the last 7 days… how often have you taken the following measures to protect yourself or others from coronavirus (COVID-19)? As a reminder, please exclude any measures that you have already taken for reasons other than coronavirus (COVID-19). Worn a face mask outside your home (e.g. when on public transport, going to a supermarket, going to a main road) Washed hands with soap and water Covered your nose and mouth when sneezing or coughing Cleaned frequently touched surfaces in the home (e.g. doorknobs, toilets, taps) Avoided contact with people who have symptoms or you think may have been exposed to the coronavirus Avoided small social gatherings (not more than 2 people) Avoided medium-sized social gatherings (between 3 and 10 people) Avoided large-sized social gatherings (more than 10 people)
Which, if any, of the following might make it hard for you to get a COVID-19 vaccine? COVID-19 vaccination is not yet available for me COVID-19 vaccination costs too much I can't go on my own The vaccination site is too far away The opening times are inconvenient People are turned away without vaccination The waiting time is too long Nothing. It would not be hard
In the last 7 days, have you personally been tested for coronavirus (COVID-19?)
How well or badly do you think the government are handling the issue of the Coronavirus (COVID-19)?
I believe government health authorities in my country will provide me with an effective COVID19 vaccine
How much confidence do you have in the healthcare system to respond to a Coronavirus (COVID-19) outbreak in your country?
I am worried about getting COVID19
I am worried about potential side effects of a COVID19 vaccine
How many of your close family and friends do you think will get a COVID-19 vaccine?
Do you think most of your close family and friends would want you to get a COVID-19 vaccine?
I have seen information online that has warned me about the negative effects of COVID-19 vaccines
Friends and family have warned me about the negative effects of the COVID-19 vaccines
I have shared my own concerns about the negative effects of the COVID-19 vaccines

Note: full list of survey questions can be found here: *https://github.com/YouGov-Data/covid-19-tracker*

CEPR PRESS

B. Figures

Figure B.1: Data availability by Country and Week

Week of (Friday)	11/11	11/18	12/16	12/23	1/1	1/8	1/15	1/22	1/29	2/5	2/12	2/19	2/26	3/5	3/12	3/19	3/26	4/2	Total
AUS	1		1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	17
CAN	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	18
DEU	1		1		1	1	1	1	1	1	1	1	1	1	1	1	1		15
DNK	1		1		1	1	1	1	1	1	1	1	1	1	1	1	1		15
ESP	1		1		1	1	1	1	1	1	1	1	1	1	1	1	1		15
FIN	1		1		1	1	1												5
FRA	1		1		1	1	1	1	1	1	1	1	1	1	1	1	1		15
GBR	1		1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		16
ISR								1	1	1	1	1	1	1	1	1	1		9
ITA	1		1		1	1	1	1	1	1	1	1	1	1	1	1	1		15
JPN	1				1	1	1	1	1		1			1	1	1	1		12
KOR	1	1			1	1	1	1	1	1	1	1	1	1	1	1	1		16
NLD	1		1		1	1	1	1	1										8
NOR	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		17
SGP	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	18
SWE	1		1		1	1	1	1	1	1	1	1	1	1	1	1	1		15
USA												1	1	1	1	1	1	1	7
Total	15	4	15	5	15	15	15	14	15	15	13	15	14	14	15	15	15	4	

Note: This figure shows survey wave dates for each country.

*Figure B.2: Mobility to Work and Number of Daily Contacts*

Mobility Relative to Pre-pandemic Baseline

Y-axis: 20, 0, -20, -40, -60, -80, -100
X-axis: Dec-19, Mar-20, Jun-20, Sep-20, Dec-20, Mar-21, Jun-21, Sep-21
Legend: – – – Predicted Mobility; —— Mobility (Data)

Mobility and Daily Contacts

Left Y-axis: 20, 0, -20, -40, -60, -80, -100
Right Y-axis (RHS): 6, 5, 4, 3, 2, 1, 0
X-axis: Dec-19, Mar-20, Jun-20, Sep-20, Dec-20, Mar-21, Jun-21, Sep-21
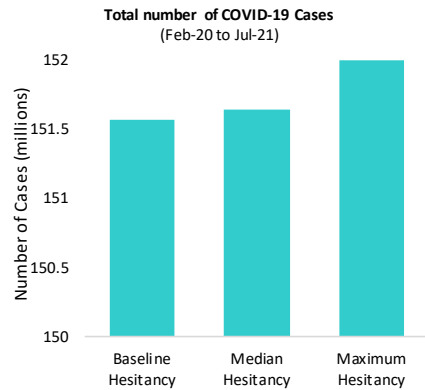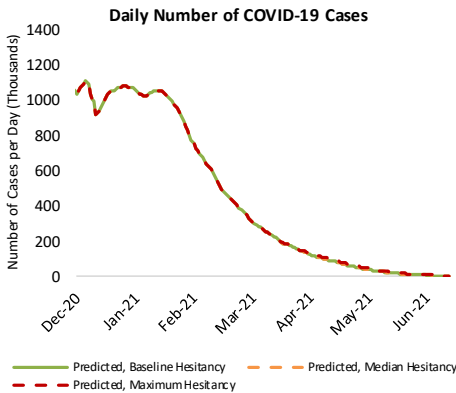Legend: – – – Preducted Mobility; —— Daily Contacts (RHS)

C Vaccine Hesitancy and the Pandemic Dynamics: COVID-19 Cases

Here we present the number of COVID-19 cases under each scenario discussed in Section 4 of the main text. As seen in the figures below, the number of cases follows the same pattern as the number of deaths (although at a larger scale) and thus the same intuition applies.

Scenario 1. Vaccine hesitancy slows down the pace of vaccination and introduces a cap at the total number of people vaccinated at any point in time.

Daily Number of COVID-19 Cases

Total number of COVID-19 Cases
(Feb-20 to Jul-21)

Predicted, Baseline Hesitancy — Predicted, Median Hesitancy — Predicted, Maximum Hesitancy
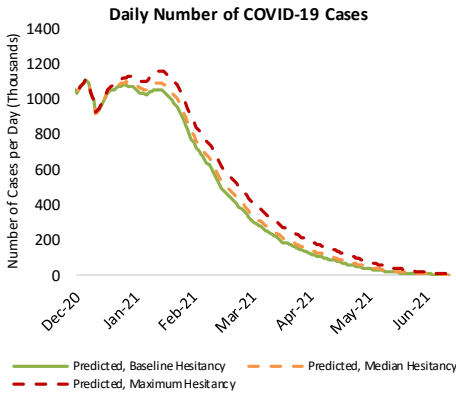
Scenario 2. Authorities adapt to hesitancy and keep the vaccination pace, but the cap remains.

Daily Number of COVID-19 Cases

Total number of COVID-19 Cases
(Feb-20 to Jul-21)

Predicted, Baseline Hesitancy — Predicted, Median Hesitancy — Predicted, Maximum Hesitancy

Incentives for accelerating the production of Covid-19 vaccines in the presence of adjustment costs[1]

Claudius Gros[2] and Daniel Gros[3]

Once new Covid vaccines were approved, vaccinating the population as quickly as possible became paramount. However, in the presence of adjustment costs firms will increase production capacity only gradually. The existing contracts specify solely that a fixed quantity is to be supplied within a specified period. At fixed prices, this kind of contract provides no incentive for an accelerated build up of vaccine production capacities within the stipulated delivery period. With adjustment costs the price is however very sensitive to the length of the delivery period (elasticity of 3). An optimal contract would specify a decreasing price schedule overtime that can replicate the social optimum. We show in particular that different forms of adjustment costs, whether they are defined in relation to absolute or proportional increases in capacity, can lead to qualitatively different production profiles over time. Evidence from Covid vaccines is not compatible with the textbook model of adjustment costs, which assumes a relation to proportional increases in production. We also show that if governments had chosen the delivery time so as to minimize social losses, they would have spent much more on vaccine production.

1. Introduction

Ramping up production became the key goal once the safety and efficacy of new vaccines against the SARS Covid-19 virus had been proven in clinical trials, which lead to the approval for widespread use by major medical agencies towards the end of 2020. However, it was not possible to quickly vaccinate the entire population because the supply of newly developed vaccines was limited in the short run [1].

The production of vaccines is undertaken by private firms under procurement contracts which specify typically that a stipulated amount of vaccines is to be delivered for a fixed price over a certain time period (usually several quarters or one year). For the firm a dose delivered today or tomorrow has thus the same value as long as it happens within the delivery period.

For society, however, a dose delivered earlier has a substantially higher value. With earlier delivery many lives could be saved and costly lockdowns could be shorter. The private and social costs of early delivery of vaccines are thus not aligned. The existence of adjustment costs leads profit maximising firms to increase capacity more gradually than it might be optimal from a social point of view. The problem for public authorities is then to find a way to accelerate the increase in production capacity. The vaccine supply contracts were mostly concluded before the vaccines had been fully developed, let alone approved for general use. It was thus impossible to impose tight delivery deadlines. The Advance Purchase Agreements of the EU, three of which have been published [2, 3, 4], thus specified only an overall price and tentative delivery schedules in terms of quarters, not months or weeks. When even these tentative schedules started to slip, the EU had little leverage to induce companies to make efforts to accelerate delivery.

We analyze the consequences of this type of contracts for the supply schedule of a vaccine and how the resulting incentives for back-loading supplies can be mitigated. We are not aware of any other contributions which analyse the time path of the production capacity for Covid vaccines. E.g., [5] models only the negotiations between manufacturers without considering the time dimension.

Limitations: A substantial part of the literature on vaccine policy focuses on how and whom to vaccinate, usually taking it as granted that the supply of vaccines is not a constraint [1, 6, 7, 8]. We do not consider this issue as we concentrate on the case of Covid vaccines, for which mass production had to start immediately after test trials were successful. Another issue we

do not consider is vaccine hesitancy [9], the situation that a certain proportion of the population refuses to be vaccinated, which may make it difficult to achieve herd immunity. The immediate problem facing policy makers is the opposite, at least initially, namely that demand for vaccination far outstrips supply. Moreover, even if full herd immunity could not be reached, there is still a considerable benefit from every person vaccinated, which reduces potential hospitalization costs correspondingly, allowing governments in the aggregate to end lockdown measures earlier [10, 11].

Likewise we do not consider the ex ante issue of uncertain efficacy of new vaccines and the related problem of ordering portfolios of potential candidates [10], as done by most major countries. Our analysis concentrates instead on the problem of ramping up production once the efficacy of a vaccine has been established [12]. The importance of this issue for the global economy has been laid out [11].

A by-product of our analysis is that the form of the adjustment cost has important implications for the time path of production. The nature of adjustment costs has never been fully clarified [13]. They are an essential element of any theory of investment, given that capital stocks would jump instantaneously to the desired level (with the consequence that investments would diverge positively or negatively), if adjustment costs were absent. Models of investment have therefore to posit the existence of a cost to increasing capital rapidly. These costs are usually specified as a convex function of the *proportional* increase in the capital stock. This specification has the advantage that it leads to tractable solutions, especially if the adjustment costs are specified in a quadratic form. Here we consider the case that adjustment costs are a function of the *absolute*, not the proportional, increase in production capacity. The proportional approach is difficult to justify when the initial capacity or capital stock is zero, the relevant scenario for the case of Covid-19 vaccine production. Proportional increases are diverging when starting with an empty factory. We also show that uncertainty concerning production costs [14] should not change our main results.

We start by analyzing the case in which firms producing the vaccines optimize the production time path with the aim of minimizing costs under the constraint of a fixed price. Next we examine how the resulting schedule differs under two different specifications of adjustment costs (proportional and absolute). The available data from the first few months of Covid-19 vaccine production are consistent with a model based on adjustment costs which depend on the proportional increase in capacity. We also calculate the

price needed for firms to induce a shorter contract period. Finally we consider the problem of building up production over time from the perspective of a social planner and show that it is equivalent to a pricing scheme with an initial high starting price, which declines subsequently linearly over time. The resulting optimal pricing scheme aligns the interests of the producer with that of the society as a whole.

2. Adjustment costs for ramping up vaccine production

The problem that the producers of a new product, like Covid-19 vaccines, face involves one key element, namely adjustment costs. It is not possible to ramp up production instantaneously. Standard economic analysis takes this into account by positing that there is a cost to increasing capacity and that this cost is convex, i.e. the costs of increasing capacity are only small when the buildup is slow [15, 13] . The implication is straightforward: it will be optimal to smooth production over time.

Consider a contract in which a certain quantity Z_T is to be delivered over a period T (say one year), at a constant price p_0 per unit. In this case the exact timing of the delivery, close to the start or to the end of the delivery period, does not matter for the revenues the producer receives. It will then be optimal for the producer to minimize costs by increasing capacity only gradually over time.

We thus start by analyzing the production path resulting from the type of contract that has been used standardly in 2020/21 for Covid-19 vaccine procurement, namely a fixed price for the delivery of a stipulated quantity over a time period specified in advance. For example, the Advanced Purchase Agreement of the European Union with Curevac specifies the delivery of certain amounts of doses for the year 2021 [3], with only *tentative* delivery schedules by quarter. This implies that the firm can distribute the supply schedule over the entire year, which is nearly an eternity in terms of a pandemic costing several percentage points of GDP at each instant, threatening at the same time uncountable lives every day. We thus focus on the inter-temporal problem of increasing production capacity over time within the overall time frame given by the contract, which could be thought of representing one year. Given this time frame (and interest rates around zero), we neglect time discounting.

In most applications [15], the convexity of the costs of adjustment is assumed to be quadratic (which would also be the result of a second order

approximation). With quadratic adjustment costs, as considered here, the marginal cost of adjusting becomes linear, allowing for explicit solutions.

2.1. Basic setup

Formally we consider a firm which has been contracted to supply an amount Z_T of doses over a given period T. The marginal cost of each dose, denoted by c, is assumed to be constant and independent of the level of capacity.

Denoting the instantaneous production capacity (the number of vaccines produced per unit time, say daily) with z_t, the adjustment costs will be a function $f(\dot{z}_t)$ of $\dot{z}_t$, which quantifies the speed at which production is ramped up. Overall adjustment costs are then determined by the integral of $f(\dot{z}_t)$ over the delivery period, subject to the constraint that a total of $Z_T = \int_0^T z_t dt$ units are produced. The total order can be thought of as the amount required to vaccinate the entire population, which would allow the lifting of all the restrictions needed before and during the vaccination campaign to keep the spread of the virus under control.

We assume that the initial capacity is low, possibly equal to zero, but definitely not large enough to satisfy the entire order within $t \in [0, T]$. This implies that the firm must scale up capacity during the contract period $[0, T]$.

Formally one can think of the production function as constant returns to scale, employing only capital so that the capacity to produce z_t units of the vaccine per unit of time is proportional to the capital stock available at that time.[1] We will mainly use the term capacity instead of capital stock, which is usually employed in literature on investment.

2.2. The evolution of the production capacity for fixed price contracts

The problem for the firm is to maximize revenues minus adjustment costs, subject to the overall production constraint:

$$\int_0^T p_0 z_t dt - a_z \int_0^T (\dot{z}_t)^2 dt - \lambda \left[\int_0^T z_t dt - Z_T \right] - \int_0^T c z_t dt \,, \qquad (1)$$

where p_0 denotes the price per vaccine and a_z encodes the scale of the adjustment costs. The production costs per unit is c, with the Lagrange multiplier

---

[1] In reality vaccine production requires not just the physical factory, which might have to satisfy specific requirements, but also the schooling of personnel, etc.

λ enforcing the constraint that the total production over the period $[0, T]$ is $Z_T = \int_0^T z_t dt$.

In (1) we use adjustment costs which do not depend on the level of capacity already reached. Substantial effort has been devoted in the literature to the study the alternative [15], namely adjustment costs which are a function of the proportional increase $\dot{z}_t/z_t$ of the capacity. However, this would lead to conceptual difficulties when starting the production of a new product (i.e. when $z_0 = 0$), the case of Covid-19 vaccines.[2]

Standard variational calculus [16, 17] establishes that the stationary solution to (1) satisfies

$$2a_z \ddot{z}_t = \lambda + c - p_0, \qquad z_t = z_0 + \gamma t + \frac{\lambda + c - p_0}{4a_z} t^2, \qquad (2)$$

where z_0 is the initial production capacity and γ the speed at which production capacity increases initially.

Only the difference between price and marginal costs enters the condition (2). In the remainder we thus assume at times, for computational convenience and without loss of generality, that c is equal to zero. Any constant marginal cost would add only a fixed amount to the overall costs of the firm. One could thus think of the price as representing the difference between the unit price contracted and any marginal cost of production.

The problem that the firm faces can be reduced to minimizing the total cost of adjustment over the delivery period T, as total revenues are fixed, being equal to the price times the quantity delivered. The production schedule that minimizes the adjustment cost is to increase capacity accordingly to (2).

A constant rate of increase in production would not be optimal, on general grounds, because an increase in capacity implemented today yields higher production over the remainder of the delivery period and is thus more valuable than an increase in capacity just before the end of the delivery period. The speed at which capacity increases should thus decline over time. This intuition is born out by (2). To be concrete, we parameterize the solution to

---

[2]The problem considered here differs from the textbook problem of investment under adjustment costs, see Blanchard and Fisher (Lectures on Macroeconomics, MIT Press), in which the firm adjusts output over time in response to a changing price or demand schedule. Here, the integral of output over a certain period is given and the main problem is how to satisfy the production constraint over time.

(2) with

$$z_t = z_0 + \gamma t + \delta t^2, \qquad \delta = \frac{\lambda + c - p_0}{4a_z}. \qquad (3)$$

The production condition $Z_T = \int_0^T z_t dt$ implies then

$$Z_T = z_0 T + \frac{\gamma}{2} T^2 + \frac{\delta}{3} T^3, \qquad \gamma = \frac{2\Delta Z}{T} - \frac{2\delta T}{3}. \qquad (4)$$

where ΔZ denotes the difference between the average capacity needed to fulfill the order and the initial one, $\Delta Z = Z_T/T - z_0$. It is assumed here that $\Delta Z > 0$, namely that the capacity needs to be increased. In the opposite case, when $Z_T < T z_0$, the company would have to shut down part of the existing production capacity - which is not the case for Covid vaccines.

The overall production constraint (4) can be satisfied by any linear combination of γ and δ. These two parameters are determined by maximizing total profit. Given that both the first and the second term in (1) are constant, together as $(p_0 - c)Z_T$, one just has to minimize the adjustment costs:

$$E_{\text{adj}} = a_z \int_0^T (\dot{z}_t)^2 dt = a_z \left[\gamma^2 T + 2\gamma\delta T^2 + \frac{4\delta^2 T^3}{3} \right], \qquad (5)$$

where $\dot{z}_t = \gamma + 2\delta t$ has been used. The relation (4) entails that $\partial\gamma/\partial\delta = -2T/3$, which leads to

$$\frac{dE_{\text{adj}}}{d\delta} = \frac{\partial E_{\text{adj}}}{\partial\delta} + \frac{\partial E_{\text{adj}}}{\partial\gamma}\frac{\partial\gamma}{\partial\delta} = \frac{\partial E_{\text{adj}}}{\partial\delta} - \frac{2T}{3}\frac{\partial E_{\text{adj}}}{\partial\gamma} \qquad (6)$$

$$= a_z \left[2\gamma T^2 + \frac{8\delta T^3}{3} - \frac{2T}{3}(2\gamma T + 2\delta T^2) \right] \qquad (7)$$

$$= a_z \left[\frac{2\gamma T^2}{3} + \frac{4\delta T^3}{3} \right] = 0.$$

The first order condition for cost minimization over the choice of γ and δ leads therefore to following simple relationships:

$$\gamma = -2\delta T, \qquad \gamma = 3\frac{\Delta Z}{T}, \qquad \delta = -\frac{3\Delta Z}{2T^2}. \qquad (8)$$

where the last two relations follow from (4). The time path z_t for the production capacity is then

$$z_t = z_0 + \frac{3\Delta Z}{T}\left[t - \frac{t^2}{2T} \right] \qquad \dot{z}_t = \frac{3\Delta Z}{T}\left[1 - \frac{t}{T} \right], \qquad (9)$$

as illustrated in Figure 1. This implies that the increase in capacity is at first approximately linear but declines over time and tends towards zero at the end of the delivery period, $\lim_{t \to T} \dot{z}_t \to 0$. The result (9) also implies that $\dot{z}_t$ is proportional to the missing average production capacity ΔZ, scaling inversely with the production period T. At the end of the contract period (t=T), the production capacity will be equal to 1.5 times the one which is needed on average (Z_T/T when $z_0 = 0$).

