Mobility under the COVID-19 Pandemic: Asymmetric Effects across Gender and Age*

Francesca Caselli, IMF
Francesco Grigoli, IMF
Damiano Sandri, IMF and CEPR
Antonio Spilimbergo, IMF and CEPR

October 20, 2020

PRELIMINARY VERSION

Abstract

Using unique mobility indicators based on anonymized data provided by Vodafone for Italy, Portugal, and Spain, we find that lockdowns, especially school closures, reduced more the mobility of women more strongly. Lockdowns also had a stronger impact on the mobility of younger cohorts. Furthermore, rising COVID-19 infections reduce more the mobility of younger people. These findings are based on regression discontinuity design and local projections, and are robust to a variety of tests. These differential effects on mobility may widen gender and inter-generational inequality.

Keywords: COVID-19, lockdown, mobility, gender, age.

JEL Codes: E1, I1, H0.

*Francesca Caselli, fcaselli@imf.org; Francesco Grigoli, fgrigoli@imf.org; Damiano Sandri, dsandri@imf.org; Antonio Spilimbergo, aspilimbergo@imf.org.

The views expressed in this working paper are those of the authors and do not necessarily represent those of the IMF, its Executive Board, or its management. Working papers describe research in progress by the authors and are published to elicit comments and to encourage debate.

Contents

List of Figures

List of Tables

1 Introduction

The COVID-19 pandemic has dramatically reduced people's mobility. This was due in part to the lockdown measures that governments adopted to reduce infections, including travel restrictions, school and business closures, and stay-at-home orders. Mobility also declined because people voluntarily reduced social interactions out of fear of the virus. The literature has documented these effects using a broad range of *aggregate* mobility indicators provided by private companies such as Google, Apple, and SafeGraph

This paper innovates relative to existing studies by showing that mobility patterns have differed considerably across gender and age groups. Our analysis is based on novel and confidential mobility indicators provided by Vodafone for Italy, Portugal, and Spain at the provincial level.[1] These data offer the unique advantage of disaggregating mobility information across gender and age groups. This makes it possible to shed light on heterogeneous reactions to the pandemic and lockdown measures. This papers makes several contributions.

First, the analysis contributes to the growing evidence about the disproportionate impact of the COVID-19 crisis on women. These could happen for a variety of reasons. For instance, during the pandemics part of home productions that in normal times can be outsourced had to be performed within the household; it could be the the burden felt disproportionately on women. Indeed, Hupkau and Petrangolo, 2020 find that in the UK within the household, women provided on average a larger share of increased childcare needs, but in an important share of households fathers became the primary childcare providers. Alon et al. (2020) show that contrary to past recessions, the current crisis has led to a stronger increase in women unemployment in the US. This is because women are more likely to care for children when schools are closed and because they are employed in sectors more severely hit by the pandemic, such as restaurants and personal care. Survey data also suggest that women face an unequal burden in caring for children when schools are closed and are at a higher risk of facing a reduction in working hours (Adams et al., 2020; Sevilla and Smith, 2020). These papers have focused on the *consequences* of differential mobility across gender; our contribution focuses on the documenting the differential mobility and in understanding the *causes*.

In this regard, this paper contributes also to the literature on the determinants of labor force participation. School closure and lockdowns by reducing mobility in a different ways across gender have also a differential effect on labor supply. Previous studies have found that

[1] We focus on Italy, Portugal, and Spain because reliable data provided by Vodafone were available for these countries.

(exogenous changes in the) length of school schedule impact female labor force participation (Berthelon et al., 2015.) We provide complementary and high frequency evidence.

Second, this paper documents considerable heterogeneity in the impact of lockdowns across age groups. These findings are quite relevant to inform the ongoing debate on the distribution of costs and benefits across generations. While lockdowns protect especially older people that are more likely to develop serious medical conditions because of COVID-19, business closures, school shutdowns, and stay-at-home orders impose economic costs especially on working age people.

Third, the heterogeneous impact of lockdowns across age can shed light on the scarring effects from the crisis. For instance, given a reduction in aggregate mobility has different long-term effects if it is concentrated only on the young or only on the older people. Knowing the age composition of the reduction in mobility is a pre-condition to study the channels of possible scarring effects. This is analogous to the fact that the type of unemployment (long vs. short-term) is crucial to start a discussion on the long-term effects of a recession.

Fourth, the heterogeneous effects across gender and age uncovered by the analysis proved raise important methodological considerations. Lumping all groups together in estimations of the effects of containment measures could lead to aggregation bias, and underestimation of the effects of containment measures.

Firth, in presence of repeated waves in several countries (as we write, Europe is experiencing a strong resurgence) authorities need to consider more nuanced containment and mitigation measures. The only way to design these measures properly is by considering the effects on several groups. This perspective is particularly important because the health risks posed by COVID-19 are very heterogeneous across age, being much more severe for people aged 65 and above. Therefore, some researches have argued for targeted measures to isolate older people without unduly limiting the mobility and employment opportunities of younger people (Acemoglu et al., 2020).

Studying the effect of containment measures and fear factors across groups poses several methodological challenges. First, the events considered (stay-at-home order, school closure, and lifting of the lockdown) typically happened just once. Second, these events are endogenous as mentioned before. Third, the effects of these events are distributed over time. For these reasons, this paper uses regression discontinuity design and local projections.

First, we use a regression discontinuity design to show that women's mobility dropped more than men's when Italy, Portugal, and Spain adopted national stay-at-home orders. This generally coincided with school closures, making identification of the channel difficult. We then narrow the analysis on a few provinces that imposed school closures before stay-at-home orders. This shows that the mobility gender gap opens already when schools are closed, highlighting the uneven role of women in caring for children.

Second, we examine the impact of lockdowns across gender using local projections that exploit more systemically the variation in the stringency of lockdowns over time, as measured by the University of Oxford's Coronavirus Government Response Tracker. As other works on the effect of containment on mobility, this paper must address the issue of endogeneity and omitted variable. In fact, governments adjust the tightness of lockdowns depending on the stage of the epidemic, for example by tightening restrictions when infections are growing. At that time, people also voluntarily reduce mobility as they fear contracting or spreading the virus. Furthermore, governments are also likely to tighten lockdowns if mobility is too elevated.

To alleviate these endogeneity concerns, the local projections analyze the impact of lockdowns while controlling for the number of COVID-19 infections and lags of the mobility indicator. This is to account for the severity of the country's epidemic and for mobility trends. Furthermore, the local projections use sub-national data to assess the impact of national lockdown measures on the mobility in regions with relatively low infections. This considerably strengthens identification since the adoption of national lockdowns in Italy, Portugal, and Spain was driven by localized major outbreaks and it was thus largely exogenous to the conditions prevailing in regions with few infections.

The local projections corroborate the results of the regression discontinuity analysis, showing that lockdowns have a disproportionate effect on the mobility of women. The differential impact is statistically significant and quantitatively relevant. A full lockdown—including all measures used by governments during the pandemic—reduces the number of women leaving home in a given day by almost 30 percent, against an impact on men of about 20 percent.

The local projections also allow to examine the effects of rising COVID-19 cases on mobility holding constant the stringency of lockdowns. This captures the extent to which people decide to voluntarily limit social interactions when the fear of contracting the virus becomes more acute. Examining this aspect is quite relevant since much of the public debate on the need for lockdowns has centered on whether people can autonomously change behavior when infection risks arise. The analysis shows that both men and women significantly reduce mobility when infections increase and they do so with equal intensity.

Using regression discontinuity designs, we find that stay-at-home orders disproportionately reduce the mobility of working age people, especially those below 4 years of age. Local projections provide additional evidence that lockdowns have a larger impact on working age people. These findings are concerning because younger workers generally rely on labor income to support consumption, while older people have access to personal savings and possibly retirement income. Furthermore, younger workers often have less stable job contracts that are more likely to be terminated during a crisis. Survey level evidence confirms that younger people have been indeed more likely to suffer an income loss during the pandemic (Belot et al., 2020; Montenovo et al., 2020). The fact that lockdowns impose a disproportionate economic burden on the young—while protecting mostly the old given the higher health risks— calls for policy intervention to ensure greater inter-generational fairness (Glover et al., 2020).

Using local projections, we also explore if people of different age respond differently to rising infections. Because COVID-19 poses much greater risks for people aged 65 and above, it would be reasonable to expect these people to be more likely to isolate themselves when infections are rising. On the contrary, the analysis shows that younger people reduce mobility more strongly when infections increase. This is consistent with survey level evidence presented by Bordalo et al. (2020) showing that younger people are more alarmed by the risk of contracting COVID-19 .

The paper is structured as follows. The next section provides an overview of the data provided by Vodafone. Section 2 provides a few background information on the covid-19 crises in Italy, Portugal, and Spain, and the relative containment measures. Section 4 and 5 present the analysis of the mobility patterns by gender and age, respectively. Section 6 examines the robustness of the results. Section 7 concludes.

2 The COVID-19 crises in Italy, Portugal, and Spain

In this section we describe the context of the COVID-19 crisis in Italy, Portugal, focusing on the impact of COVID-19 directly and the containment measures.

Italy Italy was among the first countries to be hit by COVID-19 after China. On 31 January, the Italian government declared a state of emergency and stopped flight from and to China. Apart from two Chinese tourists who were promptly isolated, there was no confirmed case until February 21 when a patient with anomalous pneumonia was diagnosed with COVID-19 in Codogno, Lombardy. Shortly after that new cases were discovered in other towns in Lombardy and Veneto. On February 22, a decree imposed the quarantine of more than 50,000 people from 11 municipalities (*comuni*) in Northern Italy (so called *zone rosse*). In other areas of Emilia-Romagna, Lombardy, and Veneto (*zone gialle*) schools, theatres, clubs, and cinemas are closed and social and sports events were suspended. On March 4, all schools and universities across

Italy were closed for two weeks and all sporting events could be played only behind closed doors until April 3. As the outbreaks continued and the number of deaths soared, on March 8, all 12 provinces in Lombardy and 14 provinces in Piedmont, Emilia-Romagna, Veneto, and Marche, were put under lockdown. Two days later, the lockdown was extended to the whole country. Steep penalties were announced for violators, including the possibility of three months of imprisonment. On March 11, the government prohibited almost all commercial activities except for supermarkets and pharmacies. On March 21, all non-essential businesses and industries were closed, and movement of people was restricted. In May, many restrictions were progressively eased, and, freedom of movement across regions and other European countries was restored on June 3.