Note that the cost minimizing production path z_t does not depend on the overall cost of the order since the price p_0 does not influence the parameters of the differential equation (9). The reason is that p_0 enters the Lagrange multiplier of the equation of motion (4) only through the difference $\lambda + c - p_0$.

The key corollary from the above considerations, regarding the effects of adjustment costs, is then:

> The level of the price does not influence the speed at which production increases – when the price is constant.

Higher prices allow the producer to obtain larger profits, however without providing incentives to accelerate the buildup of the capacity. We have not considered explicitly the cost of developing the vaccine, which would add a constant term to the costs for the firm. But this constant term would also not have any influence on the speed at which production is increased since it represents just a sunk cost when the firm starts to ramp up production.

2.3. Adjustment costs proportional to capacity increase

We now consider the case of adjustment costs which depend on the proportional increase in capital or production capacity $\dot{z}_t/z_t$.[3] As for the case considered above, firms will maximise the difference between revenues and adjustment costs, subject to the constraint that the entire order has to be fulfilled within a time span of T. Denoting the integrand with $L(z_t, \dot{z}_t)$ (the Lagrangian, in physics terminology), the objective function $\mathcal{L}$,

$$\mathcal{L} = \int_0^T L(z_t, \dot{z}_t)dt, \qquad \frac{d}{dt}\frac{\partial L}{\partial \dot{z}_t} = \frac{\partial L}{\partial z_t}, \qquad (10)$$

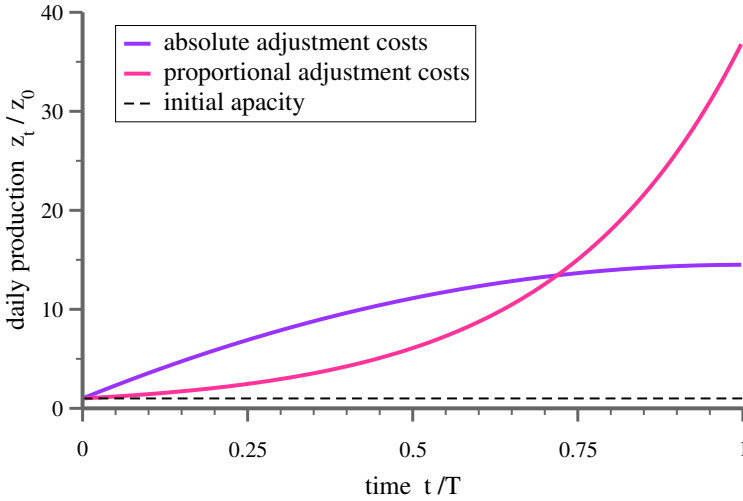[3]Another way to express assumption is that adjustments cost are homogeneous in z_t and $\dot{z}_t$.

Figure 1: The time evolution of the production capacity, z_t/z_0, the output per period relative to the starting capacity z_0. The convex (exponential) curve illustrates the equilibrium under adjustment costs based on proportional increases in capacity, see (13). The concave curve depicts the outcome with adjustment based on absolute increases, see (16). A total of $Z_t/z_0 = 10$ units have to be delivered for both cases within $[0, T]$, where T is the delivery period. The areas below the two curves are hence identical.

is mininimized by the Euler-Lagrange equation [16, 17], as given by the second equation. For the integrand, the Lagrange-function, we have

$$L(z_t, \dot{z}_t) = (p_0 - \lambda - c)z_t - a_z \left(\frac{\dot{z}_t}{z_t}\right)^2 \tag{11}$$

compare (1). Evaluating the Euler-Lagrange equations (10), we obtain

$$\left[\frac{\ddot{z}_t}{z_t^2}\right] = \frac{\lambda + c - p_0}{2a_z} + \left[\frac{\dot{z}_t^2}{z_t^3}\right] . \tag{12}$$

For an exponential solution,

$$z_t = z_0 e^{\beta t}, \tag{13}$$

one needs that

$$\frac{\beta^2}{z_0} e^{-\beta t} = \frac{\lambda + c - p_0}{2a_z} + \frac{\beta^2}{z_0} e^{-\beta t} \tag{14}$$

be fulfilled, which implies that that the Lagrange multiplier λ needs to specify a special condition, namely that $\lambda = p_0 - c$. This leaves the exponent β as a free parameter, which can then be used to satisfy the production constraint,

$$Z_T = \int_0^T z_t dt = \frac{z_0}{\beta} \left[e^{\beta T} - 1 \right], \qquad \frac{Z_T/T}{z_0} = \frac{1}{\beta T} \left[e^{\beta T} - 1 \right], \qquad (15)$$

where Z_T/T is the average production to be attained, and $Z_T/(z_0 T)$ the increase of the average production with respect to the initial capacity z_0. Solving (15) numerically, one finds, e.g., that one needs $\beta \approx 3.614/T$ in order to achieve $Z_T/(z_0 T) = 10$, viz a ten-fold increase.

For a comparison with case of absolute adjustment costs, $\sim (\dot{z}_t)^2$, we rewrite (9) as

$$\frac{z_t}{z_0} = 1 + \frac{3}{T} \left(\frac{Z_T}{z_0 T} - 1 \right) \left[t - \frac{t^2}{2T} \right], \qquad (16)$$

which shows that the production capacity remains flat if $Z_T/T = z_0$, viz when there is no need for an increase. The two time paths for the production capacity over time resulting from the two assumptions, relative of absolute, concerning the nature of the adjustment costs are illustrated in Fig. 1. It is evident that the marginal condition $\dot{z}_{t=T} = 0$ is not fulfilled for the exponential ansatz (13), which is hence valid, strictly speaking, only when $T \to \infty$, viz when the time horizon is infinite.

2.4. Realised supply paths for Covid vaccines

As mentioned above, the proportional specification of adjustment costs requires a positive finite initial capital stock or, equivalently, a finite starting level of capacity $z_0 > 0$.

One key difference between the absolute and standard adjustment costs specification is the resulting time profile for the production capacity. For the proportional specification, the capacity increases at a constant rate and is thus convex in time, t. By contrast, the time path of the capacity resulting from absolute adjustment costs is concave in time, with $\ddot{z}_t$ being always negative (and $\dot{z}_t$ positive until T) as can be seen from the relationship (9).

There is little data available on the time path of vaccine production, which companies tend to keep confidential. The same holds for the delivery schedules redacted in the contracts which the Commission has published [2, 3, 4]. However, there is ample data on the number of vaccinations implemented.
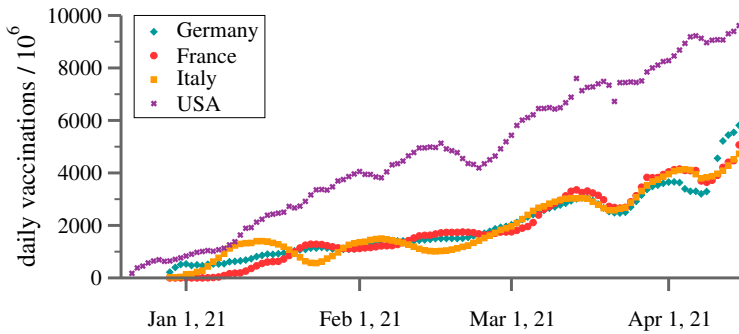
Figure 2: The evolution of daily vaccinations (proportional to population). Vaccinations actually implemented represent a good proxy for production capacity, z_t, for large entities like the US which neither import nor export significant amounts. The EU stopped exports partially after late March 2021.

For smaller countries, like Israel, most vaccines doses are imported. The availability of imports explains also the faster vaccination campaign of the UK. But for larger units, like the US, im- and exports play only a marginal role.[4] For these larger units, vaccination rates can be expected to track deliveries, with eventual organizational problems leading only to temporary delays. EU member countries agreed to place joint Advance Purchase Agreements with deliveries to be distributed on a per capita basis [18]. As a consequence, the vaccination curves of EU member countries follow similar trends. However, a substantial part of EU production was been exported until a (partial) export ban was imposed end of March 2021 [19].

Figure 2 shows the evolution of the number of daily Covid-19 vaccinations for the US and major EU countries over the first few months after approval and the start of production. The trend for the US is clearly linear, not exponential. The curve for EU countries is also linear, having however an upward kink a few weeks after the export ban went into effect in late March of 2021. This suggests that, at least in the case of a new product like Covid-19 vaccines, the adjustment costs are not proportional. We thus continue to use the absolute adjustment cost specification in the remainder of the paper.

---

[4]US regulations do not allow for exports of vaccines until the US population has been vaccinated.

2.5. Uncertainty

When the contracts for the production of Covid-19 vaccines were initially concluded in 2020, the two cost components a_z and c could not be known with certainty because vaccines in question were in part completely new products. In this section we thus introduce uncertainties with regard to the adjustment costs and the marginal cost of production.[5] The uncertainty is assumed to take a simple form:

The marginal cost, c is distributed as $\bar{c} + \varepsilon_c$, where ε_c is a suitable distribution function with finite variance and zero mean. Equivalently we assume that a_z is distributed as $\bar{a}_z + \varepsilon_{a_z}$.

Assuming risk neutrality, the problem for the firm is then to maximize the *expected value* of revenues minus adjustment costs, again subject to the overall production constraint:

$$\int_0^T p_0 z_t dt - (\bar{a}_z + \varepsilon_{a_z}) \int_0^T (\dot{z}_t)^2 dt - \lambda \left[\int_0^T z_t dt - Z_T \right] - \int_0^T (\bar{c} + \varepsilon_c) z_t dt \,. \quad (17)$$

We assume that the expectations of the firm about both cost elements are unbiased. This implies that the expected value of the functional is the same as in equation (1), only with $\bar{c}$ and $\bar{a}_z$ instead of c and a_z. This implies that the solutions derived so far should not be affected by this type of uncertainty.

We thus conclude that uncertainty about key cost parameters should not affect the time path of vaccine production capacity. For the reminder we will thus revert to the framework without uncertainty.[6]

2.6. Adjustment costs and the supply curve

We have established so far that the magnitude of the price (as long as it is constant) does not affect the time path of production (even if there is uncertainty about cost parameters). The firm will accept the contract however only if expected net revenue, $(p_0 - c)Z_T$, compensate for the adjustment

---

[5]This type of uncertainty is almost the complete opposite to the usual problem considered in the economics literature where the firm knows its own costs and faces an uncertain price over time [14].

[6]In the case of AstraZeneka the uncertainty about costs does not matter anyway since the contract specifies that AstraZeneka will supply the vaccine at production costs [2]. The EU will simply reimburse ex post the costs incurred by the company and the European Commission has to right to audit the accounts for the company.[2]

costs, E_{adj}. This condition requires that

$$(p_0 - c)Z_T = E_{\text{adj}} = 3a_z \frac{\Delta Z^2}{T} = 3a_z \frac{(Z_T - z_0 T)^2}{T^3}, \tag{18}$$

which can be obtained by substituting (8) into (5).

Neglecting the marginal cost of production c, which has been shown to be very small [20, 21], this equation also determines the total expenditure public authorities would have to sustain in a market in which firms compete for the vaccine order.[7] Total adjustment costs and thus the total expenditure for the authorities increase proportionally with the adjustment cost parameter, a_z. For $z_0 = 0$ it is apparent, ceteris paribus, that total adjustment costs incurred by the firm fall with the cube of the time it has to fulfill the entire order, T. Any initial capacity, $z_0 > 0$ reduces the cost (and this reduction falls approximately with the inverse of T. These scaling relations hold for fixed overall production Z_T.

The unit price needed to induce a firm to accept the contract is given for $z_0 = 0$ by:

$$\frac{E_{\text{adj}}}{Z_T} = 3a_z \frac{Z_T}{T^3}. \tag{19}$$

This implies that the unit price needed to compensate firms for higher adjustment costs decreases with the third power of T. The elasticity of the price with respect to the delivery time is thus equal to 3. This might explain why the contracts concluded in 2020 specified overall production targets only for the entire year 2021. A corollary of (19) is that the authorities would still have to pay four times the price when halving both the delivery time T and the total amount requested Z_T (assuming the capacity is initially negligible, $z_0 = 0$).

The overall delivery time T is not the same as the average delivery delay t_{deliver}, over the life-time of the product, which is given by

$$t_{\text{deliver}} = \frac{1}{Z_T} \int_0^T z_t \, t \, dt = \frac{1}{Z_T/T} \left[\frac{z_0}{2} + \frac{5\Delta Z}{8} \right] T. \tag{20}$$

With $\Delta Z = Z_T/T - z_0$, this result implies that for $z_0 = 0$, i.e. when the initial capacity is zero and $\Delta Z = Z_T/T$, the mean delivery delay will be

---

[7]Most producers of Covid-19 vaccines have pledged to make no profit or apply only normal profit margins [22].

$5T/8$. Instead, when $\Delta Z = 0$, viz when production is constant, the average delay would be $T/2$. The later case holds when the initial production capacity is sufficient to fulfill the entire order over time, which is however not the case for Covid-19 vaccines.

3. Social versus private costs

The social benefits resulting from vaccination campaigns differ in nature from the private costs of production. A continuing pandemic leads to elevated social costs both in terms of hospitalisations, deaths and lost output, as governments have to impose cost restrictions on mobility (so-called Non Pharmacological Interventions NPIs) [11].

For society the value of a dose depends importantly on the time it is delivered. Early delivery helps to avoid infections and allows for an earlier lifting of lockdowns. This implies that any delivery which does not occur today imposes an opportunity cost for society which is due to the economic loss of prolonged lockdowns and more infections. Each early dose thus delivers a flow of benefits in terms of avoided costs which is proportional to the time it arrives. This benefit does not materialise only when herd immunity has been reached. Every person vaccinated will reduce the potential medical costs from an infection [23].

We parametrize the per time unit costs of a continuing pandemic in terms of lost output, hospitalisations and deaths by the constant k, which has been estimated previously [24]. The opportunity cost to society of a delay in vaccine delivery (or the prolongation of the 'no vaccination' status quo) is thus k per unit of time, This cost is reduced pro rata by the part of the population which has been vaccinated. It follows that the remaining social costs are proportional to the relative number of people not vaccinated. These remaining pandemic costs E_{pandemic} are thus equal to:

$$
\begin{aligned}
E_{\text{pandemic}} &= k \int_0^T \left(1 - \int_0^t \frac{z_{t'}}{Z_T} dt' \right) dt \\
&= k \int_0^T \left(1 - \frac{3}{2} \frac{t^2}{T^2} + \frac{1}{2} \frac{t^3}{T^3} \right) dt = \frac{5}{8} kT,
\end{aligned}
\tag{21}
$$

where the solution for the time path of delivery (9) has been used, together with the simplifying assumption that $z_0 = 0$ and that the total production, Z_T, suffices to vaccinate the entire population. Over the same period the costs

to society would be kT if nobody would be ever vaccinated, which implies that ordering vaccines at a fixed price leads to a reduction of opportunity costs of 3/8 already within the delivery period, together with 100% reduction afterwards.

The aim for society should be to minimize social pandemic costs, as given by (21), taking into account the cost of ramping-up vaccine production quickly. Equation (21) shows that the social costs increase linearly with the delivery time, T; whereas the cost of production falls with the third power of T, see (18).

The overall social costs E_{social} are thus given by the sum of the pandemic costs (21) and the production adjustment costs (18),

$$E_{\text{social}} = \frac{5}{8}kT + 3a_z\frac{Z_T^2}{T^3}, \tag{22}$$

where (18) and (20) have been used for the case $z_0 = 0$. The delivery period T minimizing the sum of social and adjustment costs is hence given by

$$T_{\text{opt}} = \left[\frac{72Z_T^2}{5}\frac{a_z}{k}\right]^{1/4}. \tag{23}$$

The delivery time should in principle be set by the authorities according to this relationship. It implies that the optimal delivery time depends mainly on the ratio of two parameters, namely a_z (which denotes the cost of ramping up production quickly) and k, which denotes the cost of a continuing pandemic). It is the level of this ratio which determines the optimal delivery time. Changes in the value of the parameters involved have only a minor impact, since the optimal delivery time increases only with the fourth root of this ratio.

Equation (23) yields one observable implication, as one can rewrite it as:

$$T_{\text{opt}}^4 = T^4\frac{24}{5}\frac{3a_zZ_T^2/T^3}{kT} \qquad kT_{opt} = \frac{24}{5}p_0Z_T. \tag{24}$$

where (18) has have been used for the total expenditure $p_0Z_T = 3a_zZ_T^2/T^3$ necessary to induce firms to supply the vaccines over the time period T, which is here identical to Topt. Using (24) together with (21) yields a simple result:

$$E_{\text{pandemic}} = 3p_0Z_T. \tag{25}$$

This relationship implies that if governments had chosen the optimal delivery time, they should have spent on vaccines a sum equivalent to about one third of the cost of the pandemic during this period (equal to kT_{opt}).

We have so far considered only contracts which specify a fixed price. The optimal contract time calculated in (23) above constitutes a second best, because it is subject to this constraint. We turn to the optimal contract design after discussing the overall magnitude of the cost of the pandemic compared to the cost of production of vaccines.

4. Orders of magnitude for social and private costs

The result (25) relates two observable variables, namely the social costs and the amount spent on vaccines. Next we discuss estimates for the respective orders of magnitude.

The magnitude of the social cost of a continuing pandemic can be estimated using the available data on the economic cost of the pandemic so far [25], which have been around 4-5 percent of GDP. Reaching herd immunity allows society in consequence to avoid costs equivalent of 4-5 percent of GDP, and even more when including value of life costs [25]. This would mean that the avoided economic costs per vaccinated person would be equal to 4-5 percent of GDP per capita (or about 2600-3000 USD for the US, 1500-2000 euro for Germany).

Another approach to determine the value of a vaccination relies on surveys of the willingness to pay (WTP) expressed in standard surveys used to estimate the value of other vaccines. One study [26] concludes that the social valuation of vaccination is about 1.1 percent of the per capita gross domestic product (GDP). This would be equivalent to about 600 USD per dose for the US or 500 USD for Germany. These values constitute a lower threshold as the social value of a vaccination is likely to be substantially larger than the private value, because vaccinated individuals no longer transmit the disease to others.

The estimate of the overall cost of the Covid pandemic presented in [11] suggests a similar order of magnitude, but expressed in total amounts. It is estimated that the global total cost of the Covid pandemic is about 16 trillion USD (of which about one half is due to medical cost and the value of lives lost), resulting in a social value of about 2600 USD per vaccination

(1300 per dose if two are needed for immunity).[8] Different approaches thus yield estimates which converge to a per capita cost of the pandemic of around 1.300 USD for a single dose with a two dose treatment.

The cost of ramping up vaccine production are more difficult to pin down because manufacturers keep their costs as professional secrets. However, the prices paid by governments for Covid-19 vaccines, which have been made public only partially [29, 30], can give an indication of the costs since, as mentioned above, most manufactures have promised not to profit from the pandemic. The published prices paid by governments are generally in the region of 10-20 USD per dose for the new vaccines, see [31] for a summary of the available evidence. For example, the price of the Moderna vaccine in the US was 10 USD and 15 Euro for the Pfizer/Biontech product in Europe [29, 31]. The cost of a single dose of other vaccines is reported to be even lower, at around 3 USD per dose for AstraZeneca, for example [31].

This implies that social costs are at least between 66 and 100 times the cost of the more expensive vaccines.[9] These figure are an order of magnitude higher than the 3:1 suggested by equation (25) above, suggesting that governments should have been willing to pay for much more for a quicker delivery schedule.

These calculations have been made in per capita terms. In the light of the pandemic losses of 4 to 5 percent of GDP mentioned above, the rule of thumb incorporated in (25) could also be restated as implying that governments should have been willing to spend up to 1.5 percent of GDP on vaccines. For the US this would amount to over 300 billion USD, about ten times more than what was budgeted under the so-called Operation Warp Speed [32]. For the UK one finds similar numbers. Given a GDP of about 2 thousand billion pounds spending of up to 30 billion would have been justified, but actual spending on vaccine procurement (as opposed to administering the vaccines)

---

[8][27] provide an even higher estimate (5.800 USD), per course (two doses). [28] arrive at a similar result.

[9]When the vaccine orders were placed it was not certain which vaccines would eventually be approved and how effective they would be. Ex ante, it was thus prudent to place orders for a portfolio of candidate vaccines equivalent to 2 to 3 times the number of doses which would be needed to vaccinate the entire population. Taking into account the expenditure on vaccines candidates which, ex post, turn out not to work (or come to late), one could thus argue that the true cost of a successful vaccine is higher than the cost of the successful ones which are available today. However, even this effect would not overturn the results.

amounted to only 3.3 billion [33][10]

One can only guess at the causes for this under-spending. One reason might have been that many contracts were placed (and priced) in the summer of 2020, when the number of infections and fatalities had fallen greatly after the first wave of early 2020. Another reason might be that, ex ante, governments were afraid of a populist backlash to pay pharmaceutical companies hundreds of billions.

Fear of a populist backlash is also the reason why the producers of vaccines have been careful not to exploit their informational advantage. The producers might have been able to charge much higher prices given the high social value of the vaccines and the fact that the authorities had ex ante little information about costs. In the case of AstraZeneca the problem of asymmetric information was addressed explicitly in the contract, which allows the European Court of Auditors to audit the company's accounts [2]. Fear of a populist backlash is mentioned explicitly by [22] as a major reason why other companies also pledged to adopt 'pandemic pricing', i.e. to make no, or little profit.

5. Optimal time varying pricing

We have so far considered only contracts which specify a fixed price over the entire delivery period because this is what contracts for Covid vaccines specify in general. However, another way to draw up the contracts would have been possible. In this section we thus consider a specific contract design, namely one which takes the cost of later delivery expressively into account.

Using the expression for the opportunity costs of delay introduced above in (21), the general social planner problem, which is not constrained by a fixed price contract, is to minimize the sum of the costs of an ongoing lockdown and the adjustment costs that are necessary to accelerate production. The end point, T, represents the point in time when the entire population has been vaccinated.[11] At this point the economy would be fully back to normal

---

[10]Moreover, this sum was for close to 270 million doses. However, domestic needs can only amount to about 120 million (given a UK population of around 60 million, less if only adults are vaccinated). About one half of the total amount thus represents a surplus which is likely to be exported once the domestic vaccination campaigns has finished.

[11]Reaching herd immunity is also often mentioned as the goal of vaccination. It is usually assumed that for Covid-19 herd immunity requires that about 70 percent are vaccinated.

and the costs parametrized by k no longer arise. For our considerations we normalise $Z_T = 1$.

Denoting total social costs by W_{social}, the social planner then minimizes

$$W_{\text{social}} = k \int_0^T (1 - Z_t)\, dt + a_z \int_0^T (\dot{z}_t)^2 dt - \lambda \left[\int_0^T z_t dt - 1 \right], \qquad (26)$$

where $Z_t = \int_0^t z_{t'}\, dt'$ is the number of vaccines produced hitherto, which is assumed to correspond to the proportion of people already vaccinated at that point. Using partial integration, above expression for W_{social} can be rewritten in terms of the opportunity costs of gradual delivery, which are proportional to the time one waits for the delivery of the vaccine:

$$W_{\text{social}} = k \int_0^T t z_t dt + a_z \int_0^T (\dot{z}_t)^2 dt - \lambda \left[\int_0^T z_t dt - 1 \right]. \qquad (27)$$

We now show that the problem for the social planner can be made isomorphic to that of the firm.

The key variable for the firm is the price, or revenue per unit produced. Here we absorb the production costs per unit, c, into the price $p = p_t$. For fixed price contracts, as considered in the previous sections, the price $p_t \to p_0$ does not vary with the time the vaccine is delivered. This can be changed if the authorities offer a time varying price which declines from a certain initial level, denoted by p_0. The total revenues of the firm are then given by

$$\text{Revenue} = \int_0^T p_t z_t dt. \qquad (28)$$

One can relate the social planer problem of minimizing the cost of the pandemic, as given by (27), to the problem of a firm maximizing profits. For both, the social planer and the firm, the problem has to be solved taking into account adjustment costs.

The social optimum can be reached if the price path facing the firm coincides with the minimization of the pandemic costs, i.e. if

$$p_t = p_0 - kt, \qquad (29)$$

Here we use the stricter criterion of an essential 100% vaccination.

where p_0 now denotes the 'base price', which diminishes linearly over time. The problem facing the firm then becomes to maximize total revenues minus adjustment costs:

$$\int_0^T (p_0 - kt)z_t dt - a_z \int_0^T (\dot{z}_t)^2 dt - \lambda \left[\int_0^T z_t dt - 1 \right] , \qquad (30)$$

which can be rewritten as

$$p_0 - k \int_0^T t z_t dt - a_z \int_0^T (\dot{z}_t)^2 dt - \lambda \left[\int_0^T z_t dt - 1 \right] . \qquad (31)$$

Comparing equations (27) and (31) shows that they must lead to the same solution viz to identical time paths for z_t. The firm *maximizes* the difference between revenues and adjustment cost, with unit revenues declining linearly over time. Society *minimizes* total costs, which comprise the same adjustment costs, but taking also into account that the costs of delayed delivery are linear in time. With the pricing schedule (29), equations (31) and (30) represent hence the same problem, except for the constant term p_0, which implies that they have the same solution z_t. The size of the initial price, p_0, has no implications for the decision of the firm regarding how quickly to increased capacity. Note also, that the sign of the Lagrange parameter λ is irrelevant.

The implication is that the pricing schedule (29) can induce firms to adopt the speed of increase in production capacity which is also optimal from a social point of view. There is thus a way to align private and public interests by specifying a pricing schedule which mimics the social cost of a continuing pandemic.

The base price p_0 determines, as before, whether the firm makes a profit or a loss, taking into account adjustment costs. The optimal contract thus involves a base price which allows the firm to break even and a premium for early delivery, which declines over time.

We also note that the pricing schedule (29) remains optimal from a social welfare point of view even if there is uncertainty about adjustment costs, which would affect the optimal schedule in exactly the same way for a cost minimizing firm as for a social planner.

Given that k was shown to be large, because social costs affect the entire economy, this strategy may lead to a high premium, which would however decline rapidly. With such a pricing schedule there would be no need to

specify intermediate delivery dates (as done in existing contracts). Firms would have the incentive to ramp up production as quickly as required by society.

6. Conclusions

Our analysis starts from the observation that delays in the availability of vaccines are costly for society. A dose delivered one quarter later is substantially less valuable than a dose delivered today. The costs of delay were particularly high in the first half of 2021 as the pandemic continued in the US and most of Europe, forcing governments to implement lockdowns that depressed the economy. However, the urgency to speed up production was not recognised in the existing contracts, which specify mostly only a fixed quantity and an overall time frame, typically the entire year of 2021. In the absence of incentives to produce early, firms will tend to minimize adjustment costs, i.e. the costs resulting from ramping up production. In this case, firms will prefer to increase production capacity only gradually.

Our analysis shows that the lack of incentives to produce early does not derive from a potentially low *level* of the price offered to companies, but on its *time path*. With the existing, fixed price contracts, a dose delivered the subsequent quarter yields the same revenue for the producer as a dose delivered immediately, but for society there is a sizable difference. The practical problem is then how to provide incentives for early delivery.

The most direct approach to reduce the discrepancy between private and social incentive would be to make the price fully variable over time. We show that it is straightforward to design an optimal contract, which aligns the time paths of the price with that of the social value of a vaccination. In this case a linearly decreasing price schedules replicate the social optimum.

From our perspective there is a clear policy conclusion: Supply contracts for vaccines should contain incentives for accelerated production. Vaccines delivered early should command a higher price.

A bi-product of our analysis is that the details of the specification of adjustment costs matter. The usual specification of adjustments based on the proportional increase in capital leads to a drastically different time profile of capacity than the alternative of basing adjustment costs on the absolute increase in capacity. We argue that the latter is more appropriate for an entirely new product, like Covid-19 vaccines and show that the data on vaccine production aligns well with the prediction from our specification.

We also find that even within the sub-optimal types of contracts actually concluded governments have not sufficiently recognized the need for speed. Our framework yields the rule of thumb that governments should have been willing to spend on speedy vaccine production an amount equal to about one third of the economic losses caused by the pandemic. Actual expenditure has been an order of magnitude lower.

Funding

Acknowledgement

Declarations of interest

None.

References

[1] Nikhil Vellodi and Joshua Weiss. Optimal vaccine policies: Spillovers and incentives. *Covid Economics*, 65, 2021.

[2] European Commission. Advance Purchase Agreement with AstraZeneca, 2021.

[3] European Commission. Advance Purchase Agreement with Curevac, 2021.

[4] European Commission. Advance Purchase Agreement with Moderna, 2021.

[5] Susan E Martonosi, Banafsheh Behzad, and Kayla Cummings. Pricing the covid-19 vaccine: A mathematical approach. *Omega*, page 102451, 2021.

[6] Jens Grauer, Hartmut Löwen, and Benno Liebchen. Strategic spatiotemporal vaccine distribution increases the survival rate in an infectious disease like covid-19. *Scientific reports*, 10(1):1–10, 2020.

[7] Roy M Anderson, Carolin Vegvari, James Truscott, and Benjamin S Collyer. Challenges in creating herd immunity to sars-cov-2 infection by mass vaccination. *The Lancet*, 396(10263):1614–1616, 2020.

[8] Zhen Wang, Chris T Bauch, Samit Bhattacharyya, Alberto d'Onofrio, Piero Manfredi, Matjaž Perc, Nicola Perra, Marcel Salathé, and Dawei Zhao. Statistical physics of vaccination. *Physics Reports*, 664:1–113, 2016.

[9] Eve Dubé, Caroline Laberge, Maryse Guay, Paul Bramadat, Réal Roy, and Julie A Bettinger. Vaccine hesitancy: an overview. *Human vaccines & immunotherapeutics*, 9(8):1763–1773, 2013.

[10] Amrita Ahuja, Susan Athey, Arthur Baker, Eric Budish, Juan Castillo, Rachel Glennerster, Scott Kominers, Michael Kremer, Jean Lee, Candice Prendergast, et al. Preparing for a pandemic: Accelerating vaccine availability. Technical report, National Bureau of Economic Research, 2020.

[11] David M Cutler and Lawrence H Summers. The covid-19 pandemic and the $16 trillion virus. *Jama*, 324(15):1495–1496, 2020.

[12] Roxanne Khamsi. If a coronavirus vaccine arrives, can the world make enough. *Nature*, 580(7805):578–580, 2020.

[13] Russell W Cooper and John C Haltiwanger. On the nature of capital adjustment costs. *The Review of Economic Studies*, 73(3):611–633, 2006.

[14] Andrew B Abel and Janice C Eberly. An exact solution for the investment and value of a firm facing uncertainty, adjustment costs, and irreversibility. *Journal of Economic Dynamics and Control*, 21(4-5):831–852, 1997.

[15] Daniel S Hamermesh and Gerard A Pfann. Adjustment costs in factor demand. *Journal of Economic literature*, 34(3):1264–1292, 1996.

[16] John L Troutman. *Variational calculus and optimal control: optimization with elementary convexity*. Springer Science & Business Media, 2012.

[17] C. Gros. *Complex and adaptive dynamical systems: A primer*. Springer, 2015.

[18] Statement of the European Commission on the methodology used to determine the allocation of doses of vaccines under the Advance Purchase Agreements, 2021.

[19] Rory Watson. Covid-19 vaccines: Europe tightens export criteria in bid to secure more doses, 2021.

[20] Zoltán Kis, Cleo Kontoravdi, Antu K Dey, Robin Shattock, and Nilay Shah. Rapid development and deployment of high-volume vaccines for pandemic response. *Journal of Advanced Manufacturing and Processing*, 2(3):e10060, 2020.

[21] Zoltán Kis, Cleo Kontoravdi, Robin Shattock, and Nilay Shah. Resources, production scales and time required for producing rna vaccines for the global pandemic demand. *Vaccines*, 9(1):3, 2021.

[22] Where do COVID-19 vaccine players stand on pricing? So far, it's no profit slight profit or undecided, 2020.

[23] Claudius Gros, Roser Valenti, Lukas Schneider, Benedikt Gutsche, and Dimitrije Marković. Predicting the cumulative medical load of covid-19 outbreaks after the peak in daily fatalities. *Plos one*, 16(4):e0247272, 2021.

[24] Claudius Gros, Roser Valenti, Kilian Valenti, and Daniel Gros. Containment efficiency and control strategies for the corona pandemic costs. *Scientific Reports*, 11:6848, 2021.

[25] Claudius Gros, Roser Valenti, Lukas Schneider, Kilian Valenti, and Daniel Gros. Containment efficiency and control strategies for the corona pandemic costs. *arXiv preprint arXiv:2004.00493*, 2020.

[26] Leidy Y García and Arcadio A Cerda. Contingent assessment of the covid-19 vaccine. *Vaccine*, 38(34):5424–5429, 2020.

[27] Juan Camilo Castillo, Amrita Ahuja, Susan Athey, Arthur Baker, Eric Budish, Tasneem Chipty, Rachel Glennerster, Scott Duke Kominers, Michael Kremer, Greg Larson, et al. Market design to accelerate covid-19 vaccine supply. *Science*, 371(6534):1107–1109, 2021.

[28] Michele Kohli, Michael Maschio, Debbie Becker, and Milton C Weinstein. The potential public health and economic value of a hypothetical covid-19 vaccine in the united states: Use of cost-effectiveness modeling to inform vaccination prioritization. *Vaccine*, 39(7):1157–1164, 2021.

[29] EU agreed 15.50 euros per dose for Pfizer vaccine, 2021.

[30] Comparing COVID-19 Vaccines: Timelines, Types, 2021.

[31] Owen Dyer. Covid-19: Countries are learning what others paid for vaccines. *BMJ: British Medical Journal (Online)*, 372, 2021.

[32] Congressional Research Service. Operation Warp Speed Contracts for COVID-19 Vaccines and Ancillary Vaccination Materials, 2021.

[33] Congressional Research Service. UK spending on Covid vaccines hits nearly £12bn watchdog says, 2021.

Failing young and temporary workers: The impact of Covid-19 on a dual labour market[1]

Carolina Nunes,[2] Bruno P. Carvalho,[3] João Pereira dos Santos,[4] Susana Peralta[5] and José Tavares[6]

This paper analyses the impact of the pandemic crisis in a dual labour market, using monthly data covering the universe of individuals registered as unemployed in 278 Portuguese municipalities, between March and August 2020. Our event studies document a large causal impact of the pandemic of up to 40 percentage points increases in year-on-year growth rates of the monthly stock of unemployed. We document the asymmetric impact of the crisis by employing triple difference-in-differences. Younger workers, below the age of 35, are between 20% and 25% more likely to be unemployed, vis-à-vis the older than 55. Middle educated individuals are at least 15% more likely to be unemployed when compared to the highly educated ones. There are no differences across genders for transitions into unemployment, but women have lower job placements than men. The effects are exacerbated by the duality of the labour market: an increase of one standard deviation in the municipal share of temporary contracts causes a rise of 11.6% in the number of unemployment registries and magnifies the socio-demographic asymmetries. Our results can be interpreted as a lower bound of the impact of the crisis on the labour market, given the furlough scheme implemented in Portugal.

1   Introduction

The impact of the Covid-19 crisis on the labour market depends on the policy responses implemented in each country. While in the US, unemployment grew to over 30 million in just one month, most European countries saw a much subtler increase in the number of jobless claims, thanks to furlough policies implemented, e.g., in France, Germany, the Netherlands, Portugal, Sweden, Spain, Switzerland, and the UK (Mayhew and Anand, 2020). By May 2020, about 50 million workers were supported by such job retention schemes (OECD, 2020). In Portugal, 1.2 million workers were covered by a furlough scheme, which amounts to one fourth of the workforce (Banco de Portugal, 2020). Most countries, including Portugal, prohibit firms that benefit from furlough schemes from dismissing workers while the support lasts.[1] This avoids massive hiring costs post-crisis, provides for partial income insurance of the workers, and eases the recovery by keeping the so-called *matching capital* between workers and firms (Dias et al., 2020). The trade-off is that, in the medium run, these policies are bound to create lock-in effects of workers in unproductive firms. Therefore, job separations can only occur in firms that do not benefit from support, or through terminations of temporary contracts in the remaining ones, thus leaving the self-employed and temporary workers in a vulnerable position (Mayhew and Anand, 2020).

In this paper, we use administrative data from a Portugal, a Eurozone country severely affected by the pandemic crisis (GDP contraction of 7.6% in 2020), with a segmented labour market (with a higher share of temporary employment, 22%, which is twice the OECD average, and a very high gap in employment protection between permanent and non-permanent workers), and a generous coverage of the furlough policy.[2]

We address three related research questions. First, using the event study difference-in-differences approach borrowed from Carvalho et al. (2020b), we obtain causal estimates of

---

[1]In Portugal, other policies, such as state guaranteed emergency credit lines, also prohibit dismissals.

[2]The duality is shared by labour markets in a few other Southern European countries and Poland. In 2019, Portugal had the third largest Index of Strictness in Employment Protection (3.14), which compares with an average of 2.11 for OECD countries. Literature has shown lower job security from temporary contracts is associated with lower job satisfaction (Aleksynska, 2018), lower cumulative wages in the long-term (Fauser, 2020) and higher depression and anxiety (Virtanen et al., 2005).

the impact of the pandemic crisis on jobless claims, job placements, and the reasons that motivated the jobless claim, during the first six months of the pandemic, i.e., between March and August 2020. Second, we investigate the uneven impacts of the crisis in the workforce, using a triple-difference approach along the gender, age, and education level of the workers. Third, we analyse the extent to which the asymmetric impacts are driven by the dual nature of the labour market, by exploiting municipal differences in the pre-pandemic share of temporary work, in a triple-difference set-up. Our data covers all individuals formally registered as unemployed with the Instituto do Emprego e da Formação Profissional, which is the Portuguese Public Employment Service, aggregated at the municipal level, in all the 278 mainland municipalities, between October 2016 and August 2020.