Portugal The first cases of COVID-19 in Portugal were recorded on March 2. On March 18, the entire Portuguese territory entered in a State of Emergency, which lasted until May 2. During the Easter week (April 9 to 13), the government decreed special measures to restrict people movements between municipalities (*concelhos*) with few exceptions, closing all airports to civil transportation. On May 4, restrictions started to be eased and small stores reopened. On May 18, nurseries and the last two years of the secondary school, restaurants, cafes, medium-sized stores and some museums reopened.

Spain The first case of a patient with COVID-19 in Spain was a foreign tourist on January 31 but broader diffusion began by mid-February and, by mid-March, all 50 provinces had confirmed cases. A lockdown was imposed on March 14. Starting March 30, all non-essential workers were ordered to remain at home for the next two weeks. Covid-19 spread rapidly and by March 25, the official death toll in Spain surpassed that of China with most cases concentrated in Madrid. The number of deaths peaked in early April and progressively declined; June 1 was the first day without COVID-19 related deaths. The first local lockdown was announced on March 7 for a small municipality. On March 12, the lockdown was extended to four municipalites in Catalunya with 70,000 people affected. On March 14, the entire country entered in the state of emergency and many nonessential activities were forbidden, including large gatherings, restaurants, museums. Citizens were still permitted to travel to work and buy essential items. The authorities in some autonomous communities, including the Basque Country, Murcia, Balearic Islands, Catalunya, announced additional emergency measures. On March 28, all non-essential workers were ordered to stay home from March 30 to April 9. Progressive easing of the lockdown started at the beginning of May. On May 11, the opening of small shops, of terraces at half capacity, and of places of worship at one-third capacity was allowed in 26 provinces and territories comprising about half of the population.

3 A unique dataset

We use data on mobility kindly provided by Vodafone through a confidential agreement. By analysing connections to cell towers, Vodafone can create mobility indexes differentiated across gender and age groups using the information that customers provide when signing up for post-paid contracts.[2] The age groups include the following categories: people aged between 18 and 24, between 25 and 44, between 45 and 64, and 65 and above.[3]

More specifically, the mobility indicator used in the analysis captures the percentage of people in a given province and demographic group that leaves home in a day. The home location of each customer is identified by monitoring cell connections during the night. The top 3 cells that a phone connects to between 10pm and 5am are considered as home cells. A customer is recorded as leaving home if the phone connects to a cell different from the home cells. More details on the data construction are provided in Kariotis et al. (2020).

The mobility patterns detected by Vodafone are broadly in line with those recorded by Apple and Google data.[4] Figure 1 shows indeed that all indicators correlate fairly closely at the national level. Correlations between the Vodafone indicator and the Apple and Google indicators range between 93 and 99 percent for Italy and Portugal and between 72 and 88 percent for Spain. Moreover, the disaggregation of Vodafone data allows to appreciate the heterogeneity across provinces. In all the three countries, the interdecile range of the mobility indicator is as large as 20 percentage points. Yet, such dispersion remains broadly constant over time.

The key advantage provided by the Vodafone data is the ability to differentiate mobility

[2]These indicators were prepared by Vodafone's Big Data and Artificial Intelligence team and provided for the analysis through a confidential agreement. To protect the privacy of individuals and minority groups, the data have been provided in anonymized form, reporting the average mobility for a given gender and age group at the provincial (NUTS3) level when a minimum of 50 customers are available. Furthermore, the data sharing protocol was subject to technical and organizational controls including an ethical assessment of the analysis prior to its implementation.

[3]In a few cases, the age information is inferred. For example, in Spain, the age group 18-24 is separated from family contracts based on the amount of data used. Furthermore, in Portugal customers' age is based on the sequential number of personal identification cards that allow to infer people's age with an error of five years at most. Vodafone clients can opt-out from location tracking: 5.5 percent of post-paid customers opted out in Spain, and [...] in Italy.

[4]Apple mobility data are available at https://covid19.apple.com/mobility, and Google mobility data are available at https://www.google.com/covid19/mobility.

Figure 1: Mobility Levels from Apple, Google, and Vodafone

(a) Italy

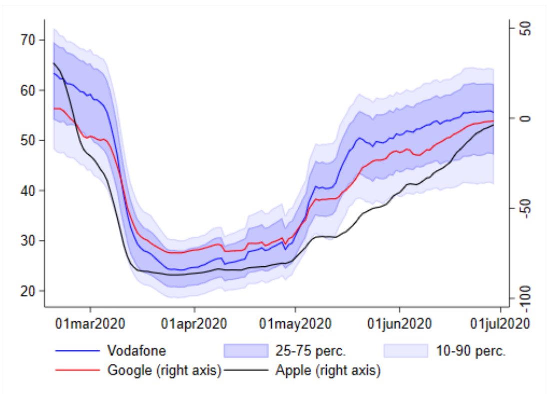

(b) Portugal

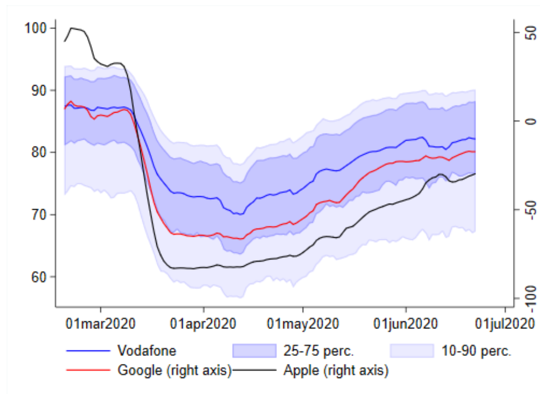

(c) Spain

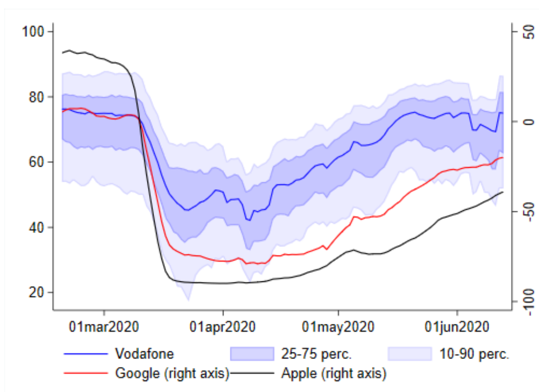

Notes: The lines denote the country-level mobility levels. In the case of Vodafone, the line corresponds to the cross-province population-weighted average of the percent of people moving, using 2018 population levels as weights; and the shaded areas denote the cross-province interquartile range (dark blue) and the cross-province interdecile range (light blue). In the case of Google, the line corresponds to the average of country-level mobility indicators at retail, grocery, parks, transit, and workplace locations, where mobility is defined relative to pre-crisis levels. And in the case of Apple, the line denotes the country-level indicator of mobility, which is computed from the number of requests made to Apple Maps for directions.

across gender and age groups. This makes it possible to address important questions, such as whether lockdowns have heterogeneous effects on people's mobility depending on gender and age. Figure 2 provides preliminary evidence in this regard. Panel 2a shows for each country how the average mobility differential between women and men correlates with the stringency of lockdowns over the period of analysis. In all countries lockdowns have been associated with a larger drop in women' mobility relative to men's. The other three panels in Figure 2 show how the stringency of lockdowns correlated with the mobility differential relative to the oldest age category of 65 and above. The charts suggests that lockdowns have an heterogeneous effect also across age groups, reducing more severely the mobility of younger people. The rest of the analysis will test more formally for these patterns using regression discountinuity approaches and local projections.

4 Heterogeneous effects on mobility across gender

In this section, we examine whether lockdowns have a different effect on the mobility of women and men. Assessing the impact of lockdowns on mobility is a challenging task since the decision to deploy lockdowns is not random. For example, governments are more likely to impose lockdowns when health risks become more acute. At that time, people voluntarily reduce social interactions because they fear being infected or infecting others. This may generate a spurious correlation between the introduction of lockdowns and the reduction in mobility. Policymakers may also impose lockdowns when mobility is too high, thus leading to a spurious association between lockdowns and high mobility.

To alleviate these endogeneity concerns we use two empirical strategies. First, we employ regression discontinuity designs that focus on high-frequency changes in mobility around specific lockdown measures, thus reducing the risk that other factors may affect mobility at the same time. Second, we use local projections that control for lagged mobility and for the severity of the country's epidemic based on the number of new infections. To further strengthen identification, local projections are estimated using data from regions that did not experience severe outbreaks and thus for which the adoption of national lockdowns was mostly an exogenous event. The use of local projections will also be instrumental to examine how people voluntarily reduce mobility in response to rising infections.