Our main findings are as follows. Covid-19 caused a rise on year-on-year growth rates of registered unemployment from 27 percentage points in April up to 38 percentage points in July, and a severe drop of 63 percentage points in new job placements in April. We also find that between March and May there was a sharp increase in the number of jobless claims filed because workers were dismissed from their (permanent or temporary) job or because their temporary contract ended. There was also a small increase in claims filed by mutual agreement on April. On the opposite direction, the registration of former inactive workers and employees that voluntarily quit their previous job declined. Thus we find evidence that both demand and supply side factors influenced job transitions during the pandemic.

We also demonstrate that the impact on unemployment was 20.8% and 25.8% larger for workers who are less than 25 years old, and between 25 and 34 years old, respectively, when compared with the impact for older than 55. We document an inverted u-shape impact of education on unemployment, with the highest impact concentrated on individuals with lower (15%) and upper secondary education (17.5%), *vis-à-vis* the highly educated ones. While we find no evidence of gender differences in unemployment, we document an additional drop in new job placements of women compared to men.

The effects are accentuated in municipalities with a higher share of temporary employ-

ment: an increase of one standard deviation (i.e., 8%) in this share causes a rise of 11.6% in the number of unemployment registries. This effect is driven by female, younger and middle educated workers. In particular, an increase of one standard deviation in the share of temporary contracts causes a rise of 12.6%, between 14% and 17%, and between 12% and 13% in the number of unemployed people who are female, younger than 34, and have secondary education. We interpret these results as evidence that furlough schemes do not insure some segments of the labour market against the negative shock of the pandemic crisis.[3]

The literature on the economic impacts of Covid-19 has uncovered large impacts on the labour market.[4] Alstadsæter et al. (2020) explore find that 12% of the labour force filed jobless claims in the first weeks of the crisis in Norway. Cajner et al. (2020) analyse US administrative payroll data and show that aggregate employment decreased by 21% through late-April, with slight signs of recovery only by late-June. Various authors documented severe drops in job posts (Bamieh and Ziegler, 2020; Hensvik et al., 2020) and a decrease in hirings on the period following lockdown (Betcherman et al., 2020). Using job vacancy data collected in real-time by the Burning Glass Technologies platform, Forsythe et al. (2020) find that, in the US, job postings collapsed by 44% between February and April 2020. Coibion et al. (2020) used scan data in April 2020 to show that job loss in the US was larger than new unemployment claims, with many workers moving into inactivity. For Canada, Jones et al. (2020) found new vacancies recovered in June, from 50% to around 80% of the pre-pandemic level.

A number of papers present convincing evidence of the unequal labour market impacts of the pandemic on temporary workers. Casarico and Lattanzio (2020) uses administrative data on a sample of contracts for the first quarter of 2020 in Italy to find that temporary workers are 8 p.p. more likely to lose their job, contrary to older and highly educated workers, who are more protected against job loss. Papers based on survey data that con-

---

[3]Ferreira et al. (2020) monitor the policies implemented in Portugal and highlight that they do not address the income losses of all vulnerable groups. Boeri and Brücker (2011) show that firms that rely heavily on temporary employment are less likely to take-up furlough policies, possibly because it is less costly for them to rely on the dismissal margin to adjust to the shock.

[4]Table B.8, in appendix, summarizes the approaches, set-up, and main findings of selected literature on the impact of the pandemic on the labour market.

firm the disproportionate effect on temporary workers include Adams-Prassl et al. (2020), Kikuchi et al. (2021) and Aum et al. (2020) in the UK, USA, and Germany, Japan, and South Korea. Younger workers are shown to be the most affected in Canada by Lemieux et al. (2020), Japan by Kikuchi et al. (2021), and the US by Cho and Winters (2020), Cortes and Forsythe (2020) and Montenovo et al. (2020). Female workers are found to be disproportionally affected by Kikuchi et al. (2021) in Japan and Cortes and Forsythe (2020) in the US. Survey data also indicates that less educated, lower income and minority workers in the US are more affected (Cho and Winters, 2020).[5] Alon et al. (2020) find that women were more struck by the crisis in the US, contrary to Hupkau and Petrongolo (2020) who find no significant job loss differences between genders in the UK.

Two papers follow an approach close to ours. Meekes et al. (2020) implements a triple difference-in-differences strategy with administrative data from Statistics Netherlands and document large impacts on the employment, working hours, and hourly wages of non-essential workers, particularly the female ones, and on employment and working hours of essential workers who are single parents. Kalenkoski and Wulff (2020) implement triple difference-in-differences specifications using the US Current Population Survey and find that the impact on employment and working hours was larger for coupled women than for coupled men, and smaller for single women than for single men.[6]

The paper offers three main contributions. Firstly, we provide causal estimates of the impact of the Covid-19 pandemic in the labour market prospects of different worker groups, along gender, education, and age. The existing causal evidence so far is focused on the differential impacts along gender and marital status (Kalenkoski and Wulff, 2020), and gender and type of occupation (Meekes et al., 2020). Secondly, we rely on administrative data that cover the universe of newly registered unemployed workers in Portugal, while most of the papers rely on survey data, with the exception of Casarico and Lattanzio (2020) and Meekes et al. (2020). Thirdly, this paper is the first to show how the duality of the labour market, as captured by the share of non-permanent workers per municipality,

---

[5] Montenovo et al. (2020) add that workers in jobs that are more compatible with remote work fared better.

[6] Cho and Winters (2020) use a similar strategy on the Current Population Survey, and show that employment decreased more in metropolitan areas.

magnifies the impact and the asymmetry of the shock. Fourthly, in terms of policy implications, we provide evidence that policies such as furlough schemes fail to insure more fragile workers, which calls for the design of more inclusive and targeted policies, that tackle the asymmetric impacts in dual labour markets.

The remainder of this paper is structured as follows. Section 2 briefly describes the institutional background and details the evolution of the pandemic and policy responses in Portugal. Section 3 clarifies the data and methodology used. Section 4 analyses the overall results and heterogeneous impacts. Section 5 combines the effects of the pandemic and the labour market segmentation. Finally, Section 6 concludes.

2   Institutional Background and Covid-19 in Portugal

In this section, we provide an overview of the main characteristics of the Portuguese labour market as well as information on the timing of, and the policy responses to, the Covid-19 pandemic in the country.

2.1   The Portuguese labour market

In the last decades, the Portuguese labour market has witnessed *(i)* an increase in the education levels of the working population, *(ii)* a higher proportion of female employment, and *(iii)* an ageing of the labour force (Portugal et al., 2018). In 2011, the country requested financial assistance to the European Commission, the European Central Bank and the IMF and, until 2015, it implemented several labour market reforms, aimed at mitigating its dual character and rigidity.

Permanent workers in Portugal benefit from one of the highest levels of employment protection across the OECD, which, according to the European Commission, increases both the reluctance of firms to hire permanent employees (European Commission, 2018) and the importance of non-permanent employment (OECD, 2017). Moreover, the strong downward nominal wage rigidity stemming from legal restrictions on nominal wage cuts, combined with the long recent period of low inflation, leave employers with little margin

to adjust real wages. As a consequence, in periods of crisis, employment (and especially temporary employment) becomes the main margin of adjustment (Martins and Portugal, 2019; Carneiro et al., 2014). Given the concentration of low skilled workers in this group, they face significant difficulties in finding a new job (Blanchard and Portugal, 2017), and are prone to be long-term unemployed. In sum, the characteristics of the labour market place some workers in a more vulnerable position when the country faces a period of economic downturn.

Albeit to different extents, the prevalence of temporary employment is also relevant in other countries, even more so since the 2008 crisis. In most countries of the Eurozone, with the exception of Greece and Spain, employment in 2017 had mostly reached the 2006 levels; however, the composition of employment changed: the share of temporary workers has increased Weel (2018).

Figure 1 shows that temporary employment accounts for 22% of all dependent employment in Portugal, above the OECD average of 12%, and only exceeded by Spain (27%) and Poland (24%).

Figure 1: Share of temporary employment (% of dependent employment), 2018

Source: OECD

The prevalence of temporary employment varies within different groups of the working population. Using data from the Labour Force Survey of 15 EU countries, between 2006 and 2009, Nunez and Livanos (2014) document that the temporary employment is higher

among the youngest (particularly aged 20-25), and middle educated individuals below 35 years old, while the prevalence across genders is similar.[7] In Section 5 we further discuss the differences in the prevalence of temporary employment across population subgroups in Portugal.

Prior to the pandemic, unemployment in Portugal had been decreasing, from a peak of 16.2% attained in 2013. The global financial crisis, followed by the international assistance program, penalized employment heavily, particularly for the younger adults. In 2019, the unemployment rate was 6.5%, the lowest since 2003.

Unemployed individuals typically register in Public Employment Services, as this is a necessary condition to receive unemployment benefits. The minimum age for registering is 16 years old. After registration, individuals have access to active labour market policies, such as professional training sessions and internships. In return, they have the obligation to comply with a personal employment plan and actively seek employment by their own means. The Public Employment Service collects job offers from firms and advertises them both online and through a vast network of local offices throughout the country. When offered a suitable or socially necessary job, unemployed individuals cannot reject the offer, otherwise they risk losing the unemployment benefit.[8] Jobless individuals have two incentives to register with the Employment Office. On the one hand, registration and abidance by the rules of the Office are necessary conditions to receive unemployment benefits; on the other hand, even individuals who do not qualify for the unemployment benefit may benefit from registering, due to the active labour market policy interventions provided by the Office that are not available to unregistered people.

According to data from Statistics Portugal, registrations in Public Employment Services cover a large majority of the unemployed population. Between 1999 and 2019, the number of individuals registered at the Public Employment Services represented an average of 94% of the unemployed population. This figure averages out asymmetries in

---

[7]The authors also discuss the four main reasons that drive temporary employment (signaling, or screening mechanisms, flexible contract or market friction) and argue that in Europe the main reason behind temporary employment is labor market frictions, particularly in the Mediterranean countries (Portugal, Greece and Spain).

[8] Socially necessary jobs are temporary work opportunities filled by unemployed individuals to satisfy social or collective needs of non-profit public or private entities.

different periods. For instance, during the crisis, in 2011 and 2012, the share of unemployed registered at the Public Employment Services is 77%.

2.2 Covid-19 in Portugal and Policy Responses

The first cases of Covid-19 in Portugal were officially confirmed on March 2[nd], in the North of the country. In the following days, the government closed schools and imposed circulation restrictions on the border with Spain. On March 18[th], the President declared the State of Emergency, which lasted until May 3[rd], when it was substituted by a less severe but still fairly constrained State of Calamity. As a consequence, all non-essential services were shutdown, with the exception of supermarkets, pharmacies, and gas stations. Restaurants were closed and only allowed to serve take-away. Further restrictions on circulation and mandatory homeworking for compatible jobs were also decreed.[9]

The Portuguese Covid-19 special furlough scheme was implemented on March 26[th].[10] This measure allows firms whose activity has been affected by the crisis to suspend employment contracts or reduce working hours of employees, with social security covering two-thirds of the wage. Administrative data from Social Security shows that more than 1 million workers were supported, a large number concentrated in the Retail and Hospitality sectors.[11]

Firms that benefited from the furlough policy *could not* dismiss workers on permanent contracts; they could, however, choose not to renew temporary contracts, stop hiring independent workers with no formal job contract to the firm, and not replace workers who depart voluntarily.[12] The bulk of the effects on unemployment should arise from these two margins that could be applied by recipient firms, or from firms that did not qualify or chose not to apply for the furlough scheme.

---

[9]Portugal is one of the European countries where self-imposed social distancing started earlier, with people avoiding to go out to restaurants eight days before the government mandated its closure (Midões, 2020).

[10] This new furlough scheme is a simplification of the pre-pandemic one, which was more complex, lengthy and restrictive.

[11] Appendix Figure D.1 presents the evolution in the number of firms and workers under the furlough policy between April and October 2020.

[12]Firm survey data from Statistics Portugal collected in 2020 shows that 77% of the firms that took up the furlough scheme claim that they would have reduced employment by an average of 19%.

Survey data from Statistics Portugal shows that there were around 1 million remote workers in the Spring of 2020, mostly highly educated and high-income ones. Their average wage is 50% higher than those of the non-remote workers and 70% of the remote workers have a higher education degree. Carvalho et al. (2021) show that the professions that are more compatible with remote working are less prevalent in the sectors that rely more on the furlough scheme, which are also that were more affected by the crisis. This may contribute to exacerbate the asymmetric effects of the crisis that we identify in this paper.

It is important to mention that the rules to qualify for unemployment benefits were eased as of July 25. Permanent or temporary employees whose contract lasted between 180 and 360 days in the 24 months prior to the unemployment date, were entitled to unemployment benefits. This relaxed the usual rule of 360 days. Since our period of analysis ends in August, this legal change is unlikely to affect the flows of unemployed individuals. Moreover, mandatory job search and training sessions were suspended between March and May, with no consequences for the benefit receivers.

The closures of businesses and services had a strong impact on the Portuguese economy. Using data from electronic purchases, Carvalho et al. (2020b) provide causal estimates of the impact of the lockdown on consumption, documenting a decrease in year-on-year growth rates of 16, 37, and 28 percentage points on overall purchases in March, April and May, respectively. As the authors also show, the impact was very uneven across sectors. The Hospitality sector (including restaurants, coffee shops and accommodation), Fashion and Beauty, and Transportation were among the most affected.[13] Importantly, as shown in Peralta et al. (2021), in 2018, the incidence of non-permanent contracts in these sectors was larger than for the overall economy. Restaurants concentrated more workers with lower education and more foreigners. Average wages in the Hospitality, and Fashion and Beauty sectors were also smaller than the national average. At the same time, Carvalho et al. (2020a), show that the sectorial composition of the municipal economies

---

[13]This is also shown in a report from SIBS Analytics and *Instituto Superior de Economia e Gestão*, available here.

leads to asymmetries in the geographic effect of the crisis.[14] The fact that some sectors, regions and workers have been affected more by the crisis paves the way for the uneven impacts that we document below.

3   Data and Empirical Methods

In this section, we describe our empirical strategy. We begin by discussing the data sources used in the paper. Then, we carefully explain the empirical methodology applied in order to obtain the causal impact of the pandemic on unemployment and new job placements.

3.1   Data

We use administrative data from *Instituto do Emprego e Formação Profissional* (IEFP), the Portuguese Public Employment Service.[15] The institution manages employment offers and mediates the matching between employers and the unemployed. IEFP provides monthly data on the number of unemployed individuals registered and the number of new job placements that take place for those registered at one of the job centers distributed across the country.

   Our sample comprises data on the 278 municipalities of the Portuguese mainland between October 2016 and August 2020. The month of September of the four years is not used, given our identification strategy explained below.[16] Data on the number of registered unemployed individuals is split into several dimensions, namely gender, age group and education level. Data for job placements is disaggregated by gender. We also use data on the main reasons for registration at the job centers[17], i.e., dismissals, voluntary quits, mutual agreement dismissals, end of temporary jobs, self-employment

---

[14] 2020a results are reminiscent of Cho et al. (2020), who find that the employed-at-work rate decreases more in larger metropolitan areas than in non-metropolitan areas in the US

[15] Unemployment can be measured from Labour Force Surveys or administrative data from registrations in Public Employment Services. As the former are not representative at the municipal level, we rely on the latter. As discussed in Subsection 2.1, data from registrations on Public Employment Services covers 94% of the unemployed population.

[16] Portugal is divided in 308 municipalities, 278 in the Portuguese mainland and 30 in the Autonomous Regions of Madeira and Azores. IEFP only provides data at the municipality level for mainland Portugal.

[17] A broad category of "other reasons" was dropped from the data as it represents residual situations, such as re-registrations after non-compliance with requirements, being an ex-migrant or reaching the end of military service.

or former inactivity.[18] It is worth noticing that while data on unemployment refers to the situation at the end of each month (*stock*), data on job placements and the motives for registering at IEFP refers to the movement throughout the month (*flow*). Summary statistics of all variables for the average municipality are provided in Table 1.

Table 1: Descriptive statistics

	Obs.	Mean	St. Deviation	Min.	Max.
Unemployment (stock)					
Total	12232	1273.6	2359.5	23	25796
by gender					
Male	12232	567.4	1091	12	13000
Female	12232	706.2	1274.8	7	13895
by age					
Less than 25 years old	12232	137.1	218.6	1	2961
Between 25 and 34 years old	12232	238.2	442.4	1	5953
Between 35 and 54 years old	12232	549.8	1069.5	8	12163
More than 55 years old	12232	348.5	650.7	5	6437
by education					
Primary education ($1^{st} - 4^{th}$ grade) or less	12232	315.4	585.3	5	7341
Basic Education ($5^{th} - 6^{th}$ grade)	12232	187.6	356.5	2	4690
Lower Secondary ($7^{th} - 9^{th}$ grade)	12232	250.5	441.7	4	4928
Upper Secondary ($10^{th} - 12^{th}$ grade)	12232	342	635.2	5	7181
Higher Education	12232	178.1	412.1	0	6157
Job Placements (flow)					
Total	12232	23.9	35.3	0	494
by gender					
Male	12232	11.1	17.9	0	234
Female	12232	12.8	19.2	0	273
Motive to register at IEFP (flow)					
Dismissed from previous job	12232	18.8	36.9	0	831
Voluntarily quit previous job	12232	7.3	11.8	0	132
Mutual agreement dismissal	12232	5.3	12.2	0	189
End of temporary job	12232	73.2	135.3	0	2625
Former inactive worker	12232	15	26	0	369
Self-employed	12232	1.5	3.4	0	48
Share of temporary contracts (2018)	12232	0.33	0.08	0.16	0.69

The average number of adults registered in Public Employment Services per munic-ipality is 1274, with a minimum of 23 and a maximum of around 26 thousand. The number is higher for females and adults between 35 and 54 years old, when compared to males and other age categories, respectively. Registrations also vary according to the level of education, with those with upper secondary education having the highest number

---

[18] Former inactive workers are workers who were out of the labour force for a period of time and start to actively seek employment again.

of registrations. The average number of job placements per month and municipality is 24.

In terms of motives to register at the IEFP, end of temporary job is the most frequent (on average 73 people per municipality), followed by dismissal from previous job (19) and the registration of formerly inactive worker (15). Finally, we note that job placements are very low when compared to the flow of registrations at the Public Employment Service.

We complement the unemployment data with the share of workers with temporary contracts in the private sector of each municipality in 2018, retrieved from PORDATA, based on data from *Quadros de Pessoal*, a linked employer-employee dataset covering all private-sector firms based in Portugal with at least one wage earner. The average share of temporary employment contracts is 33%, with a standard deviation of 8%.

3.2 Methodology

We follow the identification strategy in Carvalho et al. (2020b) and implement a difference-in-differences (DiD) event study. On Subsection 4.3 and Section 5, we use triple difference-in-differences to assess heterogeneous effects.

Our identification strategy is easily explained analysing Figure 2. The treatment and comparison groups are sets of months. The treatment group is represented by the blue lines in both panels, i.e., it comprises the months between October 2019 and August 2020. The comparison group comprises the remaining lines, i.e., the same sequence of months lagged one, two, or three years, respectively. The treatment period includes the months between March and August for all the years between 2017 and 2020. The identifying assumption is that, absent the pandemic, the year-on-year growth rate between the month of March (resp., April, May, June, July, and August) 2020 and the corresponding month in 2019 would be equal to a weighted geometric mean of the year-on-year growth rates for the same month, lagged between 1 and 3 years.