Figure 2: Mobility and Lockdown Stringency

(a) Women-men differential

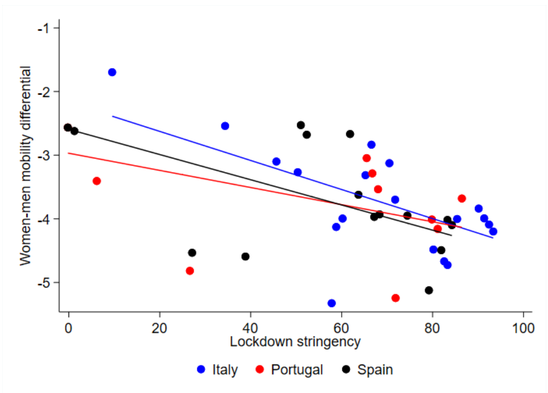

(b) 18–24, differential with 65+

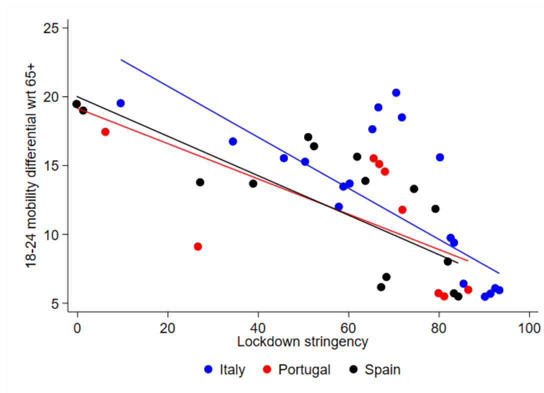

(c) 25–44, differential with 65+

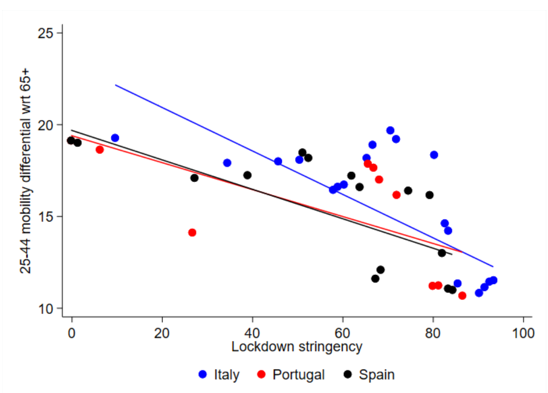

(d) 45–64, differential with 65+

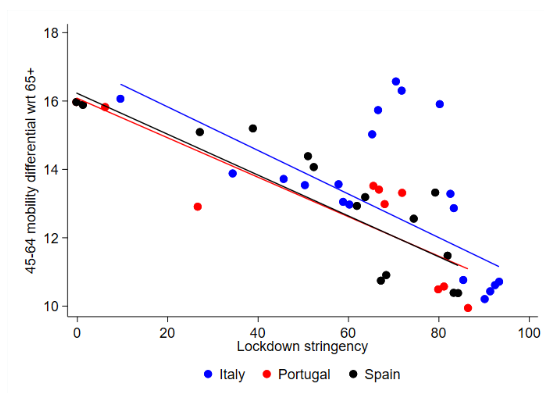

Notes: The figure presents a binned scatter plots showing the association between the mobility differential between women and men and the stringency of lockdowns over the period of analysis. Each dot denotes the cross-province average at any given time. The percent of people moving is residualized with respect to days of the week fixed effects. The lines denote the linear fit.

4.1 Regression discontinuity

To test whether lockdowns have unequal effects across gender, we first use a regression discontinuity (RD) approach in a similar spirit to Davis (2008), Anderson (2014), and Chetty et al. (2020). With respect to a standard cross-sectional RD setting, in this case the running variable is time and the treatment date is the threshold, making this approach akin to an event study exercise. As in more standard RDs, endogeneity is addressed by considering a narrow bandwidth (in this case a time window) around the introduction of the treatment. The identification assumption is that, within this interval, unobserved confounding factors affecting the outcome variable are likely to be similar. In our context, this means that no other factors affecting mobility should change close to national stay-at-home orders.[5]

Figure 3 uses a bin scatter plot to present preliminary evidence that lockdowns are associated with a discontinuity in the mobility of women relative to men. Each dot represents the average mobility levels of men and women calculated within 20 equally sized bins around the introduction of national stay-at-home orders. We start by considering people aged between 25-44. Mobility data are residualized with respect to province and day-of-the week fixed effects. The figure shows that the introduction of stay-at-home orders led to a sharp drop in the mobility of both men and women. The percentage of people living their homes in a day declined by about 15 percentage points. Yet, the impact on women was stronger, as their mobility declined by about 3 percentage points more than for men (see also Table 1).

In most provinces, the adoption of stay-at-home orders coincided with or rapidly followed the decision to close schools. Therefore, the gender gap in Figure 3 could be driven by women carrying a disproportionate burden in caring for children when they are at home. To shed light on this aspect, we take advantage of the fact that five regions in Northern Italy closed schools well in advance of the national stay-at-home order.[6] Using mobility data from provinces in Northern Italy, Figure 4a presents an RD exercise with two discontinuities: the first is set on February 23rd, the day when local schools closed, and the second one on March 9th, when the national lockdown was implemented. The divergence in mobility between men and women starts already as soon as schools were closed. Men's mobility declines very marginally when schools were closed, while women's mobility has a clear discontinuity. This corroborates the hypothesis

[5]Of course, the number of contagions surged in the weeks leading to the stay-at-home orders, which were given because of this surge. However, the identifying assumption is that there was no discontinuity in the number of cases in the day of the orders.

[6]Schools in Northern Italy closed on February 23rd. On March 4th, the Italian government imposed the shutdown of all schools and universities nationwide. The national stay-at-home order was announced with a Presidential Decree on March 9th.

Figure 3: Impact of Stay-at-Home Orders on Mobility, by Gender
(Age group 25-44, percent of people leaving home in a day)

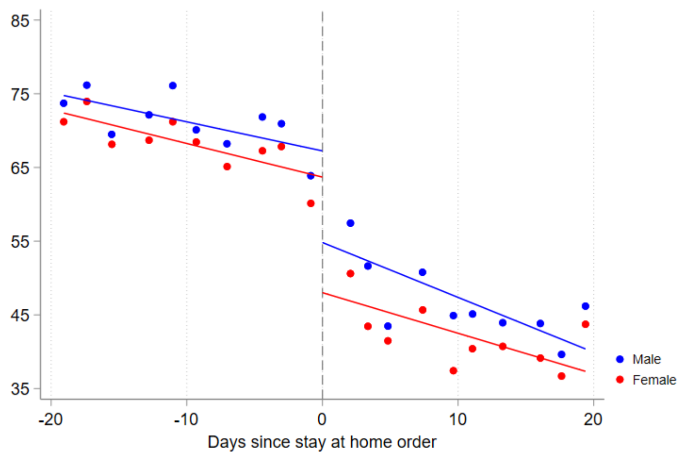

Notes: Percent of people moving is divided into 20 equal-sized bins. The series are residualized with respect to province and day-of-the-week fixed effects.

that women carry uneven responsibilities in looking after the children when schools were closed. The gender gap widens further at the time of the stay-at home order, even though the effect appears to be short-lived.

Figure 4b provides additional evidence about the role of school closures by examining the mobility gender gap across all provinces in Northern Italy. The heat map reports the difference between the mobility of men and women through time, with darker colors representing lower women's mobility. The mobility of men and women was similar before February 23rd. When schools closed, a disproportionate reduction in women's mobility occurred as clear from the darker colors across all provinces in the heat map. The adoption of stay-at-home orders led to further widening of the mobility gap. [7]

[7]Figure 4b provides additional information. The provinces in figure 4b are listed in order of decreasing frequency of COVID-19 cases as a share of the province population. The absence of a clear vertical pattern indicates that the effects of school closing and lockdowns on the difference in mobility was not correlated with the local intensity of the COVID-19 crisis.

Figure 4: Impact of school closures and stay-at-home orders in Northern Italy
(Age group 25-44, percent of people living home in a day)

(a) Impact on average mobility

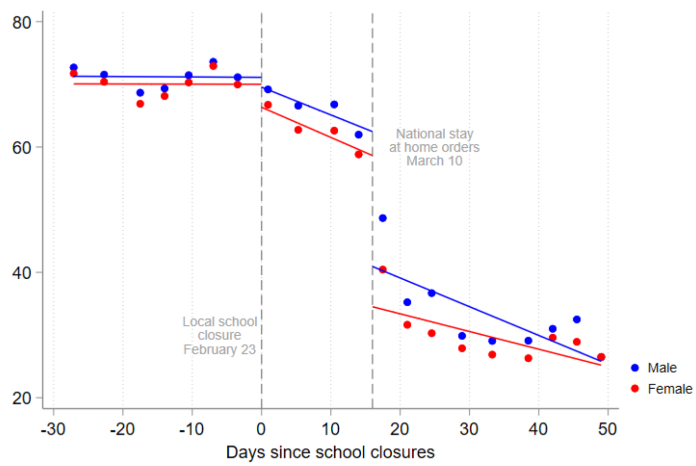

(b) Impact across all provinces

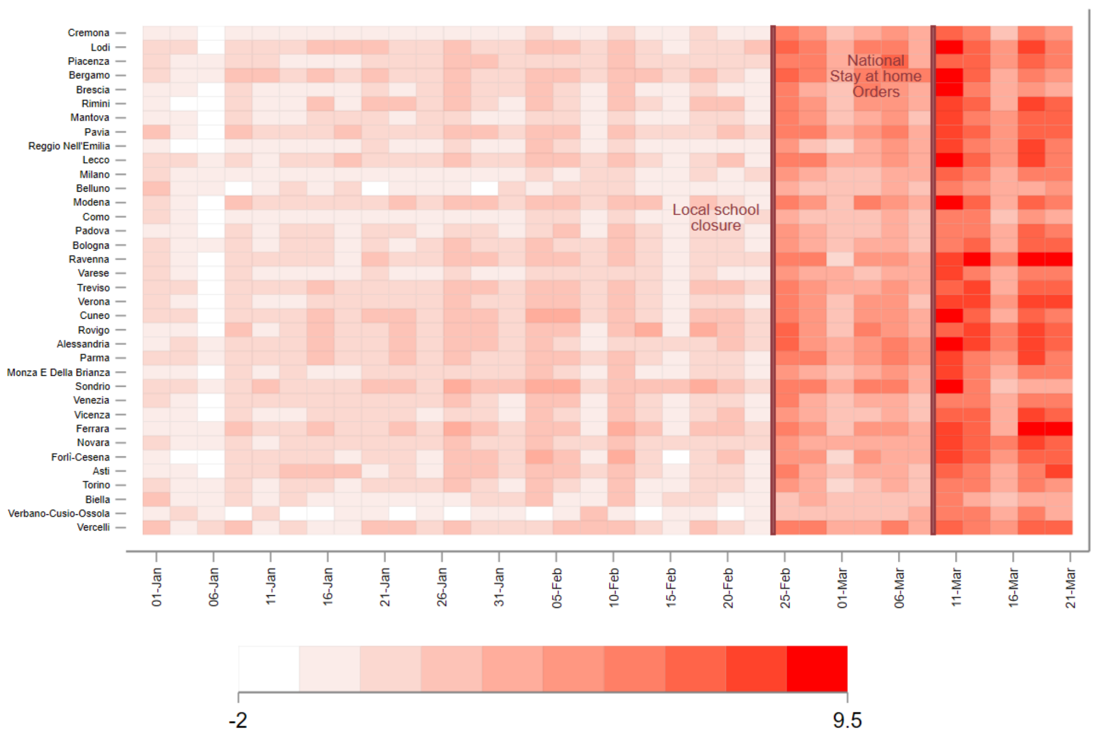

Notes: Panel (a) reports the difference between men and women mobility in Friuli-Venezia-Giulia, Emilia-Romagna, Lombardia, Piemonte, and Veneto. The timeline of different lockdown measures is as follows: local school closures were introduced on February 23rd, national stay-at-home orders on March 9th, further restrictions on non-essential production activities on March 21st. Selected stores reopened on April 14th and national stay-at-home orders were lifted on May 4th. Panel (b) reports a binned scatter plot where percent of people moving is divided into 20 equal-sized bins. The series are residualized with respect to province and day-of-the-week fixed effects.