Figure 2: Identification Strategy

(a) Unemployment

(b) New Job Placements
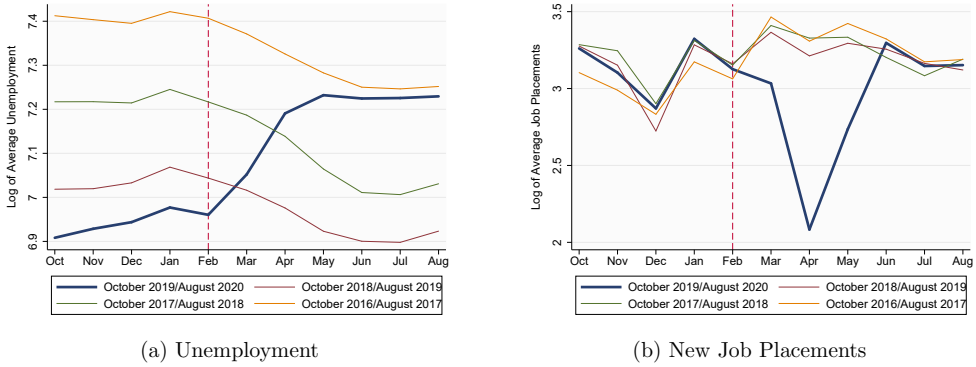
We begin by estimating the following event study equation:

$$ln(y)_{imt} = \gamma_i + \delta_m + \lambda \, \mathbb{1}_T + \sum_{m \in \{1,3,...,12\}} \beta_m \mathbb{1}_T + \epsilon_{imt} \qquad (1)$$

where $ln(y)_{imt}$ corresponds to natural log of unemployment, new job placements or the motives for new registrations at the job centers of municipality i, in month m and year $t \in \{2016, 2017, 2018, 2019, 2020\}$. Municipal, γ_i, and month, δ_m, fixed effects are also included. Additionally, $\mathbb{1}_i$ is an indicator variable for the municipality, where $i \in \{1, ..., 278\}$, and $\mathbb{1}_m$ is an indicator for the month m. February is the omitted month, since it is the one just before the start of the crisis.

The treatment indicator, $\mathbb{1}_T$ takes the value one for the months between October 2019 and August 2020. As explained above, the identifying assumption for the estimation of (1) is that, if the pandemic had not occurred, the monthly year-on-year change between March 2020 and March 2019 would have been parallel to a weighted geometric mean of the year-on-year change of the previous three years for the same month, and analogously for the remaining months between April and August. Thus the parallel trend assumption implies that $\hat{\beta}_1, \hat{\beta}_{10}, \hat{\beta}_{11}$, and $\hat{\beta}_{12}$ must not be statistically different from zero.

The error term is given by ϵ_{imt}. Standard errors are clustered at time period (month, year) and NUTS II level.[19]

---

[19] Mainland Portugal is divided in 5 NUTS II regions: Norte, Centro, Área Metropolitana de Lisboa, Alentejo and Algarve.

We show in the Appendix that $\hat{\beta}_m$ can be written as

$$ln\left(\frac{1+g_m^{20,19}}{1+g_2^{20,19}}\cdot\sqrt[3]{\frac{(1+g_m^{19,18})^2}{(1+g_2^{19,18})^2}\cdot\frac{1+g_m^{18,17}}{1+g_2^{18,17}}}\right), \qquad (2)$$

where $g_m^{t,t-1}$ stands for the month m YoY growth rate between years t and $t-1$. In order to provide estimates of the causal impact of Covid-19 in the YoY growth rates for each month from March 2020 onward, we use (2) to correct for seasonality. More specifically, we use the empirically observed YoY growth rates between 2019 and 2018, and between 2018 and 2017, to replace for the cubic root term.[20]

We then explore the possibility that the impact of Covid-19 for the different groups in each dimension of our data (gender, age and education) is not homogeneous. We use a triple difference-in-differences strategy. The following equation is estimated for the gender dimension:

$$ln(y)_{kimt} = \alpha+\gamma_i+\delta_m+\lambda\mathbb{1}_T+\beta_0\mathbb{1}_f+\beta_1\mathbb{1}_T\mathbb{1}_f+\beta_2\mathbb{1}_{m\geq3}\mathbb{1}_f+\beta_3\mathbb{1}_{m\geq3}\mathbb{1}_T+\beta_4\mathbb{1}_{m\geq3}\mathbb{1}_T\mathbb{1}_f+\epsilon_{kimt} \quad (3)$$

Where $ln(y)_{kimt}$ is the (log of) number of unemployed people of gender $k\in\{$female, male$\}$, in municipality i, month m, and year t, $\mathbb{1}_{m\geq3}$ is an indicator for the months of March and subsequent, $\mathbb{1}_f$ is female indicator, and the remaining variables have the same meaning as in (1). We estimate a similar equation for three age categories (in which the reference category is above 55), and for the four education levels (in which the reference category is higher education). Our coefficient of interest is β_4, and it gives the differential causal impact of the pandemic crisis on females.

Lastly, on Section 5 we again use a triple difference-in-differences specification (4). This time, we interact the indicators described above with $temp_i$, the share of workers with temporary contracts in the private sector of each municipality i in 2018. We begin by estimating (4) for the whole sample, and then for sub-samples according to gender, age, and education level. [21]

---

[20]Please refer to the Appendix A for details.

[21]This data is collected every year in October, and may not account for seasonality in the regional distribution of temporary jobs. In order to mitigate this concern, we conducted a robustness test about our Dual Market results in Section 5.

$$ln(y)_{imt} = \gamma_i + \delta_m + \lambda \mathbb{1}_T + \alpha_0 \mathbb{1}_T \, temp_i + \alpha_1 \mathbb{1}_{m\geq 3} \, temp_i + \alpha_2 \mathbb{1}_{m\geq 3} \mathbb{1}_T + \alpha_3 \mathbb{1}_{m\geq 3} \mathbb{1}_T \, temp_i + \epsilon_{imt} \quad (4)$$

In this case, our coefficient of interest is α_3, and it gives us the causal impact of the Covid-19 pandemic on municipal unemployment when the share of temporary workers increases by 1 p.p. Since we control for municipality fixed effects and $temp_i$ is time invariant, we do not include it alone in the regression.

4 Main results

In this section, we present our main results. We begin by estimating the overall impact of the pandemic on unemployment, new job placements and the motives for new registrations at the job centers. We then exploit the potential heterogeneous effects for different demographic groups.

4.1 The size of the shock

We start by using (1) to estimate the impact of the pandemic crisis on registered unemployment and new job placements. The coefficient estimates for β_m are depicted in Figure 3. All our coefficient plots display the 95% confidence intervals.

The first important remark is that in both cases the estimates for β_m, where $m \in \{1, 10, 11, 12\}$, are not statistically different from zero; in other words, the parallel trends assumption is verified, which validates our identification strategy and shows that our results reflect the causal impact of Covid-19 on the variables of interest, as explained in Subsection 3.2. This is also shown in Figure 2, where we plot the municipal average of (natural log) of unemployment and new job placements across time, for different periods. In both cases, it is clear that the trends are parallel in the pre-treatment period.

Figure 3: Event study aggregate effects

(a) Unemployment

(b) New Job Placements
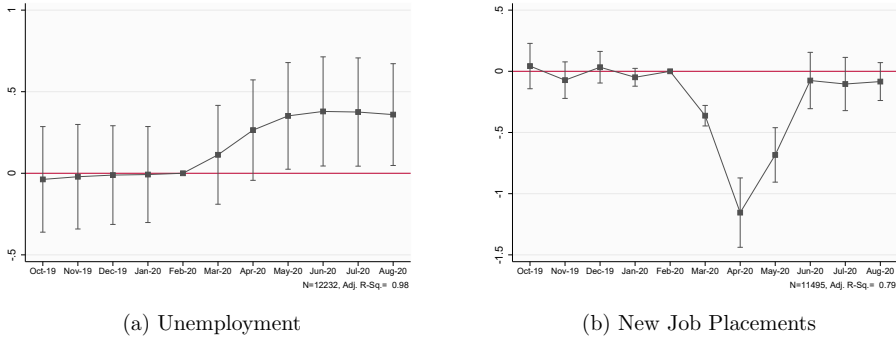
Panel (a) in Figure 3 shows a strong impact on unemployment following the lockdown period that began in March 2020. The increase is persistent but more pronounced until June, and stabilizes thereafter. In terms of job placements, Panel (b) presents a colossal drop of new placements, especially in April, followed by a recovery in May and June and a subsequent stabilization. Although between June and August the point estimates are not statistically different from zero, they are still negative.

Table 2 displays the net causal impact of the pandemic on the YoY growth rates, computed as explained in Appendix A.

Table 2: Event study aggregate effects: magnitudes

Dep.Var.:	Log of Unemployment			Log of Job Placements		
	Point Estimate	t-test	Effect (pp)	Point Estimate	t-test	Effect (pp)
	(1)	(2)	(3)	(4)	(5)	(6)
Mar-20	0.113	1.04	10.67	-0.363	-12	-23.55
Apr-20	0.265	2.39	26.92	-1.155	-11.31	-62.74
May-20	0.352	2.99	37.25	-0.683	-8.53	-43.14
Jun-20	0.379	3.15	39.06	-0.075	-0.91	-3.37
Jul-20	0.375	3.14	38.42	-0.104	-1.33	-8.48
Aug-20	0.360	3.2	35.5	-0.084	-1.5	-0.23

Notes: Point estimates are the coefficients β_m from (1). The effect is given by $(1 + g_2^{20,19})(\vartheta_m - 1)$. Please refer to Appendix A for more information.

The YoY growth rates of unemployment increased gradually over time, from 27 p.p. in April, up to 39 p.p. and 38 p.p. in June and July, respectively. The sharp decline of new job placements shown in Panel (b) on Figure 3 corresponds to a 63 p.p. drop in April. From June onward, the impact has been attenuated but is still negative. These effects are

consistent with the deep lockdown in April and the slow restart of the economic activity during the summer.

To provide additional evidence of the strength of our results, we perform robustness tests to *(i)* further assert that the parallel trend assumption holds and *(ii)* show that the remaining coefficient estimates are stable across different specifications. Appendix Figure D.2 shows the baseline results for unemployment and new job placements when we replace the municipality fixed effects by NUTS III fixed effects (in red) and NUTS III x month fixed effects (in green). This last option controls for unobserved regional seasonality not accounted for on our baseline specification. Reassuringly, the results are very similar across specifications.

We also use (1) to analyse regional differences on unemployment across the five Portuguese mainland NUTS II regions. Event studies and the causal impacts are shown in the Appendix (Figure D.3 and Table C.1, respectively). The Southern region of Algarve was by far the most hit by the pandemic, with YoY growth rate increases of 166 p.p., 187 p.p. and 180 p.p. in May, June and July, respectively. This effect is likely a consequence of the Algarve region being highly dependent on tourism and hospitality services, which suffered a severe downturn due to the restrictions imposed in the country. The second most affected regions were Lisboa e Vale do Tejo and Alentejo. By August, all the regions remained far from recovery.

4.2   Labour market transitions in a pandemic

We now turn to the investigation of the reasons that led the individuals to register with the Public Employment Service. Our data splits the motives into four categories that are explicit job separations (dismissal, voluntary quit, mutual agreement dismissal, end of temporary contract) and two additional categories for transition from inactivity and registration from self-employed.

Figure 4 shows the event study for each motive. Since these are flow variables, the results measure the impact on new unemployment each month, net of composition effects. The impact for self-employed individuals is positive, but not significant. We now exploit

the remaining ones.

Figure 4: Motives to register with IEFP

(a) Dismissed from previous job

(b) Voluntarily quit previous job

(c) Mutual agreement dismissal

(d) End of temporary contract

(e) Former inactive worker

(f) Self-employed

Panels (a) and (d) exhibit very clear spikes in dismissals and terminations of temporary contracts, respectively, lasting until the end of the period of analysis. The effects were particularly strong in April, with YoY growth rate increases of 216 p.p. and 87 p.p. for dismissals and end of temporary contracts, respectively.[22] After April, the impact was less pronounced, but always above the baseline levels before the pandemic. Although the April

---

[22] The causal impacts of Covid-19 on all variables are shown in Appendix Table C.2.

spike is higher for dismissals, it is important to highlight that the bulk of job separations is due to the termination of temporary contracts. As shown in Figure D.4, in appendix, they represent in previous years, on average, around 58% of total new registrations between March and August, and 64% in 2020.

Dismissals may include both temporary and permanent contracts, provided the former are terminated before the end of the contract.[23] Despite the fact that we cannot disentangle permanent and temporary workers in this category, given the evidence from the Labour Force Survey presented in Section 5, it is unlikely the former were the most penalized. Indeed, the number of temporary contracts decreased sharply in the first and second quarters of 2020, compared to 2019, while there were no noticeable changes on permanent contracts. At the same time, as discussed earlier, since the furlough system in place in Portugal prohibited dismissals, any effect from panel (a) must have come from firms that do not benefit from the program.

The transition from inactivity, panel (e), was also severely affected. By April, the YoY growth rate was down by 53 p.p., implying that individuals refrained from actively seeking employment during the stricter lockdown period. We also find an impact on dismissals by mutual agreement in April, and a small drop of voluntary quits in May and June, in panels (c) and (b), respectively.

Taken together, the results show that the labour market impact was driven both by demand and supply side mechanisms. On the demand side, firms responded by letting go of employees to reduce costs, particularly workers with temporary contracts. On the supply side, workers responded to the lack of job prospects induced by the crisis by refraining from quitting jobs voluntarily or by mutual agreement, and by decreasing the transitions from inactivity. The lower transition from inactivity implies that our results on unemployment are a lower bound of the effect of the pandemic on the labour market, since many individuals remain hidden in the inactive population.

---

[23]Dismissals generally require compensation, which is higher for permanent than for temporary contracts.

4.3 Gender, Age and Education

In this section, we use a triple difference-in-differences strategy to explore the heterogeneous effects of the Covid-19 shock on the outcomes of interest for different *(i)* gender, *(ii)* age groups, and *(iii)* education levels. On Tables 3, 4 and 5, we report the estimates of β_3 and β_4 from (3). In all specifications, the estimates of β_3 represent the impact for the reference group.

Table 3 shows the results for gender, using male as the omitted group. Column (1) presents the impact on registered unemployment while column (2) presents the impact on new job placements. The results indicate that Covid-19 increased male unemployment by 33.8% and decreased new job placements by 24.1% between March and August 2020. In terms of unemployment, there is no statistically significant difference between men and women. Taken together, our results show that women were most severely hit by the pandemic: while there is no statistically significant difference in unemployment, women suffer an additional drop of 17.5% in placements after March, when compared to men.

Table 3: Triple DD on unemployment and new job placements, by gender

Dep. Var.:	Log of Unemployment	Log of New Job Placements
	(1)	(2)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T$	0.338**	-0.241*
	(0.08)	(0.11)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T \times \mathbb{1}_{\text{female}}$	-0.026	-0.175*
	(0.02)	(0.08)
Number of Obs.	24,464	21,265
R-squared	0.968	0.725

Notes: Standard errors (in parenthesis) are clustered at NUTS II and time period (month, year) level.
Significance levels: * p < 0.10, ** p < 0.05, *** p < 0.01

The absence of gender differences in unemployment is in line with the findings of Casarico and Lattanzio (2020) for Italy, and Hupkau and Petrongolo (2020) for the UK. The negative effect that we identify on placements suggests that women are less likely to find a job following an unemployment episode during the pandemic; this adds to the long list of differential gender impacts due to a higher proportion of female workers in the most affected industries (ILO, 2017), disproportionate take up of household chores

and childcare after school closures and work from home restrictions (Farré et al., 2020; Del Boca et al., 2020).

As mentioned in Subsection 3.1, data on new job placements is only disaggregated by gender. Hence, for the remainder of our analysis we will focus exclusively on unemployment.

To study the impact on different age groups, we use the unemployed aged more than 55 as the reference group. Our findings in Table 4 show this is the least affected group. There is a very strong impact on youth unemployment after March 2020, amounting to an additional increase of 20.8% and 25.8%, for individuals younger than 25, and between 25 and 34 years old, respectively.

Table 4: Triple DD on unemployment, by age

Dep. Var.:	Log of Unemployment
	(1)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T$	0.177**
	(0.05)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T \times \mathbb{1}_{\text{less than 25}}$	0.208**
	(0.06)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T \times \mathbb{1}_{25-34}$	0.258***
	(0.04)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T \times \mathbb{1}_{35-54}$	0.179***
	(0.03)
Number of Obs.	48,928
R-squared	0.953

Notes: Standard errors (in parenthesis) are clustered at NUTS II and time period (month, year) level.
Significance levels: * p < 0.10, ** p < 0.05, *** p < 0.01

These results are consistent with the prevalence of non-permanent contracts among younger workers in Portugal that we document on Section 5. Younger workers are also less unionized than older ones (Portugal and Vilares, 2013) and more vulnerable to precarious working conditions.

Table 5 shows the heterogeneity in terms of education levels. The reference group is individuals with higher education, which experienced an increase of 23.9% in unemployment after March 2020. The differential impact of the crisis was 17.5% higher for individuals with upper secondary education, 15% for individuals with lower secondary

education, and 9.6% for individuals with basic education, than for those with higher education. Individuals with less than four years of formal education do not behave differently than the highly educated ones. This may stem from the fact that these workers are concentrated in essential jobs that kept working during the lockdown; however, the lack of statistical significance may also be explained by the relatively small number of workers with this level of education.

Table 5: Triple DD on unemployment, by education level

Dep. Var.:	Log of Unemployment
	(1)
$\mathbb{1}_{m\geq3} \times \mathbb{1}_T$	0.239**
	(0.07)
$\mathbb{1}_{m\geq3} \times \mathbb{1}_T \times \mathbb{1}_{\text{Primary or less}}$	-0.037
	(0.03)
$\mathbb{1}_{m\geq3} \times \mathbb{1}_T \times \mathbb{1}_{\text{Basic}}$	0.096**
	(0.02)
$\mathbb{1}_{m\geq3} \times \mathbb{1}_T \times \mathbb{1}_{\text{Lower Secondary}}$	0.150***
	(0.03)
$\mathbb{1}_{m\geq3} \times \mathbb{1}_T \times \mathbb{1}_{\text{Upper Secondary}}$	0.175***
	(0.02)
Number of Obs.	61,156
R-squared	0.930

Notes: Standard errors (in parenthesis) are clustered at NUTS II and time period (month, year) level.
Significance levels: * p < 0.10, ** p < 0.05, *** p < 0.01

The inverted u-shaped relationship between education levels and the labour market impact of the pandemic can be explained by the fact that workers with secondary education are usually not employed in the the subset of service sectors compatible with home working, and are thus in a more vulnerable position in this crisis. Indeed, data from Statistics Portugal shows that during the second quarter of 2020, 4.7% of the employed population with lower secondary education or less was working from home, compared with 53.8% of the population with higher education degrees (INE, 2020).

5 Dual labour market and the Covid-19 crisis

We already documented an increase in registrations at the Public Employment Services and a sharp decrease in job placements (Subsection 4.1), and we showed that this is mostly

due to non-voluntary dismissals and temporary contract terminations (Subsection 4.2). Then, we analysed the causal impact of the shock on different groups of the population, and estimate a sizeable causal impact of the pandemic crisis on younger workers, on the one hand, and middle educated ones, on the other hand (Subsection 4.3).

We now explore whether the duality of the labour market explains the transitions into unemployment in this period. We do so by exploiting the differences in local labour markets. More specifically, we analyse the possibility that municipalities with a higher share of temporary contracts are more impacted by the crisis, and study whether this effect is stronger for some groups of workers. We use the municipal share of temporary contracts in 2018, the last year available. As we show in Figure D.5, in appendix, this share is strongly correlated with that of previous years, which suggests that we capture a structural feature of the local labor markets.

In Figure 1 we established that Portugal has one of the highest shares of temporary employment of the OECD. We now analyse this further by type of worker. This is done in panels (a) to (c) of Figure 5, that display the share of temporary employment as a percentage of total employment between the first quarter of 2018 and the second quarter of 2020, for each group of workers split by gender, age, and education level. In addition, panel (d) shows the year-on-year change of the number of permanent and temporary workers between the first quarter of 2018 and the second quarter of 2020.

According to panel (a), the share of temporary contracts decreased in 2020, without any consistent difference between female and male workers. The age differences, shown in panel (b), are the most striking. The share of workers aged less than 25 years old with temporary contracts is around four times that of those who are older than 35. In the 25-34 age interval, the prevalence of temporary work is twice as much as that of the older individuals. Finally, from panel (c) it is clear that temporary employment is more prevalent among individuals with upper secondary education. The fact that individuals with basic, primary education, or less, represent the group with the lowest share of temporary contracts is mostly driven by age. These education levels are more common in the oldest cohorts, who also have a more stable relationship with the labour

market.

Figure 5: Temporary employment in Portugal

(a) by gender

(b) by age

(c) by education level

(d) YoY change
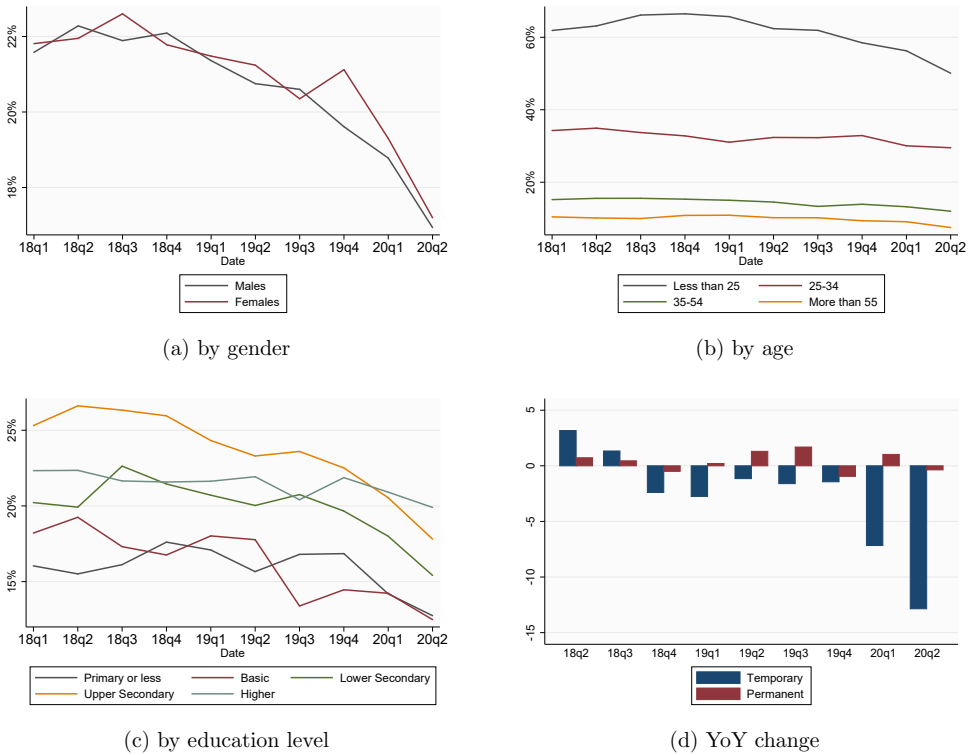
Source: Labour Force Survey (Statistics Portugal)

The three panels display a sharp decrease in the share of temporary jobs for both males and females, for individuals younger than 25, and those with upper or lower secondary education, as of the first quarter of 2020. This drop may be explained by a conversion of temporary contracts into permanent ones or by an increase in job separations that hits temporary contracts and spares the permanent ones. Panel (d) shows that the YoY quarterly change in employment in 2020 was positive in the first quarter and only marginally negative in the second, for permanent workers, in sharp contrast with the strongly negative for temporary ones. Therefore, the explanation is clearly the latter.

The combined evidence in Subsection 4.3 and Figure 5 strongly suggests that the young and middle educated workers who are more hit by the crisis are also the temporary

ones. We assess this hypothesis more formally by testing if the increase in unemployment was larger for young and middle-educated workers in municipalities with a higher share of temporary contracts.

Table 6 shows the estimate of α_3 in equation 4 for all the unemployed. The estimate indicates that there are more registrations at the Public Employment Services in municipalities with a higher share of temporary employment. The number of newly registered unemployed increases by 1.5% with an increase of 1pp in the share of temporary workers in a municipality. Alternatively, a one standard deviation (8%) increase in the share of temporary contracts amounts to a 11.6% change in the number of registries. This effect is sizeable, particularly given that the average share of temporary workers is 33%.

Table 6: Share of temporary contracts and the Covid-19 crisis

Dep. Var.:	Log of Unemployment
	(1)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T \times temp_i$	1.452*
	(0.57)
Number of Obs.	12,232
R-squared	0.976

Notes: Standard errors (in parenthesis) are clustered at NUTS II and time period (month, year) level.

Significance levels: * p < 0.10, ** p < 0.05, *** p < 0.01

We now present the effects across types of workers. The estimates of α_3 from (4) for each group of workers are presented in Table 7. We report the tests for the null hypothesis that the estimated coefficients are equal in Tables C.3 and C.4 in the Appendix. As regards gender, the value of the χ^2 statistic is 4.03, i.e., the coefficients are statistically different at 5%. More specifically, a one standard deviation increase in the share of temporary contracts in a dual labour market increases the number of newly registered female workers in response to the pandemic shock by 12.6%, and that of males by 10.7%.

We now analyse the results with respect to age. Interestingly, the estimated coefficients are monotonic in the age of the individuals. The tests of the difference of the coefficients in Table C.3 indicate that there is no statistical difference between the two groups who are younger than 34. Conversely, the coefficients for both these groups are different from the ones of older workers. The negative impact of the pandemic shock in dual labour markets

is therefore more concentrated in younger workers. More specifically, a one standard deviation increase in the share of temporary workers leads to an increase of between 14.2% and 16.5% in the number of newly registered young workers in the Employment Office. Conversely, the impact on older workers amounts to between 7.4% and 11.4%.

Regarding education, the effect follows an inverted U-shaped pattern. In municipalities with a higher share of temporary contracts, Covid-19 impacts less severely individuals with primary education or less, and individuals with higher education. The strongest impact falls on the those with lower and upper secondary education, as the results of the tests of the equality of the coefficients reported in Table C.4 confirm. A one-standard deviation increase in the share of temporary workers in the municipality increases the number of registered unemployed people with upper and lower secondary education by 12.6% and 13.7%, respectively. The impact on highly educated workers is non significant.

Table 7: Share of temporary contracts and the Covid-19 crisis

Dep.Var.:	Log of Unemployment					
Dimension:	**Gender**		**Age**			
	Male	Female	< 25	25-34	35-54	> 55
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T \times temp_i$	1.331**	1.569*	2.065*	1.771*	1.426*	0.919*
	(0.47)	(0.66)	(0.35)	(0.36)	(0.29)	(0.19)
Number of Obs.	12,232	12,232	12,232	12,232	12,232	12,232
R-squared	0.971	0.976	0.940	0.961	0.971	0.983

Dimension:	**Education**				
	Primary or less	Basic	Lower Sec.	Upper Sec.	Higher
	(7)	(8)	(9)	(10)	(11)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T \times temp_i$	1.156*	1.338*	1.712*	1.572*	1.219
	(0.47)	(0.58)	(0.62)	(0.60)	(0.58)
Number of Obs.	12,232	12,232	12,232	12,232	12,232
R-squared	0.969	0.961	0.963	0.970	0.970

Notes: Standard errors (in parenthesis) are clustered at NUTS II and time period (month, year) level.

Significance levels: * p < 0.10, ** p < 0.05, *** p < 0.01

These results strongly suggest that the effect of the crisis in female, young and middle-educated workers is driven by the duality of the labour market. The crisis has asymmetric effects depending on the workers' ties to the labour market, and measures like the furlough scheme do not seem to be enough to protect certain groups of the labour force.

As we discussed in Subsection 3.1, the share of temporary workers is collected, every year, in October. We address these concern in Table C.6, in appendix, which presents our results excluding the municipalities in the top quartile of the distribution of tourist overnight stays, which are bound to have a peak of jobs in the Summer. The results are similar to baseline, with the exception of the gender differences, which are less precisely estimated.

6   Conclusion

In the beginning of 2020, the Coronavirus pandemic hit the world economy, and rapidly turned into the biggest shock since the Second World War. Labour markets were particularly hit by the crisis, given the massive disruption in supply chains and production processes, that brought many sectors of activity to an almost complete halt.

In this paper, we analyse the effects of the crisis on unemployment, using administrative data from *Instituto do Emprego e Formação Profissional*, covering the universe of unemployed individuals registered at job centers from October 2016 to August 2020. Using event study difference-in-differences, we rely on the assumption that, in the absence of the Covid-19 outbreak, the monthly year-on-year change between March/August 2020 and March/August 2019 would have been parallel to a weighted geometric mean of the year-on-year change of the previous 3 years.

We document a large causal impact of the pandemic on registered unemployment, with YoY growth rate increases from 27 percentage points in April up to 39 and 38 percentage points in June and July, respectively. New job placements were also severely affected, i.e., the YoY growth rates were below pre-crisis levels from March to August, with a negative peak of 63 percentage points in April.

We show that the transitions into unemployment were driven by dismissals and termination of temporary contracts. While the April spike is higher for dismissals, the bulk of job separations is due to the termination of temporary contracts. We also document a strong decrease in transitions from inactivity into unemployment, meaning that inactive individuals were discouraged from seeking a job, which suggests that the impact on

unemployment is a lower bound of the true effect of the pandemic on the labour market.

We then decompose the impact across demographic groups. We perform a triple difference-in-differences analysis and show that the impact on unemployment is more pronounced for individuals who are less than 25 years old (additional 20.8%) and between 25 and 34 years old (additional 25.8%) than for those with less than 55 years old. In terms of education, the bulk of the effect is concentrated in individuals with lower and upper secondary education (additional 15% and 17.5%, respectively *vis-à-vis* those with higher education). We find no evidence of gender differences on registered unemployed, but women are more affected in terms of new job placements, with an additional decline of 17.5%, when compared to men.

Finally, we show that the disproportionate impact of the crisis on females, younger, and middle educated workers is explained by dual labour markets. Using the Labour Force Survey, we document a sizeable decrease in the number of temporary contracts for the young and the middle educated workers, in the first and second quarters of 2020. We then show that the impact of Covid-19 on total unemployment is higher in municipalities with a higher share of temporary workers. Duality also amplifies the asymmetry of the impact: a one standard deviation increase in the share of temporary workers in a municipality leads to an increase of 12.6% for females, between 14.2% and 16.5% in the registrations of workers who are younger than 25, and an increase between 12.6% and 13.7%, for those with secondary education.

These results likely reflect the lower cost to dismiss temporary workers compared to permanent ones. While this is the case in normal times, the policies implemented with the aim of mitigating the impact of the pandemic exacerbated the different protection levels of the two types of workers, as they explicitly forbid firms that receive support from terminating labour contracts. This created an additional layer of job protection, that increased the dual character of the market during the pandemic. This is particularly problematic during a crisis where job postings are severely constrained. The dual job protection should have been compensated by generous income support policies targeting the least protected part of the market.

Furlough policies maintain the matching capital between firms and workers in the short run, but longer periods of support can be problematic because they lock-in production factors in zombie firms. Our results suggest an unanticipated effect of the prohibition of dismissals linked to Covid-19 support policies, applied in several countries, including the United Kingdom, Spain, and Italy. These measures exacerbate the labour market duality and put temporary workers and service providers in a vulnerable position in the face of a large crisis, particularly given the few job offers during these periods. Moreover, whether the matching capital of permanent workers is more valuable than that of permanent ones is an open policy question. The optimal design of policies protecting matching capital in dual labour markets is an open avenue of research.

References

Adams-Prassl, Abigail, Teodora Boneva, Marta Golin, and Christopher Rauh. 2020. "Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys." *Journal of Public Economics* 189:104245. https://doi.org/10.1016/j.jpubeco.2020.104245.

Aleksynska, Mariya. 2018. "Temporary employment, work quality, and job satisfaction." *Journal of Comparative Economics* 46 (3): 722–735. https://doi.org/10.1016/j.jce.2018.07.004.

Alon, Titan, Matthias Doepke, Jane Olmstead-Rumsey, and Michèle Tertilt. 2020. "This Time It's Different: The Role of Women's Employment in a Pandemic Recession." *National Bureau of Economic Research,* Working Paper Series, no. 27660, https://doi.org/10.3386/w27660.

Alstadsæter, Annette, Bernt Bratsberg, Gaute Eielsen, Wojciech Kopczuk, Simen Markussen, Oddbjorn Raaum, and Knut Røed. 2020. "The First Weeks of the Coronavirus Crisis: Who Got Hit, When and Why? Evidence from Norway." *National Bureau of Economic Research,* Working Paper Series, no. 27131, https://doi.org/10.3386/w27131.

Aum, Sangmin, Sang Yoon (Tim) Lee, and Yongseok Shin. 2020. "COVID-19 Doesn't Need Lockdowns to Destroy Jobs: The Effect of Local Outbreaks in Korea." *Centre for Economy Policy Research,* CEPR Discussion Papers, no. 14822.

Bamieh, Omar, and Lennart Ziegler. 2020. "How Does the Covid-19 Crisis Affect Labor Demand? An Analysis Using Job Board Data from Austria." *Institute of Labor Economics,* IZA Discussion Papers, no. 13801.

Banco de Portugal. 2020. "Boletim Económico - Outubro 2020."

Betcherman, Gordon, Nicholas Giannakopoulos, Ioannis Laliotis, Ioanna Pantelaiou, Mauro Testaverde, and Giannis Tzimas. 2020. "Reacting Quickly and Protecting Jobs: The Short-Term Impacts of the COVID-19 Lockdown on the Greek Labor Market." *Institute of Labor Economics,* IZA Discussion Papers, no. 13516.

Blanchard, Olivier, and Pedro Portugal. 2017. "Boom, Slump, Sudden stops, Recovery, and Policy Options. Portugal and the Euro." *Banco de Portugal,* Working Papers 2017, no. 11.

Boeri, Tito, and Herbert Brücker. 2011. "Short-Time Work Benefits Revisited: Some Lessons from the Great Recession." *Economic Policy* 26:697–765. https://doi.org/10.1111/j.1468-0327.2011.271.x.

Cajner, Tomaz, Leland D. Crane, Ryan A. Decker, John Grigsby, Adrian Hamins-Puertolas, Erik Hurst, Christopher Kurz, and Ahu Yildirmaz. 2020. "The U.S. Labor Market during the Beginning of the Pandemic Recession." *National Bureau of Economic Research,* Working Paper Series, no. 27159, https://doi.org/10.3386/w27159.

Carneiro, Anabela, Pedro Portugal, and José Varejão. 2014. "Catastrophic Job Destruction during the Portuguese Economic Crisis." *Journal of Macroeconomics* 39:444–457. https://doi.org/10.1016/j.jmacro.2013.09.018.

Carvalho, Bruno, Susana Peralta, and Mariana Esteves. 2021. "Portugal, Balanço Social 2020." *Nova School of Business and Economics,* https://doi.org/10.34619/LLCR-OTR2.

Carvalho, Bruno, Susana Peralta, and João Pereira dos Santos. 2020a. "Regional and Sectorial Impacts of the Covid-19 Crisis: Evidence from Electronic Payments." *ULB – Universite Libre de Bruxelles,* Working Papers ECARES, nos. 2020-48.

———. 2020b. "What and how did people buy during the Great Lockdown? Evidence from electronic payments." *Centre for Economic Policy Research,* Covid Economics: Vetted and Real-Time Papers. Issue 28:119-158.

Casarico, Alessandra, and Salvatore Lattanzio. 2020. "The heterogeneous effects of COVID-19 on labor market flows: Evidence from administrative data." *Centre for Economic Policy Research,* Covid Economics: Vetted and Real-Time Papers. Issue 52:152-174.

Cho, Seung Jin, Jun Yeong Lee, and John V. Winters. 2020. "Employment Impacts of the COVID-19 Pandemic across Metropolitan Status and Size." *Institute of Labor Economics,* IZA Discussion Papers, no. 13468.

Cho, Seung Jin, and John V. Winters. 2020. "The Distributional Impacts of Early Employment Losses from COVID-19." *Global Labor Organization,* GLO Discussion Paper Series, no. 554.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2020. "Labor Markets During the COVID-19 Crisis: A Preliminary View." *National Bureau of Economic Research,* Working Paper Series, no. 27017, https://doi.org/10.3386/w27017.

Cortes, Guido Matias, and Eliza C. Forsythe. 2020. "The Heterogeneous Labor Market Impacts of the Covid-19 Pandemic." *W.E. Upjohn Institute for Employment Research,* Upjohn Working Papers and Journal Articles, nos. 20-327, https://doi.org/10.17848/wp20-327.

Del Boca, Daniela, Noemi Oggero, Paola Profeta, and Maria Cristina Rossis. 2020. "Women's Work, Housework and Childcare, before and during COVID-19." *Institute of Labor Economics,* IZA Discussion Papers, no. 13409.

Dias, Mónica Costa, Robert Joyce, Fabien Postel-Vinay, and Xiaowei Xu. 2020. "The Challenges for Labour Market Policy during the COVID-19 Pandemic." *Fiscal Studies* 41. https://doi.org/10.1111/1475-5890.12233.

European Commission. 2018. "European Semester: Country Specific Recommendation/ Commission Recommendation - Portugal," https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52018DC0421.

Farré, Lídia, Yarine Fawaz, Libertad Gonzalez, and Jennifer Graves. 2020. "How the COVID-19 Lockdown Affected Gender Inequality in Paid and Unpaid Work in Spain." *Institute of Labor Economics,* IZA Discussion Papers, no. 13434.

Fauser, Sophia. 2020. "Career trajectories and cumulative wages: The case of temporary employment." *Research in Social Stratification and Mobility* 69:100529. https://doi.org/10.1016/j.rssm.2020.100529.

Ferreira, Priscila, João Cerejeira, and Miguel Portela. 2020. "IZA COVID-19 Crisis Response Monitoring - Portugal (November 2020)." *Institute of Labor Economics.*

Forsythe, Eliza, Lisa B. Kahn, Fabian Lange, and David G. Wiczer. 2020. "Labor Demand in the time of COVID-19: Evidence from vacancy postings and UI claims." *Journal of Public Economics* 189:104238. https://doi.org/10.1016/j.jpubeco.2020.104238.

Hensvik, Lena, Thomas Le Barbanchon, and Roland Rathelot. 2020. "Job Search during the COVID-19 Crisis." *Journal of Public Economics,* 104349. https://doi.org/10.1016/j.jpubeco.2020.104349.

Hupkau, Claudia, and Barbara Petrongolo. 2020. "Work, care and gender during the Covid-19 crisis." *Institute of Labor Economics,* IZA Discussion Papers, no. 13762.

ILO. 2017. "The Covid-19 Response: Getting Gender Equality Right for a Better Future for Women at Work." *International Labour Organization,* ILO brief.