The interpretation that school closures impacted disproportionately women's mobility is further confirmed by comparing the drop in mobility across age groups. In fact the analysis presented so far focuses on people aged between 25 and 44. These cohorts are more likely to have young kids and so they are most probably to have to look after the children when schools are closed.[8] A natural test of our conjecture is to test if the mobility gap was different for other age ranges.

To analyze the impact of lockdowns on the gender gap across all age groups and to test for the statistical significance of the discontinuity, we formalize the RD graphical evidence following Anderson (2014) and estimate this local linear regression:

$$
\begin{aligned}
mob_{i,g,a,t} &= \alpha_p + \tau_{dow} + \beta stay_{j,t} + \phi women_i + \gamma date_{j,t} \\
&+ \theta stay_{j,t} \times date_{j,t} + \lambda women_{i,a} \times stay_{j,t} + \nu date_{j,t} \times women_i + \varepsilon_{i,g,a,t} \qquad (1)
\end{aligned}
$$

where $mob_{i,g,a,t}$ is the mobility indicator provided by Vodafone capturing the percentage of people moving in province i, of gender g and age group $a = \{[18,24];[25,44];[45,64];[65+]\}$, at time t; $stay_{j,t}$ is the treatment variable for country j (with $i \in j$), equal to one when the national stay-at-home orders are in place; $women_{i,a}$ is a dummy variable equal to one when the dependent variable refers to the mobility of women; $date_{j,t}$ is the number of days since the introduction of the stay-at-home order; and α_p and τ_{dow} are province and day-of-the-week fixed effects. The coefficient β captures the effect of the stay-at-home orders on men's mobility, while $\lambda + \beta$ traces the effect on women's mobility. Standard errors are clustered at the province level. The identification assumption is that the term $stay_{j,t} \times date_t$ should absorb any smooth relationship between the $date_{j,t}$ and the error term $\varepsilon_{i,g,a,t}$ in the days around the introduction of the lockdown (Anderson (2014)). This means that no other factor affecting mobility should change close to national stay-at-home orders. Consistent with Figure 3, we estimate equation (1) using a relatively narrow window of 20 days around the adoption of stay-at-home orders since our identification strategy aims at estimating β and $\lambda + \beta$ by considering the mobility changes close to the introduction of lockdowns.

Table 1 reports the results for the baseline model. Column (2) shows that the mobility of women aged 25–44 declined by 3 percentage points more than men's. Column (1) and (3) present the results of the same specification for the age groups 18–24 and 45–64: the gap between

[8] In Italy, as in other countries, grandparents often play a big role in taking care of the kids while parents work. However, because COVID-19 affects disproportionately old people, social contacts between old and young people were discouraged. Therefore, the traditional arrangement was possibly less used, magnifying the effect on parents' mobility.

women and men mobility is still present, statistically significant, but smaller, equal 2.3 and 1.7 percentage points respectively. Finally, Column (4) shows that lockdowns no longer have a disproportionate effect on the mobility of women in the age group 65+.

Table 1: RD Estimate of the Gender Gap by Age Group

	18–24 (1)	25-44 (2)	45–64 (3)	65+ (4)
Stay-at-home	-19.60***	-12.95***	-12.52***	-11.76***
	(0.59)	(0.46)	(0.41)	(0.46)
Women × stay at home	-2.31***	-3.24***	-1.74***	1.55***
	(0.50)	(0.35)	(0.34)	(0.57)
Observations	13,909	14,102	14,151	13,102
R-squared	0.87	0.86	0.87	0.82

Notes: The table reports the coefficient on the stay-at-home variable and the coefficient on the interaction between the gender dummy and the stay at home variable. All regressions include the gender dummy, a variable for the number of days since the introduction of the stay-at-home order, the interaction terms of the latter with the stay-at-home variable and with the gender dummy, and province and day-of-the-week fixed effects. Standard errors are clustered at the province level. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

4.2 Local projections

The regression discontinuity approach used in the previous section shows that the adoption of specific lockdown measures reduced the mobility of women more forcefully. We now check if similar results hold when using local projections that exploit the entire variation in the stringency of lockdowns over the period of analysis.

To alleviate endogeneity concerns about lockdowns—namely that they are more likely imposed when the epidemic is worsening and mobility is too high—the local projections control for the number of COVID-19 infections and for lagged mobility levels to capture pre-existing trends. Furthermore, we rely on an identification strategy that takes advantage of the NUTS-3 level disaggregation of the Vodafone mobility data. The idea is that countries generally imposed lockdowns on a *national* scale in reaction to *localized* outbreaks. For example, in Italy—one of the first countries severely hit by the pandemic after China—the government imposed a national lockdown in early March even though most of the infections were concentrated in a few provinces in Lombardy. Therefore, the adoption of national lockdowns was largely exogenous to the conditions prevailing in provinces with relatively low infections. We leverage on this fact and exclude from the regression sample of each country the provinces that (i) registered the first 100 (cumulative) cases, (i) with the largest number of COVID-19 cases by the end of June 2020, and (iii) that had more than five percent of the country's total confirmed cases when the lockdown stringency index reached its maximum.[9] The regression thus retrieves the mobility response in those regions less affected by the virus for which the national lockdown was an exogenous event triggered by conditions elsewhere in the country.

Formally, to assess the differential impact of lockdowns on women's mobility, we estimate the following local projection regressions (Jordà, 2005) using data for a particular age group:

$$mob_{i,g,a,t+h} = \alpha_i^h + \kappa_g^h + \tau_t^h + \sum_{p=1}^{P} \rho_p^h mob_{i,g,a,t-p} + \sum_{p=0}^{P} \delta_p^h lock_{j,t-p} + \sum_{p=0}^{P} \beta_p^h ln\Delta cases_{i,t-p}$$
$$+ women_{i,a} \times \Big(\sum_{p=0}^{P} \gamma_p^h lock_{j,t-p} + \sum_{p=0}^{P} \psi_p^h ln\Delta cases_{i,t-p} \Big) + \varepsilon_{i,g,a,t+h} \quad (2)$$

where variable $mob_{i,g,a,t+h}$ denotes the percent of people moving in province i, of gender g and age a, at time $t+h$, with h being the horizon; $ln\Delta cases_{i,t-p}$ is the log of daily COVID-19 cases, which is used to track the stage of the pandemic, with p being the lag length (set to a week to control for the persistence of the variable); and $lock_{j,t-p}$ is an index measuring the stringency of lockdowns for country j (with $i \in j$), which also enters the specification with p lags to account for its persistence.[10] The specification also features lags of the dependent variable to account for pre-existing trends; province and gender fixed effects to control for time-invariant characteristics specific to provinces, men, and women; and time-fixed effects to control for those factors that are common to all provinces. Standard errors are clustered at the province level.

To uncover the differential impact of lockdowns on women, we include an interaction term between the lockdown stringency index and a gender dummy $women_{i,a}$, which is equal to one

[9] These criteria lead to the exclusion of Bergamo, Brescia, Lodi, Milan, Torino, and Rome in Italy; Barcelona and Madrid in Spain; and Área Metropolitana do Lisboa, Área Metropolitana do Porto, Cávado, and Região de Aveiro, Tâmega e Sousa in Portugal. Adding these areas back into the sample does not affect the results.

[10] We use the lockdown stringency index provided by the University of Oxford's Coronavirus Government Response Tracker. This index is a simple average of nine sub-indicators capturing school closures, workplace closures, cancellations of public events, gatherings restrictions, public transportation closures, stay-at-home requirements, restrictions on internal movement, controls on international traveling, and public information campaigns. Since we want to measure the impact of *actual* restrictions, we re-construct the index excluding public information campaigns as they aim to promote voluntary social distancing. The results, however, are similar when public information campaigns are included in the index.

when the dependent variable refers to the mobility of women. Thus, the coefficient δ_0^h isolates the impact of lockdowns on men's mobility and $\delta_0^h + \gamma_0^h$ the one on women's mobility. The regressions are estimated on a sample of 163 provinces in Italy, Spain, and Portugal between January 1 and June 29, 2020.

Figure 5 shows the impact of a full lockdown that includes all measures used during the pandemic on the mobility of men and women aged 24 to 45. The responses in panel 5a show that a full lockdown leads to a very significant decline in mobility for both men and women. Mobility starts to decline in the first two weeks to then sharply contracts in the third week. These dynamics follow the underlying pattern of the lockdown stringency that, as shown in Figure A.1 of Appendix A, declines gradually in the first two weeks and then temporarily increases in the third week.

Most importantly, panel 5a reveals that lockdowns have an uneven effect on mobility across gender, impacting women more strongly. Women's mobility falls by 28 percent three weeks after the introduction of lockdowns, while that of men declines by about 21 percent. Panel 5b shows that the differential between the mobility of women and that of men is statistically significant throughout the projection horizon.