INE. 2020. "Trabalho a partir de casa – Módulo ad hoc do Inquérito ao Emprego (2.º trimestre de 2020)."

Jones, Stephen R.G., Fabian Lange, W. Craig Riddell, and Casey Warman. 2020. "Waiting for Recovery: The Canadian Labour Market in June 2020." *Institute of Labor Economics,* IZA Discussion Papers, no. 13466.

Kalenkoski, Charlene Marie, and Pabilonia Sabrina Wulff. 2020. "Initial Impact of the COVID-19 Pandemic on the Employment and Hours of Self- Employed Coupled and Single Workers by Gender and Parental Status." *Institute of Labor Economics,* IZA Discussion Papers, no. 13443.

Kikuchi, Shinnosuke, Sagiri Kitao, and Minamo Mikoshiba. 2021. "Who suffers from the COVID-19 shocks? Labor market heterogeneity and welfare consequences in Japan." *Journal of the Japanese and International Economies* 59:101117. ISSN: 0889-1583. https://doi.org/10.1016/j.jjie.2020.101117.

Lemieux, Thomas, Kevin Milligan, Tammy Schirle, and Mikal Skuterud. 2020. "Initial impacts of the COVID-19 pandemic on the Canadian labour market." *Canadian Labour Economics Forum, University of Waterloo,* CLEF Working Paper Series, no. 26, https://doi.org/10.3138/cpp.2020-049.

Martins, Fernando, and Pedro Portugal. 2019. "De que forma a rigidez à baixa dos salários afetou o desemprego durante a crise económica?" In *O Crescimento Económico Português: uma visão sobre questões estruturais, bloqueios e reformas,* 97–106. Banco de Portugal.

Mayhew, Ken, and Paul Anand. 2020. "COVID-19 and the UK labour market." *Oxford Review of Economic Policy* 36 (Supplement$_1$): S215–S224. https://doi.org/10.1093/oxrep/graa017.

Meekes, Jordy, Wolter Hassink, and Guyonne Kalb. 2020. "Essential Work and Emergency Childcare: Identifying Gender Differences in COVID-19 Effects on Labour Demand and Supply." *Institute of Labor Economics,* IZA Discussion Papers, no. 13843.

Midões, Catarina. 2020. "Social distancing: did individuals act before governments?" (April 7, 2020). Accessed December 29, 2020. https://www.bruegel.org/2020/04/social-distancing-did-individuals-act-before-governments/.

Montenovo, Laura, Xuan Jiang, Felipe Lozano Rojas, Ian M. Schmutte, Kosali I. Simon, Bruce A. Weinberg, and Coady Wing. 2020. "Determinants of Disparities in Covid-19 Job Losses." *National Bureau of Economic Research,* Working Paper Series, no. 27132, https://doi.org/10.3386/w27132.

Nunez, Imanol, and Ilias Livanos. 2014. "Temps "by choice"? An Investigation of the Reasons Behind Temporary Employment Among Young Workers in Europe." *Journal of Labor Research* 36:44–66. https://doi.org/10.1007/s12122-014-9195-3.

OECD. 2017. "Labour Market Reforms in Portugal 2011-15: A Preliminary Assessment." *OECD Publishing,* https://doi.org/10.1787/9789264269576-en.

———. 2020. "Job retention schemes during the COVID-19 lockdown and beyond." *OECD Policy Responses to Coronavirus (COVID-19).*

Peralta, Susana, Bruno P. Carvalho, and Mariana Esteves. 2021. "Portugal, Balanço Social 2020: Um retrato do país e dos efeitos da pandemia." *Nova School of Business and Economics,* Social Equity Initiative.

Portugal, Pedro, Pedro Raposo, and Hugo Reis. 2018. "The distribution of wages and wage inequality." *Bank of Portugal Economic Studies* (4(1):1-14).

Portugal, Pedro, and Hugo Vilares. 2013. "Labor unions, union density and the union wage premium." *Banco de Portugal Economic Bulletin,* 61–71.

Virtanen, Marianna, Mika Kivimäki, Matti Joensuu, Pekka Virtanen, Marko Elovainio, and Jussi Vahtera. 2005. "Temporary employment and health: a review." *International Journal of Epidemiology* 34 (3): 610–622. https://doi.org/10.1093/ije/dyi024.

Weel, Bas. 2018. "The Rise of Temporary Work in Europe." *De Economist* 166. https://doi.org/10.1007/s10645-018-9329-8.

A Derivation of the causal impacts

Departing from Equation 1, we can write the coefficients $\hat{\beta}_m$ as an estimate of a function of growth rates:

$$ln\left(\sqrt[3]{\frac{1+g_m^{20,19}}{1+g_2^{20,19}}\frac{1+g_m^{20,18}}{1+g_2^{20,18}}\frac{1+g_m^{20,17}}{1+g_2^{20,17}}}\right)$$

where $g_m^{20,19}$, $g_m^{20,18}$ and $g_m^{20,17}$ represent the YoY growth rate of the outcome variable in month m, where $m \in \{3,...,8\}$, from 2019, 2018 and 2017 to 2020, respectively. This expression can be further simplified to:

$$ln\left(\frac{1+g_m^{20,19}}{1+g_2^{20,19}}\sqrt[3]{\frac{\left(1+g_m^{19,18}\right)^2}{\left(1+g_2^{19,18}\right)^2}\frac{1+g_m^{18,17}}{1+g_2^{18,17}}}\right) \tag{5}$$

with $g_m^{19,18}$ representing the YoY growth rate of month m from 2018 to 2019 and $g_m^{18,17}$ representing the YoY growth rate from 2017 to 2018. As such, to estimate the causal impact of the pandemic crisis on the gross YoY growth rates $\frac{1+g_m^{20,19}}{1+g_2^{20,19}}$, we can compute:

$$\vartheta_m = Exp(\hat{\beta}_m) \times \sqrt[3]{\frac{\left(1+g_2^{19,18}\right)^2}{\left(1+g_m^{19,18}\right)^2}\frac{1+g_2^{18,17}}{1+g_m^{18,17}}}$$

Hence, we use the growth rates observed in the data to correct for any possible seasonal differences between the YoY growth rates of each month m and February. Finally, we estimate the net YoY growth rates by computing $(1+g_2^{20,19})(\vartheta_m - 1)$, with which we obtain the net impact of the crisis on the outcome variables in percentage points.

B   Overview of the literature

Table B.8: The effects of the pandemic on the labor market: summary

Authors	Country	Dependent Variable	Results
Alstadsæter et al. (2020)	Norway	Individual unemployment benefits applications	During the first few weeks after the initial government measures, approximately 12% of the labour force signed up for unemployment benefits.
Cajner et al. (2020)	US	Number of paychecks issued in the period under analysis by the company ADP	US aggregate employment fell by 21% through late-April. As of late June, employment was still 13% below February levels.
Adams-Prassl et al. (2020)	US, UK, Germany	Percentage of survey respondents reporting to have lost their jobs within the last four weeks	In early April, 18%, 15% and 5% of respondents report having lost their jobs due to the coronavirus outbreak in the US, the UK and Germany, respectively.
Lemieux et al. (2020)	Canada	Employment using data from the Canadian labour Force Survey	The impact of Covid-19 represents a 15% decline in employment.
Baek et al. (2020)	US	Cumulative weekly unemployment insurance claims by state normalized by total employment for each state	An additional week of exposure to stay-at-home policies increased unemployment insurance claims by approximately 1.9% of a state's employment level.
Cerqua and Letta (2020)	Italy	Log of overall employment using administrative data of the universe of Italian private non-financial sector firms	By the end of the third quarter of 2020, the pandemic had entailed a 1.86% decrease in overall employment in Italy, compared to what employment levels would have been had the pandemic never reached the country.
Ramos (2020)	Spain	Registered unemployment using administrative data from Public Employment Services records	Registered unemployment increased by 21.1% in April, 25.3% in May and 28.1% in June compared to the same month of the previous year.

C   Additional Tables

Table C.1: NUTS II: magnitudes

Dep.Var.:	Log of Unemployment									
	Norte		**Centro**		**Lisboa VT**		**Alentejo**		**Algarve**	
	P.E.	Eff. (pp)	P. E.	Eff. (pp)	P.E.	Eff. (pp)	P.E.	Eff. (pp)	P.E.	Eff. (pp)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mar-20	0.079	6.77	0.099	9.24	0.126	11.88	0.126	14.18	0.284	33.22
	(0.75)		(0.9)		(0.98)		(0.99)		(6.9)	
Apr-20	0.210	19.15	0.206	19.39	0.323	35.03	0.257	27.88	0.683	101.17
	(1.99)		(1.93)		(2.51)		(1.99)		(9.46)	
May-20	0.284	27.18	0.274	27.25	0.413	46.09	0.330	35.57	0.963	166.11
	(2.59)		(2.48)		(3.16)		(2.5)		(11.02)	
Jun-20	0.300	27.03	0.287	27.01	0.442	49.01	0.371	39.22	1.042	186.71
	(2.7)		(2.56)		(3.43)		(2.85)		(10.97)	
Jul-20	0.305	28.04	0.292	28.64	0.462	50.06	0.325	31.27	0.994	179.54
	(2.74)		(2.59)		(3.66)		(2.57)		(11.26)	
Aug-20	0.296	26.71	0.280	25.62	0.448	46.50	0.322	30.92	0.889	146.78
	(2.78)		(2.64)		(3.71)		(2.72)		(12.12)	

Notes: t-statistics in parenthesis. Point estimates are the coefficients β_m from (1). The effect is given by $(1 + g_2^{20,19})(\vartheta_m - 1)$. Please refer to Appendix A for more information.

Table C.2: Motives for registration: magnitudes

Dep.Var.:	Log of New Unemployment					
	Dismissed from previous job		Voluntarily quit previous job		Mutual Agreement Dismissal	
	P.E.	Eff. (pp)	P. E.	Eff. (pp)	P.E.	Eff. (pp)
	(1)	(2)	(3)	(4)	(5)	(6)
Mar-20	0.654	112.52	-0.126	-6.10	-0.004	3.99
	(9.45)		(-3.85)		(-0.07)	
Apr-20	1.145	215.70	-0.031	-6.37	0.179	26.08
	(22.62)		(-0.57)		(4.98)	
May-20	0.513	72.80	-0.324	-23.31	-0.062	-10.59
	(9.4)		(-8.94)		(-1.49)	
Jun-20	0.322	55.72	-0.147	0.53	-0.110	-7.79
	(4.49)		(-3.57)		(-2.62)	
Jul-20	0.301	33.27	-0.125	-14.68	-0.080	-8.96
	(8.68)		(-1.71)		(-1.08)	
Aug-20	0.296	40.17	-0.148	-6.25	-0.002	-2.55
	(5.24)		(-2.66)		(-0.03)	
	End of temporary job		Former inactive worker		Self-employed	
	P.E.	Eff. (pp)	P. E.	Eff. (pp)	P.E.	Eff. (pp)
	(7)	(8)	(9)	(10)	(11)	(12)
Mar-20	0.424	69.58	-0.424	-25.97	0.078	24.58
	(8.73)		(-3.41)		(0.59)	
Apr-20	0.675	87.26	-0.932	-53.48	0.392	50.53
	(13.14)		(-13.89)		(1.94)	
May-20	0.455	58.00	-0.462	-29.07	0.148	18.40
	(12.34)		(-5.22)		(0.91)	
Jun-20	0.274	36.88	-0.136	-2.25	0.199	37.38
	(5.03)		(-1.52)		(1.79)	
Jul-20	0.220	19.64	-0.131	-16.53	0.067	11.82
	(3.37)		(-1.96)		(0.42)	
Aug-20	0.186	22.94	0.065	13.40	0.165	27.36
	(3.19)		(0.65)		(1.17)	

Notes: t-statistics in parenthesis. Point estimates are the coefficients β_m from (1). The effect is given by $(1 + g_2^{20,19})(\vartheta_m - 1)$. Please refer to Appendix A for more information.

Table C.3: Means tests of coefficient equality, age

	25-34	35-54	> 55
< 25	4.44	17.28	33.78
	(0.04)	(0.00)	(0.00)
25-34	-	8.81	26.05
		(0.00)	(0.00)
35-54	-	-	22.03
			(0.00)

Notes: χ^2 statistics. Probability $> \chi^2$ in parenthesis.

Table C.4: Means tests of coefficient equality, education level

	Basic	Lower Sec.	Upper Sec.	Higher
Primary or less	1.98	14.77	9.67	0.12
	(0.16)	(0.00)	(0.00)	(0.73)
Basic	-	9.38	3.56	0.39
		(0.00)	(0.06)	(0.53)
Lower Sec.	-	-	1.40	5.19
			(0.24)	(0.02)
Upper Sec.	-	-	-	3.53
				(0.06)

Notes: χ^2 statistics. Probability $> \chi^2$ in parenthesis.

Table C.5: Share of temporary contracts and the Covid-19 crisis (sub-sample)

Dep. Var.:	Log of Unemployment
	(1)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T \times temp_i$	0.962**
	(0.29)
Number of Obs.	9,152
R-squared	0.980

Notes: Standard errors (in parenthesis) are clustered at NUTS II and time period (month, year) level.

Significance levels: * p < 0.10, ** p < 0.05, *** p < 0.01

Table C.6: Share of temporary contracts and the Covid-19 crisis (sub-sample)

Dep.Var.:	Log of Unemployment					
Dimension:	**Gender**		**Age**			
	Male	Female	< 25	25-34	35-54	> 55
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T \times temp_i$	1.036**	0.906*	1.429*	1.107**	0.922**	0.608*
	(0.26)	(0.36)	(0.56)	(0.34)	(0.25)	(0.23)
Number of Obs.	9,152	9,152	9,152	9,152	9,152	9,152
R-squared	0.973	0.980	0.946	0.965	0.974	0.985

Dimension:	**Education**				
	Primary or less	Basic	Lower Sec.	Upper Sec.	Higher
	(7)	(8)	(9)	(10)	(11)
$\mathbb{1}_{m \geq 3} \times \mathbb{1}_T \times temp_i$	0.703*	0.767**	1.080**	1.151**	0.873*
	(0.29)	(0.25)	(0.28)	(0.33)	(0.37)
Number of Obs.	9,152	9,152	9,152	9,152	9,152
R-squared	0.971	0.965	0.968	0.974	0.970

Notes: Standard errors (in parenthesis) are clustered at NUTS II and time period (month, year) level.
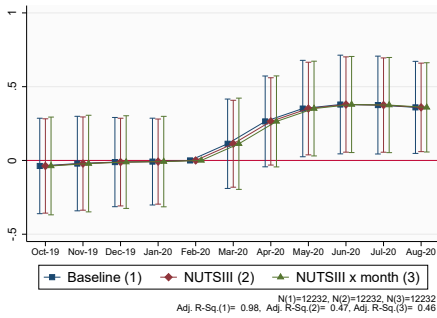
Significance levels: * p < 0.10, ** p < 0.05, *** p < 0.01

D Additional Figures

Figure D.1: Total number of firms and workers under the furlough scheme
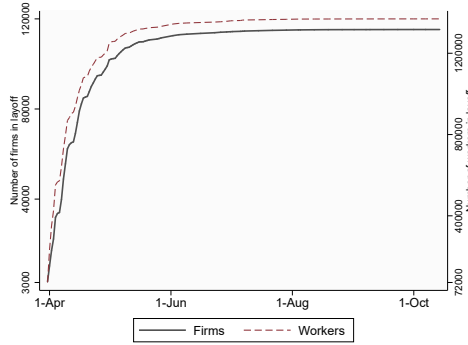
Source: GEP/MTSSS

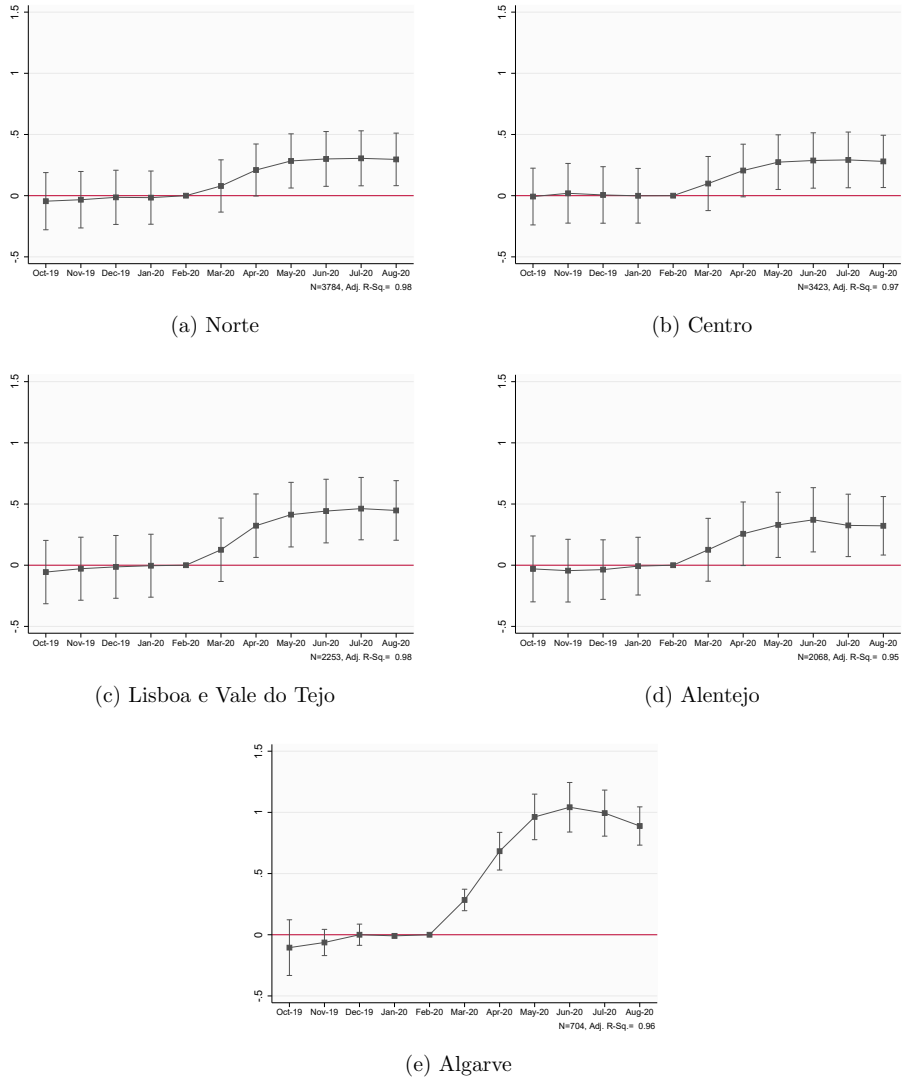Figure D.2: Event study aggregate effects: different fixed effects

(a) Unemployment

(b) New Job Placements

Figure D.3: Event study aggregate effects: unemployment by NUTS II region

(a) Norte

(b) Centro

(c) Lisboa e Vale do Tejo

(d) Alentejo

(e) Algarve

Notes: Standard errors are clustered at the municipality (instead of NUTS II) and time period level.