Figure 5: Impact of a Full Lockdown on Mobility, by Gender
(Age group 25-44, percent)

(a) Men vs Women (b) Women-men differential

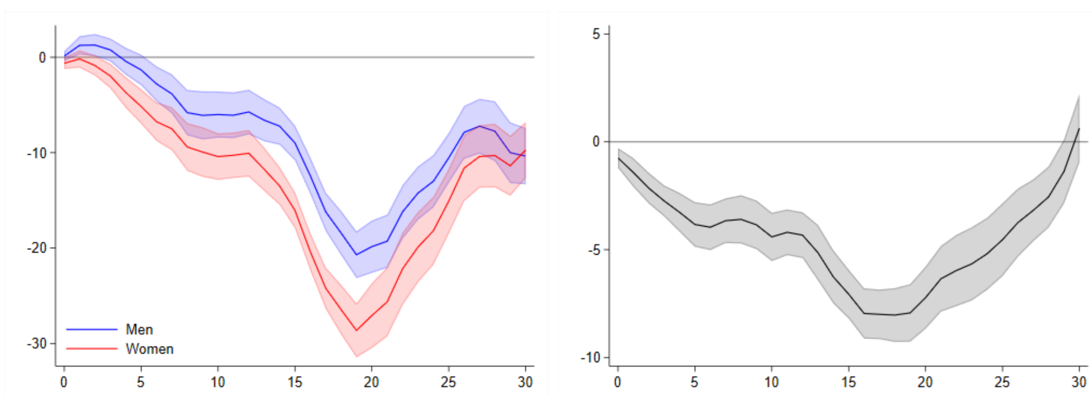

Notes: The x-axes denote the number of days, the lines denote the point estimates, and the shaded areas correspond to 90 percent confidence intervals computed with standard errors clustered at the province level.

These results corroborate the findings from the regression discontinuity analysis showing that lockdowns tend to disproportionately impact women. The regression discontinuity analysis also suggests that the gender differential is the largest for people aged 25–44, probably because they are more likely to have young children that have to be surprised at home when schools are closed. To check for the robustness of these findings, we re-estimate the local projections for the age groups 18–24, 45–64, and 65+. Table 2 reports the estimated mobility gap between women and men at the trough of the estimated response over the 30 days following the introduction of a full lockdown. We confirm that women's mobility falls the most relative to men's for those aged 25–44, with a differential of 7.3 percentage points. The gender gap declines to 6.6 percent for people aged 18–24 and to 2.7 percent for those aged 45–64. For people in the age group 65+, the fall in women's mobility is statistically indistinguishable from the men's one.

Table 2: Gender Gap at the Trough of the Estimated Response

	18–24 (1)	25-44 (2)	45–64 (3)	65+ (4)
Lockdown stringency	4.10***	-3.78***	-0.50	-0.11
	(0.77)	(0.61)	(0.59)	(0.74)
Women × lockdown stringency	-6.48***	-7.96***	-3.72***	-0.39
	(0.77)	(0.59)	(0.71)	(1.13)
Days after the shock	21	16	17	12
Observations	11,712	11,950	11,912	11,193
Provinces	163	163	163	157
R-squared	0.96	0.96	0.95	0.90

Notes: The table reports the coefficient on lockdown stringency and the coefficient on the interaction term between the gender dummy and lockdown stringency at the trough of the estimated response. All regressions include the contemporaneous value and/or seven lags of the stringency index, the log of daily cases, the interaction between a gender dummy and the stringency index, and province, gender, and time fixed effects. Standard errors are clustered at the province level. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Besides capturing the impact of lockdowns, the local projection in equation (3) also measures how mobility responds to an increase in COVID-19 infections. This is an important issue because during the pandemic people have voluntarily reduced exposure to each other as they feared contracting the virus. For example, Aum et al. (2020), Goolsbee and Syverson (2020), and Maloney and Taskin (2020) document that mobility has been tightly correlated to the spread of COVID-19 even after controlling for government lockdowns. In line with this literature, the specification in equation (2) sheds light on the strength of voluntary social distancing by capturing the response of mobility to rising COVID-19 infections for a given lockdown stringency.[11] The interaction term between daily COVID-19 infections and the gender dummy also reveals if the extent of voluntary social distancing differs between men and women. Specifically, the coefficient β_0^h measures the extent of voluntary social distancing for men, while the coefficients $\beta_0^h + \psi_0^h$ reflect the response of women.

Figure 6 shows how mobility responds to rising COVID-19 infections for a given lockdown stringency. An increase in COVID-19 cases has a negative effect on mobility of both men and women. A doubling of daily COVID-19 cases leads to a contraction in mobility by about 0.3 percent two weeks after the introduction of the lockdowns. Panel 6b indicates that the impact of COVID-19 cases on mobility is not statistically different between men and women.

Figure 6: Impact of a Doubling of COVID-19 Cases on Mobility, by Gender
(Age group 25-44, percent)

(a) Men vs Women — (b) Women-men differential

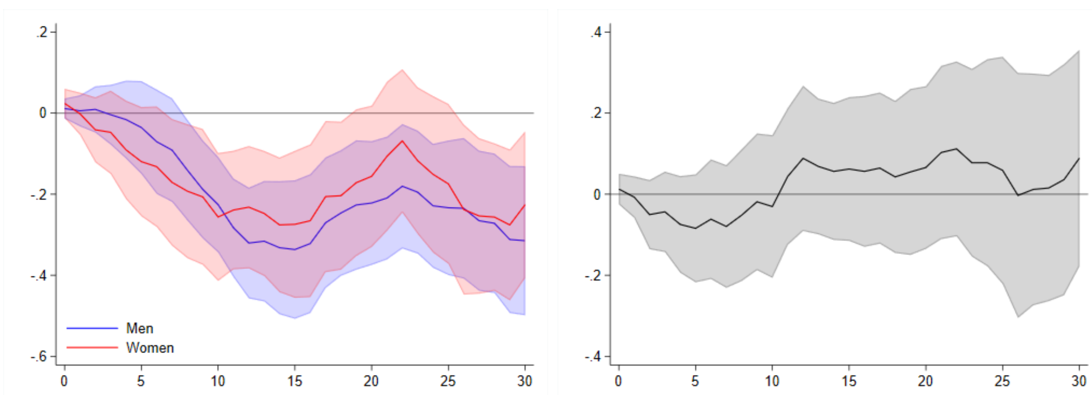

Notes: The x-axes denote the number of days, the lines denote the point estimates, and the shaded areas correspond to 90 percent confidence intervals computed with standard errors clustered at the province level.

This result rules out that the difference in the mobility across gender after lockdown is due to a fear factor—measured as the increase in the COVID-19 cases. If anything the fear factor should be more relevant for men, for whom COVID-19 is more dangerous.

[11] Besides reacting to the spread of COVID-19, people may opt to voluntarily self distance also in response to other factors, such as public health announcements, news about celebrities being infected, or even the adoption of government lockdowns. As such, the analysis may underestimate the true extent of voluntary social distancing. Also, as shown by Adda (2016), higher mobility and economic activity might lead to faster spread of viral diseases, generating some reverse causality between the outcome variables and COVID-19 infections. The dynamic structure of the estimation should alleviate this endogeneity concern.

5 Heterogeneous effects on mobility across age groups

In this section, we examine if lockdowns have a different impact on mobility depending on people's age. In line with the analysis on the effects across gender, we first examine the data using a regression discontinuity approach. We then revisit the evidence using local projections which also allow us to examine if people respond differently to rising infections depending on their age.

5.1 Regression discontinuity

We study the impact of stay-at-home orders on different age groups using the RD framework described in section 4.1. Panel 7a shows graphical evidence of the impact of stay-at-home orders on the mobility of each age group. Each dot captures the average mobility of both women and men in a given age group from 20 days before to 20 days after the adoption of stay-at-home orders. We see that lockdowns drastically reduced people's mobility across all age groups.

Yet, the mobility drop was significantly stronger for younger cohorts. This is clearly illustrated in panel 7b which reports the estimated drops from an RD specification akin to 1 and the associated 90 percent confidence intervals. Because of stay-at-home orders, the mobility of people below 25 years of age declined by more than 20 percentage points. The mobility drop becomes progressively smaller for older people, being equal to only 11 percent for people aged 65 and above (see also Table 4).

These patterns clearly show how lockdowns tend to disproportionally impact the mobility of younger cohorts. This is not surprising once considering that the mobility of people aged 65 and above—most of whom are retirees—was already significantly lower prior to lockdowns, as illustrated in panel 7a. Younger people have instead to leave their homes on a daily basis to join the work places and bring children to schools. Therefore, they are much more affected by lockdown measures that impede movements with adverse effects on their employment opportunities. The largest response of the younger generations could also reflect the different ways in which different generations receive information (and so perception of the risk).