Figure D.4: Average new unemployment between March and August (% of total new registrations) by motive of registration
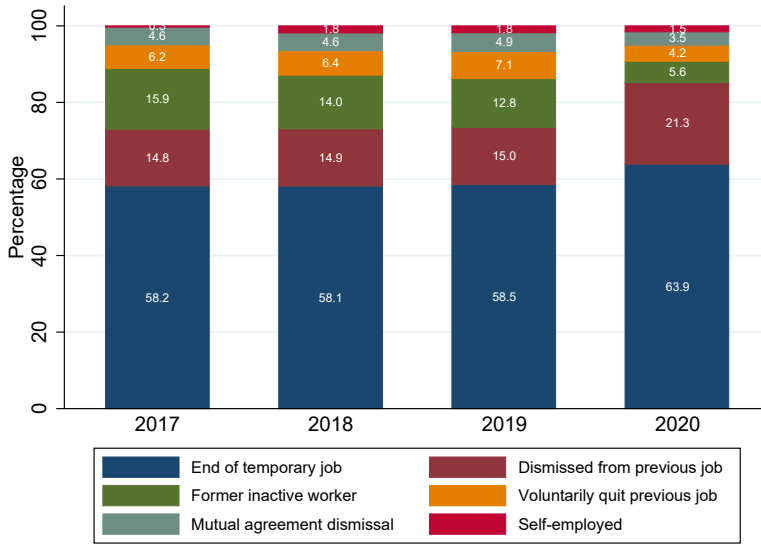

Figure D.5: Scatterplot share of temporary workers

(a) 2018 versus 2016

(b) 2018 versus 2017

Learning loss during COVID-19: An early systematic review[1]

Robin Donnelly[2] and Harry Anthony Patrinos[3]

With COVID-19 having caused significant disruption to the global education system, researchers are beginning to become concerned with the impact that this has had on student learning progress and, in particular, if learning loss has been experienced. To evaluate this, we conduct a thorough analysis of recorded learning loss evidence documented between March 2020 and March 2021. This systematic review aims to consolidate available data and document what has currently been reported in the literature. Given the novelty of the subject, eight studies were identified; seven of these found evidence of student learning loss amongst at least some of the participants, while one of the seven also found instances of learning gains in a particular subgroup. The remaining study found increased learning gains in their participants. Additionally, four of the studies observed increases in inequality where certain demographics of students experienced learning losses more significant than others. It is determined that further research is needed to increase the quantity of studies produced, their geographical focus, and the numbers of students they observe.

1. Introduction

Due to the emergency nature of the COVID-19 pandemic, education systems around the world are facing extreme disruption. At its peak, UNESCO (2020) reported that nearly 1.6 billion learners in more than 190 countries, or 94 percent of the world's student population, were impacted by educational institution closures. Given the abruptness of the situation, teachers and administrations were unprepared for this transition and were forced to build emergency remote learning systems almost immediately. In response to this disruption, education researchers are beginning to analyze the impact of these school closures on student learning progress or lack thereof.

The term "learning loss" is commonly used in the literature to describe declines in student knowledge and skills (Pier et al. 2021). Historic data provides researchers with information regarding where student learning should be year over year and is often measured through regular testing. Learning loss occurs when educational progress does not occur at the same rate it historically has compared to previous years (Pier et al. 2021).

Outside of the classroom, these losses may translate to greater long-term challenges. Currie and Thomas (2001) put this into perspective as they observe that a 0.20 decrease in standardized test scores could decrease future employment probability by 0.86 percent. Additionally, Chetty et al. (2014) observe that increasing student achievement by 0.20 standard deviations results on average, a 2.6 percent increase in annual lifetime earnings (Maldonado and De Witte 2020). Likewise, on average, another year of schooling is associated with an 8-9 percent gain in future earnings (Psacharopoulos and Patrinos 2018).

While many researchers have established learning loss prediction models (e.g., Azevedo et al. 2020), formal research and documentation of the actual impact that COVID-19 has had on student learning progress is just beginning to emerge. As the global education system continues to face pandemic related disruption, a strong understanding of how COVID-19 school closures are impacting student learning progress can better equip educators, policy makers, and researchers going forward.

Our focus in this paper is on the important research question: have COVD-19 education closures resulted in recorded student learning losses? To answer this, we conduct a thorough analysis of recorded learning loss evidence documented between March 1, 2020 and March 18, 2021. This systematic review aims to consolidate such data and document what has currently been reported in the literature. To the best of our knowledge, this is the first systematic review of its kind.

Our paper makes three important contributions. First, we develop a comprehensive review that consolidates the research that has been presented related to the impact of COVID-19 on student learning progress. Second, we answer the research question: has pandemic related learning loss been recorded in the literature? Third, based on our review, we identify significant gaps in the literature and provide relevant guidance for further research.

We begin by describing the methods used in this review to identify and collect the articles analyzed. We then present an analytical review where each article was categorized by its geographical region studied, length of school closure, education level of students analyzed, subject

analyzed, documented learning impact, presence of increased inequality, and sample size. Next, we move on to our discussion where we review the findings of the analytical review. We then end by discussing areas for future research and summarizing the main ideas of this paper.

2. Methodology

The initial search was performed using English-language articles published between March 1, 2020 and March 18, 2021. To ensure a comprehensive, multidisciplinary search, included electronic databases were ECONLIT, Google Scholar, PubMed, Education Resources Information Center (ERIC), and Cochrane Library. To conduct the search, the key words "covid-19", "coronavirus", "2019-ncov," "sars-cov-2," or "cov-19," were used in combination with "learning loss," "learning slide," "education gap," or "achievement gap." Along with this, some studies were identified by reaching out to colleagues and researchers.

When conducting the search analysis, thousands of articles were identified; however, the majority of these pertained to hypothesized or predicted learning loss. To narrow this down to studies with recorded results, article abstracts were then screened. Studies which conducted student analyses and reported impacts on learning progress (either positive, negative, or insignificant) as a result of COVID-19 school disruptions were included. After this screening process, eight articles remained (Table 1). Factors for rejecting studies from our review included the absence of a student analyses and/or recorded impacts on learning progress, analyses occurring before the onset of COVID-19, or hypothesized results.

Table 1: Studies Included

Title	Authors
Learning inequality during the COVID-19 pandemic	Engzell, Frey, Verhagen 2020
Influence of COVID-19 confinement on students' performance in higher education	Gonzalez, Rubia, Hincz, Comas-Lopez, Subirats, Fort, Sacha 2020
The impact of COVID-19 on student learning in New South Wales primary schools: an empirical study	Gore, Fray, Miller, Harris, Taggart 2021
Collaborative for student growth. Learning during COVID-19: Initial findings on students' reading and math achievement and growth	Kuhfeld, Tarasawa, Johnson, Ruzek, Lewis 2020
The effect of school closures on standardised student test outcomes	Maldonado, De Witte, 2020
Learning during the COVID-19 pandemic: It is not who you teach, but how you teach	Orlov, McKee, Berry, Boyle, DiCiccio, Ransom, Reese-Jones, Stoye 2020
Educational gains of in-person vs. distance learning in primary and secondary schools: a natural experiment during the COVID-19 pandemic school closures in Switzerland	Tomasik, Helbling, Moser 2020
Did students learn less during the COVID-19 pandemic? Reading and math competencies before and after the first pandemic wave	Schult, Mahler, Fauth, Lindner 2021

3   Analytical Review

Once articles were selected, they were then coded using the classifications in Table 2.

Table 2: Classifications Used to Analyze Studies

Classification Term	Description
Country	The residing nation of the study's participants.
Closure Length	The number of days that the participants were out of in person traditional schooling prior to assessment.
Education Level	Education level of participants.
Subject	Course subject of participants.
Learning Loss	Documented level of learning loss experienced by participants. If gains were experienced, "Improved" was listed.
Equality Impact	Documented differences in the level of loss experienced by certain groups of students.
Sample Size	Sample size of students analyzed.

As indicated in Table 3, we find that seven out of the eight studies identified learning loss amongst at least some of the students analyzed. For example, Maldonado and De Witte (2020) found Grade 6 students in Belgium experienced losses of 0.19 SD in math and 0.29 SD in Dutch. Engzell et al. (2021) find that overall, Grade 4-7 students in the Netherlands have encountered an average 0.08 SD learning loss in math, spelling, and reading. Tomasik et al. (2020) found learning progress of primary school students in Switzerland during in-person learning to be more than twice as high compared to the progress made during the eight-week school closure. Orlov et al. (2020) determined that economics students at four USA universities were 0.19 SD behind. Kuhfeld et al. (2020) found that Grade 3-8 students in the USA scored 5-10 percentile points below historic levels in math. Gore et al. (2021) found Year 3 students studying math in low ICSEA schools (Index of Community Socio-Educational Advantage schools) to be two months behind the progress students made in 2019 in Australia. Lastly, Schult et al. (2021) find learning losses of 0.07 SD in reading comprehension, 0.09 in operations, and 0.03 in numbers for Grade 5 students in Germany. At the university level, in a single university, for 458 students in STEM faculties, learning outcomes actually improved (Gonzalez et al. 2020).

Table 3: Results of Literature Classification

Source	Country	Closure Length	Education Level	Subject	Learning Loss	Equality Impact	Sample Size
Maldonado & De Witte, 2020	Belgium	9 weeks	Primary, Grade 6	Math	0.19 SD	Yes	Not specified
			Primary, Grade 6	Dutch	0.29 SD	Yes	Not specified
			Primary, Grade 6	Social Science	Insignificant	Not specified	Not specified
Engzell et al.,2021	Netherlands	8 weeks	Primary (Age 8)	Math	0.063 SD	Not specified	92180 students
			Primary (Age 8)	Reading	0.05725 SD	Not specified	76397 students
			Primary (Age 8)	Spelling	0.09375 SD	Not specified	90403 students
			Primary (Age 9)	Math	0.07325 SD	Not specified	93417 students
			Primary (Age 9)	Reading	0.0975 SD	Not specified	79016 students
			Primary (Age 9)	Spelling	0.07075 SD	Not specified	91567 students
			Primary (Age 10)	Math	0.0935 SD	Not specified	93769 students
			Primary (Age 10)	Reading	0.08425 SD	Not specified	68412 students
			Primary (Age 10)	Spelling	0.0755 SD	Not specified	91315 students
			Primary (Age 11)	Math	0.05025 SD	Not specified	73263 students
			Primary (Age 11)	Reading	0.07425 SD	Not specified	48537 students
			Primary (Age 11)	Spelling	0.07575 SD	Not specified	69841 students
			Primary (Grade 4-7)	Math, Spelling, Reading	0.08 SD	Yes	350 000 students
Tomasik et al., 2020	Switzerland	8 weeks	Primary (Grade 3-6)	Math, German	2X	Not specified	13134 Students
			Secondary (Grade 7-9)	Math, German	Insignificant	Not specified	15551 Students
Gonzalez, et al. 2020	Spain	10 weeks	Higher ed	Applied Computing, Metabolism, Design of Water Treatment Facilities	Improved	Not specified	458 Students
Orlov et al., 2020	USA	3.5 weeks	Higher ed	Economics	0.185 SD	No	4 Universities
Kuhfeld et al., 2020	USA	Not specified	Primary (Grade 3-8)	Math	5-10 percentile points	Inconclusive	4.4 million students
			Primary (Grade 3-8)	Reading	Insignificant	Inconclusive	4.4 million students
Gore et al., 2021	Australia	8-10 weeks	Primary (Year 3, all schools)	Math	Insignificant	Yes	1427 students
			Primary (Year 3, low ICSEA schools)	Math	2 months less growth	Yes	334 students
			Primary (Year 3, mid ICSEA schools)	Math	Improved, 2 months additional growth	Yes	813 students
			Primary (Year 3)	Reading	Insignificant	No	1429 students
			Primary (Year 4)	Math	Insignificant	No	1498 students
			Primary (Year 4)	Reading	Insignificant	No	1515 students
			Primary (Year 3-4)	Math and Reading	Insignificant	No	3030 students
Schult et al., 2021	Germany	8.5 weeks	Primary (Grade 5)	Math	0.09 SD 0.03 SD	Yes Yes	80000 students 80000 students
				Reading	0.07 SD	Yes	80000 students

4.   Discussion

Learning loss is being experienced. As Table 3. indicates, the early findings of seven studies provide evidence of experienced learning losses amongst students. These observed losses are occurring across a range of subjects, grade levels, and geographical regions. This signals that although robust and empirical research on COVID-19 related student learning loss is limited, learning loss itself may not be.

Not all students are experiencing learning loss. While the majority of the literature analyzed indicates that students are experiencing some level of learning loss, there were also instances where this was not the case. For example, both Maldonado and De Witte (2020) as well as Kuhfeld et al. (2020) found learning losses in certain subjects but insignificant impacts in others. Likewise, while Tomasik et al. (2020) found primary students to be impacted, they found no impact on secondary students. This is consistent with the literature showing that students in the early grades may be more vulnerable than secondary students because of their inability to seek learning on their own, due to the differences in developmental and cognitive abilities. In their Australian study, Gore et al. (2021) found overall there to be no evidence of learning loss in Year 3 and 4 students in math and reading with the exceptions being Year 3 students in math in low ICSEA schools who experiences losses and while mid ICSEA students experienced small gains. Lastly, in the case of Gonzalez et al. (2020) who studied university students in Spain, it was determined that student learning progress actually improved rather than declined during the COVID-19 learning disruption period, but this was for university students in STEM subjects at one university.

Some students are experiencing more learning loss than others. Of the eight studies, four found instances of in inequality, while only one exclusively found demographics to have no impact on learning loss. Gore et al. (2021) found instances of increased inequality as well as instances of no change. The other studies did not specify in this area or in the case of Kuhfeld et al. (2020) found inconclusive and minor differences between ethnic/racial groups. In the four studies where increases in inequality were observed, certain demographics of students experienced losses more significant than others. Maldonado and De Witte (2020) observed inequality within schools rise by 17 percent for math and 20 percent for Dutch. Engzell et al. (2021) determined that losses were up to 60 percent larger amongst students from uneducated homes. Gore et al. (2021) found the only losses to be amongst students from low ICSEA schools where the lower the ICSEA level the lower the educational advantage attending students have due to their parents' occupation and education, their geographical location, and the school's proportion of indigenous students (Australian Curriculum, Assessment and Reporting Authority, 2016). Schult et al. (2021) found losses in math amongst grade 5 students to be more severe in low achieving students. In reading comprehension Schult et al. (2021) found more severe losses amongst middle to high achieving students.

More research is needed. In general, the literature representing the impact that COVID-19 has had on student learning progress is limited in the quantity of studies available, geographical regions analyzed, and number of participating students.

Given the novelty of the subject, it is understandable why education researchers are only just beginning to analyze the learning losses that students have experienced. However, a stronger

understanding of how COVID-19 school disruptions have impacted student learning is still needed. To support this, more studies are needed.

Along with this, the current studies that are available are limited in their geographical span. The only limited information that is currently available is from Belgium, the Netherlands, Switzerland, Spain, the United States, and Australia. Given the differences in educational institutions between countries, in terms of quality, length of school closures, and remote learning strategies, it is crucial that researchers continue to investigate COVID-19 related learning loss in countries where limited research exists.

Lastly many of the studies themselves that were analyzed in this systematic review had limited numbers of participants. For example, Gonzalez et al. (2020) analyzed just 458 students at 1 university. Similarly, Orlov et al. (2020) observed economics students in just 7 classes across 4 universities. While the information these studies presented remains relevant to their observed samples, research that can more accurately represent larger groups of students remains crucial to policy makers. As such, there is a demand for studies that analyze representative groups of students.

5. Conclusion

Through conducting a thorough analysis of recorded learning loss evidence documented between March 2020 and March 2021, this systematic review provided a consolidated audit of available research on COVID-19 related learning loss. Given the novelty of the subject, eight studies were identified; seven of the eight found evidence of student learning loss amongst participants, while one of these found instances of learning gains in a particular subgroup. The remaining study observed learning gains amongst university students. Along with this, four of the studies observed increases in inequality where certain demographics of students experienced learning losses more significant than others. Further research is needed to increase the quantity of studies produced, their geographical focus, and the numbers of students they observe.

CEPR PRESS

References

Australian Curriculum, Assessment and Reporting Authority. (2015). *What does the ICSEA value mean?* https://docs.acara.edu.au/resources/About_icsea_2014.pdf

Azevedo, J.P., Hasan, A., Goldemberg, D., Iqbal, S.A., & Geven, K. (2020). Simulating the potential impacts of COVID-19 school closures on schooling and learning outcomes. World Bank Policy Research Working Paper No. 9284. World Bank, Washington, DC. https://openknowledge.worldbank.org/handle/10986/33945

Currie, J., & Thomas, D. (2001). Early test scores, school quality and SES: Longrun effects on wage and employment outcomes. *Research in Labor Economics, 20,* 103–132. https://doi.org/10.1016/S0147-9121(01)20039-9

Chetty, R., Friedman, J. N., & Rockoff, J. E. (2014). Measuring the impacts of teachers II: Teacher value-added and student outcomes in adulthood. *American Economic Review, 104*(9), 2633–2679. https://doi.org/10.1257/aer.104.9.2633

Engzell, P., Frey, A., & Verhagen, M. (2020). Learning inequality during the COVID-19 pandemic. https://doi.org/10.31235/osf.io/ve4z7

Gonzalez, T., de la Rubia, M.A., Hincz, K.P., Comas-Lopez, M., Subirats, L., Fort, S., & Sacha, G.M. (2020). Influence of COVID-19 confinement on students' performance in higher education. *PloS One*. https://doi.org/10.1371/journal.pone.0239490

Gore, J., Fray, L., Miller, D., Harris, J., & Taggart, W. (2021). The impact of COVID-19 on student learning in New South Wales primary schools: an empirical study. *The Australian Educational Researcher.* https://doi.org/10.1007/s13384-021-00436-w

Kuhfeld, M., Tarasawa, B., Johnson, A., Ruzek, E., & Lewis, K. (2020). Collaborative for student growth. Learning during COVID-19: Initial findings on students' reading and math achievement and growth. *NWEA Research.* https://www.nwea.org/content/uploads/2020/11/Collaborative-brief-Learning-during-COVID-19.NOV2020.pdf

Maldonado, J., & De Witte, K. (2020). The effect of school closures on standardised student test outcomes. *KU Leuven – Faculty of Economics and Business.* https://feb.kuleuven.be/research/economics/ces/documents/DPS/2020/dps2017.pdf

Orlov, G., McKee, D., Berry, J., Boyle, A., DiCiccio, T., Ransom, T., Reese-Jones, A., & Stoye, J. (2020). Learning during the COVID-19 pandemic: It is not who you teach, but how you teach. Working Paper No. 28022. *National Bureau of Economic Research.* https://doi.org/10.3386/w28022

Pier, L., Hough, H.J., Christian, M., Bookman, N., Wilkenfeld, B., & Miller, R. (2021, January

25). *COVID-19 and the educational equity crisis. Evidence on learning loss from the CORE data collaborative.* PACE. https://www.edpolicyinca.org/newsroom/covid-19-and-educational-equity-crisis

Psacharopoulos, G., & Patrinos, H. (2018). Returns to investment in education: a decennial review of the global literature. Policy Research Working Paper No. 8402. World Bank, Washington, DC. https://openknowledge.worldbank.org/handle/10986/29672

Schult, J., Mahler, N., Fauth, B., & Lindner, M. (2021). Did students learn less during the COVID-19 pandemic? Reading and math competencies before and after the first pandemic wave. https://doi.org/10.31234/osf.io/pqtgf

Tomasik, M., Helbling, L., & Moser, U. (2020). Educational gains of in-person vs. distance learning in primary and secondary schools: a natural experiment during the COVID-19 pandemic school closures in Switzerland. *International Journal of Psychology.* https://doi.org/10.1002/ijop.12728

UNESCO. (2020, June 6). *UN Secretary-General warns of education catastrophe, pointing to UNESCO estimate of 24 million learners at risk of dropping out.* https://en.unesco.org/news/secretary-general-warns-education-catastrophe-pointing-unesco-estimate-24-million-learners-0