5.2 Local projections

We now assess the impact of lockdowns across age groups using local projections based on the same identification strategy discussed in Section 4.2. Thus, the local projections control for the number of COVID-19 infections, lagged mobility levels, and are estimated over provinces

Figure 7: Impact of a Stay-at-home Orders on mobility, by Age
(Percent of people leaving home in a day)

(a) Regression discontinuity

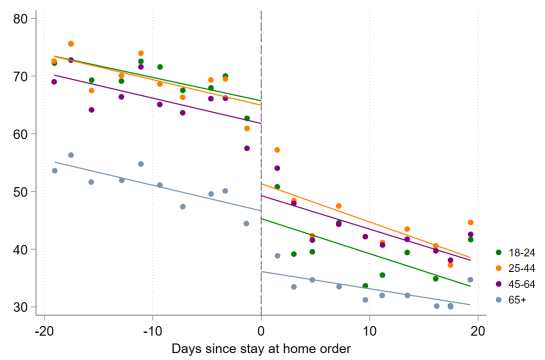

(b) Estimated impacts

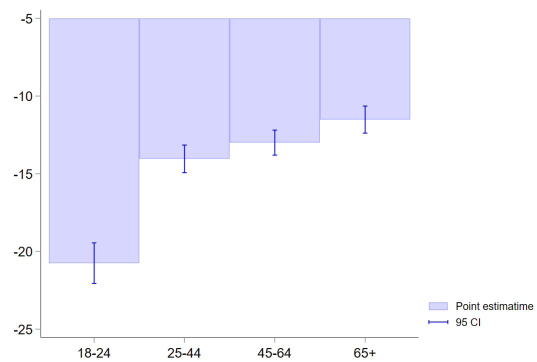

Notes: Panel 7a presents a binned scatterplot where the people moving is divided in 20 equally sized bins. The series is residualized with respect to province, gender day-of-the-week fixed effects. Panel 7b reports the estimates of the percentage drop in people moving by age group. Standard errors are clustered at the province level. See also table 4.

for which the national lockdown was largely exogenous as they did not experience early and/or major outbreaks. Formally, we estimate the following specification:

$$mob_{i,g,a,t+h} = \alpha_i^h + \kappa_a^h + \tau_t^h + \sum_{p=1}^{P} \rho_p^h mob_{i,g,a,t-p} + \sum_{p=0}^{P} \delta_p^h lock_{j,t-p} + \sum_{p=0}^{P} \beta_p^h ln\Delta cases_{i,t-p}$$
$$+ \sum_{a=1}^{3} agegroup_i^{a,g} \times \left(\sum_{p=0}^{P} \gamma_p^{s,h} lock_{j,t-p} + \sum_{p=0}^{P} \psi_p^{s,h} ln\Delta cases_{i,t-p} \right) + \varepsilon_{i,g,a,t+h} \quad (3)$$

The specification features interaction terms between the lockdown stringency index and age group dummies $agegroup_{i,g}^a$, with $a = \{1 = [18,24]; 2 = [25,44]; 3 = [45,64]\}$ and where the excluded category is the age group 65+. Hence, the impact of lockdowns on the mobility of people aged 65+ at horizon h is captured by δ_0^h, while the impact on the other age groups a is given by $\delta_0^h + \gamma_0^{a,h}$. The specification also includes interaction terms between COVID-19 infections and age groups to test whether the strength of voluntary social distancing differs across age. For a given level of lockdown stringency, the impact of rising COVID-19 cases on the mobility of people aged 65+ is measured by β_0^h, and the one on other age groups by $\beta_0^h + \psi_0^{a,h}$.

Figure 8 shows the effects of lockdowns on the mobility of different age groups. As shown in panel 8a, mobility collapses across all age categories. For the first ten days after the introduction of a lockdown, the mobility decline is similar across all age groups, with the 65+ age group experiencing a slightly larger fall. Yet, the younger cohorts experience a considerably stronger drop in mobility thereafter, reaching a trough of 37 percent 20 days after the introduction of lockdowns for people aged 18-24.[12]

To illustrate more clearly the differences across age groups, panel 8b shows the mobility differential of each age group relative to people aged 65+. During the first two weeks, the mobility of people aged 65+ declines marginally more than for the other age cohorts. However, most of the differentiation in mobility materializes after a couple of weeks, when the drop in mobility for the age groups 18–24, 25-44, and 45–64 becomes considerably larger and statistically different relative to people 65+. These findings corroborate the results from the regression discontinuity analysis in the previous section, showing that lockdowns disproportionally impact the mobility of younger cohorts.

Figure 8: Impact of a Full Lockdown on Mobility, by Age Group
(Percent)

(a) Age groups | (b) Differential with 65+

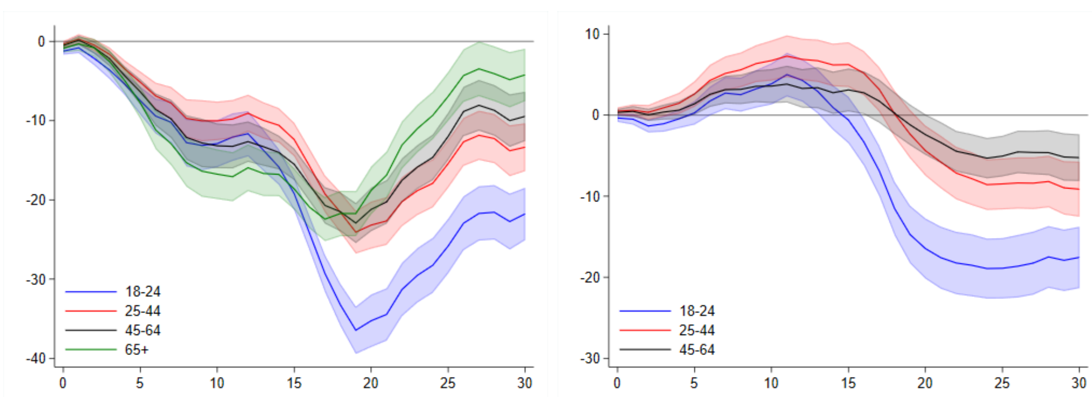

Notes: The x-axes denote the number of days, the lines denote the point estimates, and the shaded areas correspond to 90 percent confidence intervals computed with standard errors clustered at the province level.

The local projections also shed light on whether the extent of Voluntary social distancing differs across age groups. Panel 9a shows that the youngest cohorts are unequivocally the ones that react more forcefully to the rise in infections. A doubling of COVID-19 cases leads to a

[12]The mobility dynamics reflect the underlying stringency of lockdowns that as illustrated in figure of Appendix A weakens during the first 2 weeks and then temporarily strengthens in the third week.

fall in mobility by almost one percent for those aged 18–24 two weeks after the introduction of a lockdown. The size of the decline is not even half for the older age groups, reaching 0.4 for the 25–44 age group, and 0.2 for the 45–64 age group. In the case of the 65+, mobility remains broadly unchanged for about 20 days, and it increases somewhat after that. Panel 9b confirms that the declines observed for the three younger age groups are statistically different from the mobility dynamics observed for the oldest age group.

These results are somewhat surprising because people aged 65+ face much greater health risks from COVID-19 and should thus be more prone to isolate themselves when infections increase. Two considerations may explain our findings. First, the rise in infections reduces business activity in contact-intensive businesses, such as bars and restaurants, as people fear becoming infected. This in turn reduces employment in those sectors where many young people tend to work. Second, the stronger response in the mobility of younger people may reflect their stronger concerns about the virus. This is consistent with the evidence presented in Bordalo et al. (2020). Based on a a survey of 1,500 Americans in May 2020, they find that perceptions about the health risks posed by COVID-19 decline sharply with age. The fact that younger generations seem more sensitive to the fear factor—measured as doubling COVID-19 cases—could also reflect that younger generations use more media and social media which emphasize the danger.

Figure 9: Impact of a Doubling of COVID-19 Cases on Mobility, by Age Group (Percent)

(a) Age groups (b) Differential with respect to 65+

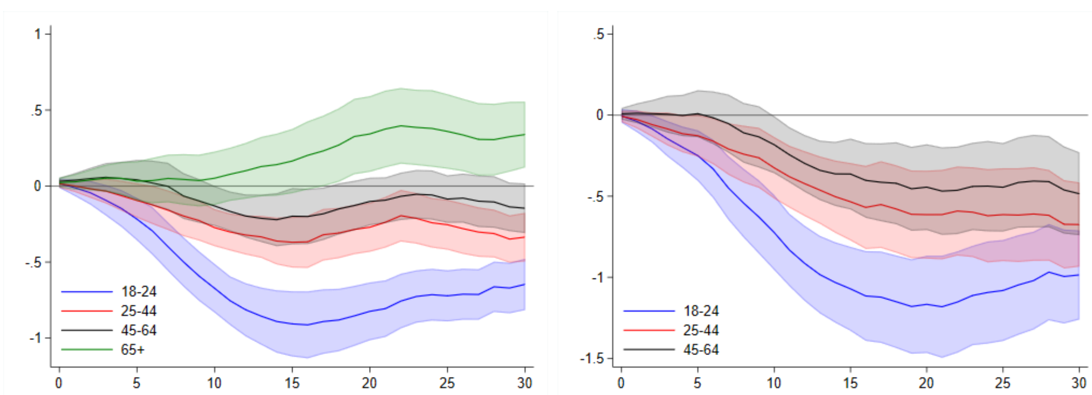

Notes: The x-axes denote the number of days, the lines denote the point estimates, and the shaded areas correspond to 90 percent confidence intervals computed with standard errors clustered at the province level.

6 Robustness

In this section we test the robustness of our results along several dimensions. We start by re-examining in Table 3 the evidence on the effect of lockdowns on mobility across gender for peopled aged 25-44 based on the regression discontinuity analysis. Column (1) reports the differential impact between women and men from the baseline analysis. It shows that lockdowns lead to a disproportionate decline in women's mobility. These results are based on a 20-day window before and after the adoption of stay-at-home orders. In column (2) we show that similar results are obtained if we shrink the regression window to 10 days to further limit possible bias from unobservable confounders. The results are also robust in column (3) to excluding the regions in Northern Italy that introduced lockdown measures, such as school closures, before the national stay-at-home order, as shown in Figure 4a and 4b. In columns (4) to (6) we include only two countries at the time in the estimation to check whether the results are driven by a particular country. The differential impact on women's mobility is confirmed on these different samples. The estimation excluding Italy however shows a smaller gap of less than 2 percentage points. Finally, in Column (7) we also control for the moving average of daily COVID-19 at the province level to reduce concerns of omitted confounders that may have a discontinuous effect on the mobility. The inclusion of this control, beyond province and day of the week fixed effects does not alter the results.

Following a similar set of robustness tests, Table 4 corroborates the regression discontinuity findings that lockdowns have a stronger impact on the mobility of younger cohorts.

Turning to the results of the local projections, as discussed in Section 4.2, the estimation is based on a sample that excludes provinces with early and/or large outbreaks. While this approach mitigates endogeneity concerns regarding the introduction of lockdowns, it may also affect our estimates if lockdowns or voluntary social distancing had different effects on mobility in regions more impact by the virus. Thus, we test if our findings are robust to the inclusion of those provinces. Figure 10 summarizes the results showing the coefficients on the interaction terms that capture the differential impact on the mobility across gender and age groups. Panels 10a and 10b confirm that lockdowns hit women's mobility disproportionately and that voluntary social distancing was broadly similar across gender. Panels 10c and 10d also corroborate the baseline results that lockdowns and voluntary social distancing take a larger toll on the youngsters. In terms of magnitudes, the estimated effects are virtually identical from those in the baseline.

Another concern is that the results may be specific to one of the three countries in the sample. Thus, we re-estimate the local projections excluding each of them at a time. Figure

Figure 10: Unrestricted Sample
(Percent)

(a) Impact of a full lockdown,
women-men differential

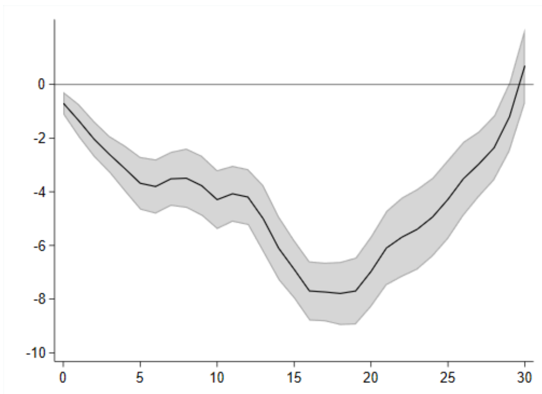

(b) Impact of a doubling of COVID-19 cases,
women-men differential

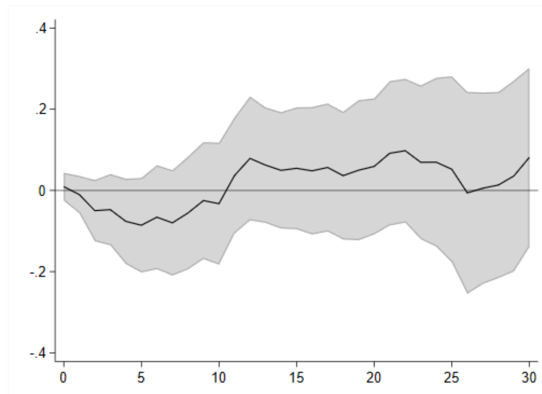

(c) Impact of a full lockdown,
differential with respect to 65+

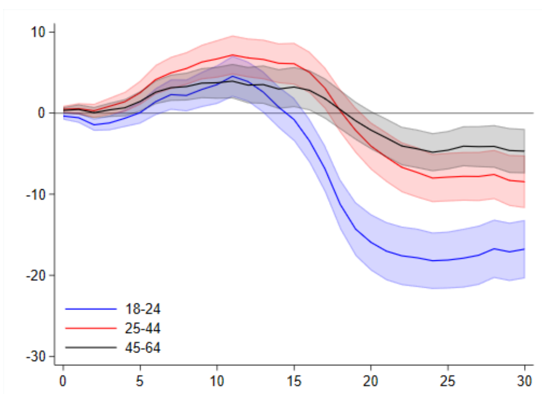

(d) Impact of a doubling of COVID-19 cases,
differential with respect to 65+

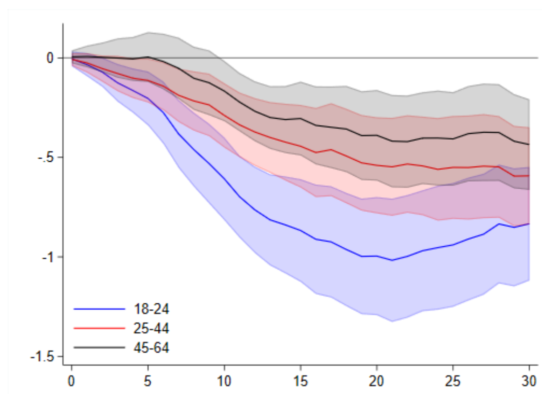

Notes: The x-axes denote the number of days, the lines denote the point estimates, and the shaded areas correspond to 90 percent confidence intervals computed with standard errors clustered at the province level.

Table 3: Robustness Exercises for Gender Gap with Regression Discontinuity

	Baseline (1)	10-day window (2)	No North. Italy (3)	Italy and Spain (4)	Italy and Portugal (5)	Portugal and Spain (6)	20-day window (7)
Stay at home	-12.95***	-8.99***	-12.82***	-14.42***	-13.04***	-10.71***	-11.69***
	(0.459)	(0.424)	(0.548)	(0.440)	(0.454)	(0.960)	(0.469)
Women × stay at home	-3.24***	-3.07***	-3.07***	-3.58***	-3.79***	-1.48*	-3.25***
	(0.353)	(0.431)	(0.442)	(0.403)	(0.149)	(0.869)	(0.352)
Observations	14,102	7,228	11,150	12,052	10,742	5,410	14,098
R-squared	0.86	0.85	0.85	0.84	0.92	0.69	0.86

Notes: The table reports the coefficients of an interaction term between the gender dummy and and stay-at-home variable. All regressions include the gender dummy, a variable for the number of days since the introduction of the stay-at-home order, the interaction terms of the latter with the stay-at-home variable and with the gender dummy, and province and day-of-the-week fixed effects. Column (7) considers a 20-day window and controls for the moving average of daily COVID-19 cases. Standard errors are clustered at the province level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

11 presents the results of the impact of lockdowns and voluntary social distancing on women's mobility compared to men's. Panels 11a, 11c, and 11e unambiguously confirm the larger effect of lockdowns on women's mobility compared to men's, which at the trough varies between 5 percent (when Italy is dropped) and 9 percent (when Portugal is dropped). Panels 11b and 11f show that the gender differential in the mobility response to rising COVID-19 cases remains statistically indistinguishable from zero. However, when Spain is excluded from the sample, women appear to curtail mobility more than men in the first week and by less subsequently, as shown in panel 11d. However, the point estimates are quantitatively very small.

Table 4: Robustness Exercises for Regression Discontinuity by Age Group

	Baseline (1)	10-day window (2)	No North. Italy (3)	Italy and Spain (4)	Italy and Portugal (5)	Portugal and Spain (6)	20-day window (7)
Stay at home	-11.51***	-9.11***	-11.02***	-12.35***	-12.71***	-8.08***	-10.49***
	(0.441)	(0.429)	(0.537)	(0.459)	(0.354)	(0.989)	(0.425)
18–24 × stay at home	-9.24***	-6.99***	-8.93***	-10.85***	-9.33***	-5.23***	-9.12***
	(0.505)	(0.461)	(0.630)	(0.468)	(0.464)	(1.092)	(0.502)
25–44 × stay at home	-2.53***	-0.75*	-3.03***	-3.13***	-1.35***	-3.58***	-2.45***
	(0.445)	(0.447)	(0.550)	(0.485)	(0.244)	(1.095)	(0.435)
45–64 × stay at home	-1.48***	0.24	-2.36***	-1.96***	0.71***	-4.81***	-1.42***
	(0.489)	(0.504)	(0.597)	(0.542)	(0.196)	(1.140)	(0.478)
Observations	55,264	28,258	43,456	47,672	42,360	20,496	55,248
R-squared	0.85	0.85	0.84	0.83	0.92	0.67	0.86

Notes: The table reports the coefficients of an interaction term between the age groups dummies and the stay-at-home variable. All regressions include the age group dummy, a variable for the number of days since the introduction of the stay-at-home order, the interaction terms of the latter with the stay-at-home variable and with the age group dummy, and province and day-of-the-week fixed effects. Column (7) considers a 20-day window and controls for the moving average of daily COVID-19 cases. Standard errors are clustered at the province level. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

Figure 11: Excluding One Country at a Time, Women-Men Differential
(Percent)

(a) Impact of a full lockdown,
excluding Italy

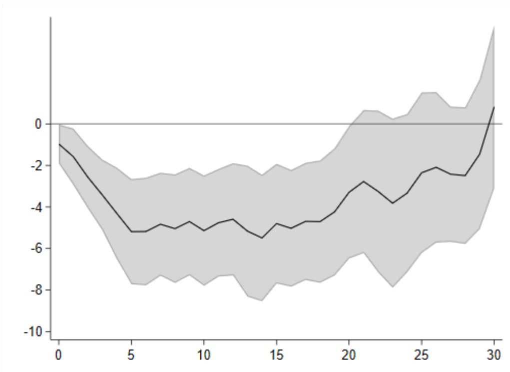

(b) Impact of a doubling of COVID-19
cases,
excluding Italy

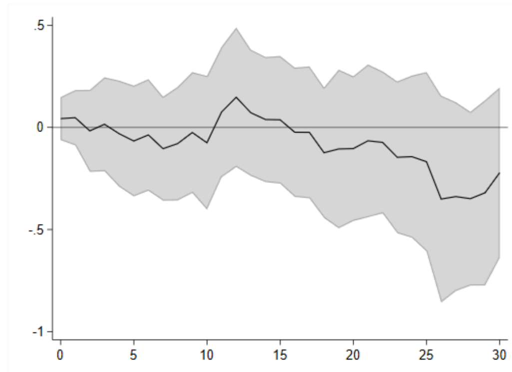

(c) Impact of a full lockdown,
excluding Spain

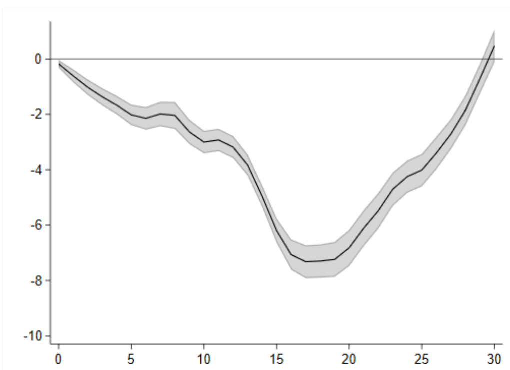

(d) Impact of a doubling of COVID-19
cases,
excluding Spain

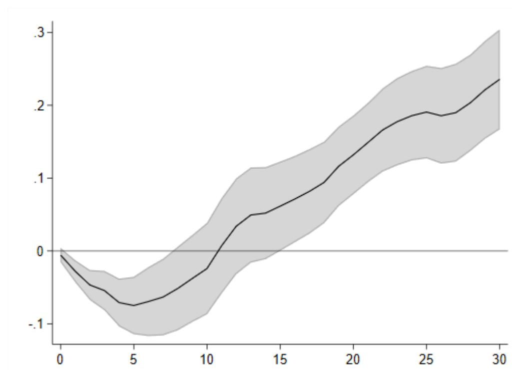

(e) Impact of a full lockdown,
excluding Portugal

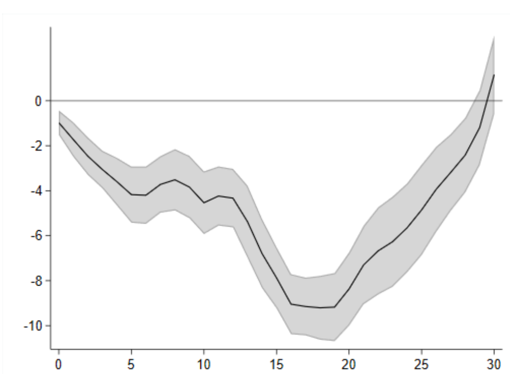

(f) Impact of a doubling of COVID-19 cases,
excluding Portugal

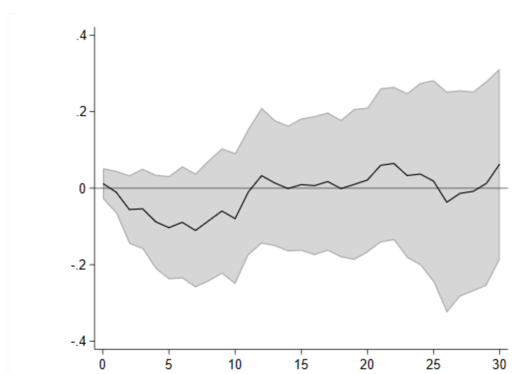

Notes: The x-axes denote the number of days, the lines denote the point estimates, and the shaded areas correspond to 90 percent confidence intervals computed with standard errors clustered at the province level.

Figure 12 shows the results of the impact of lockdowns and voluntary social distancing of different age groups compared to the age group 65+. Panels 12a, 12c, and 12e corroborate the baseline results for which lockdowns have a larger impact on the mobility of younger cohorts, regardless of which country is dropped. The declines are only marginally smaller when Portugal is excluded from the sample, as shown in panel 12e. We also confirm that the impact of voluntary social distancing on the mobility of younger cohorts is larger compared to the age group 65+, as shown in panels 12b, 12d, and 12f. Yet, panel 12f shows that magnitudes are generally larger when Portugal is not part of the sample. All in all, we find that the baseline results are generally robust to the exclusion of any country in the sample.

Figure 12: Excluding One Country at a Time, Differential with Respect to 65+ (Percent)

(a) Impact of a full lockdown, excluding Italy

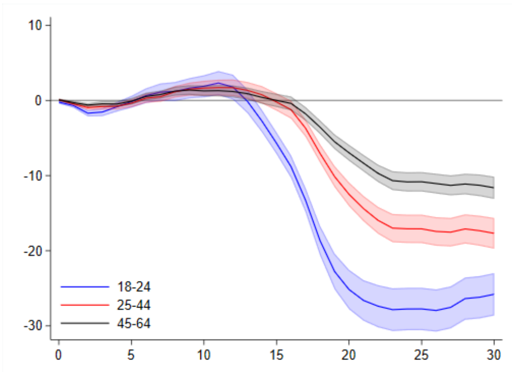

(b) Impact of a doubling of COVID-19 cases, excluding Italy

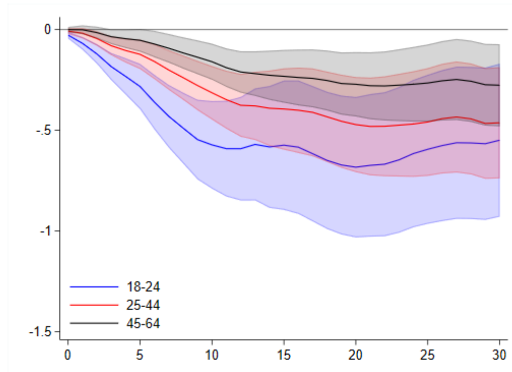

(c) Impact of a full lockdown, excluding Spain

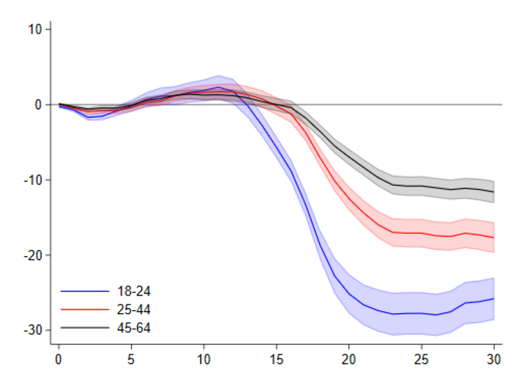

(d) Impact of a doubling of COVID-19 cases, excluding Spain

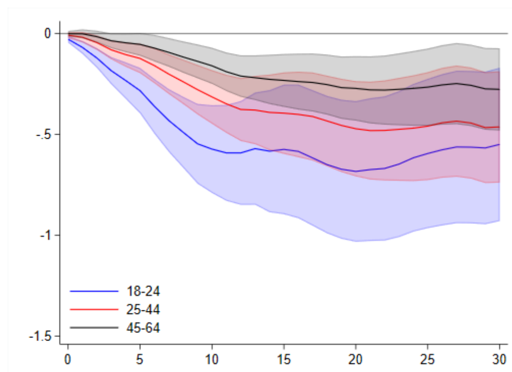

(e) Impact of a full lockdown, excluding Portugal

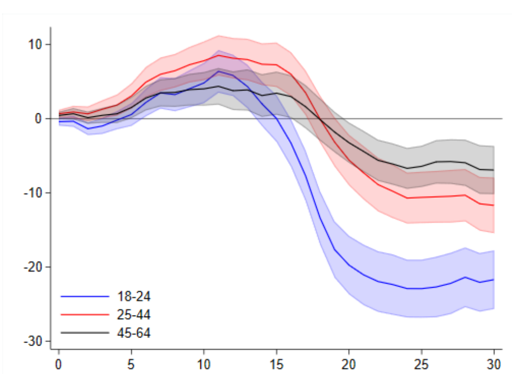

(f) Impact of a doubling of COVID-19 cases, excluding Portugal

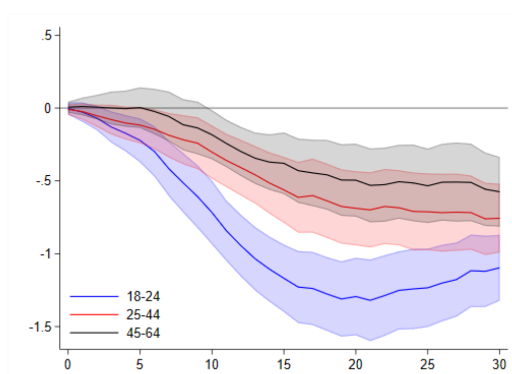

Notes: The x-axes denote the number of days, the lines denote the point estimates, and the shaded areas correspond to 90 percent confidence intervals computed with standard errors clustered at the province level.

We also examine if the results are robust to excluding time fixed effects which are used to capture movements in mobility that are common across provinces but are unrelated to the dynamics of lockdown stringency and COVID-19 infections.[13] One could argue that controlling for them in the local projections may saturate the specification given the high-frequency of the data. We thus replace time fixed effects with day-of-the-week fixed effects. Panels 13a and 13b of Figure 13 confirm the disproportional impact of lockdowns on women's mobility compared to men's, as well as that rising infection do not have a statistically significant different effect across gender. Without time fixed effects, panel 13c shows that, in response to a full lockdown, the mobility of the oldest age group appears to decline by more than for other age groups. However, such differential becomes statistically insignificant 22 days after the introduction of lockdowns and the youngest age group shows the largest mobility decline by the end of the projection horizon. Finally, panel 13d shows that the baseline findings about the impact of rising COVID-19 cases on the mobility of different age groups are robust to replacing time fixed effects with day-of-the-week fixed effects.

[13]Factors causing such movements could include, for example, public announcements of the government, public health officials, or international organizations, news about celebrities being infected, among others.

Figure 13: Replacing Time Fixed Effects with Day-of-the-Week Fixed Effects (Percent)

(a) Impact of a full lockdown, women-men differential

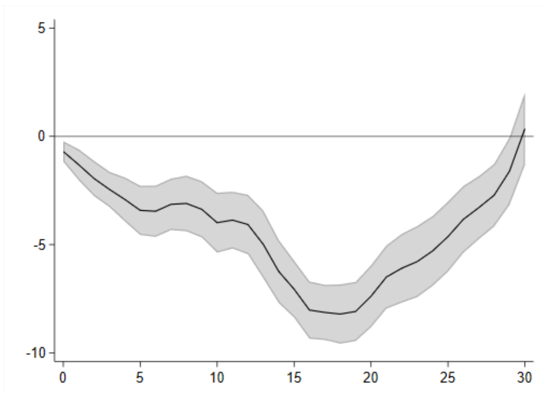

(b) Impact of a doubling of COVID-19 cases, women-men differential

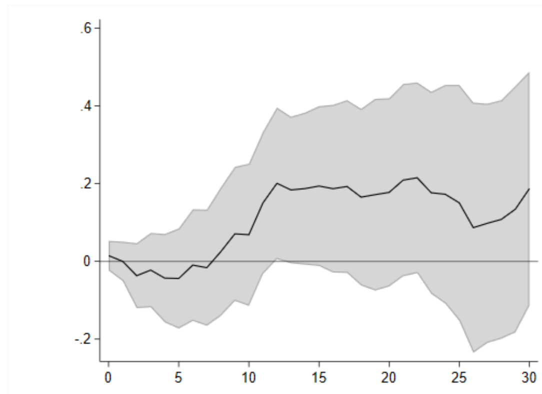

(c) Impact of a full lockdown, differential with respect to 65+

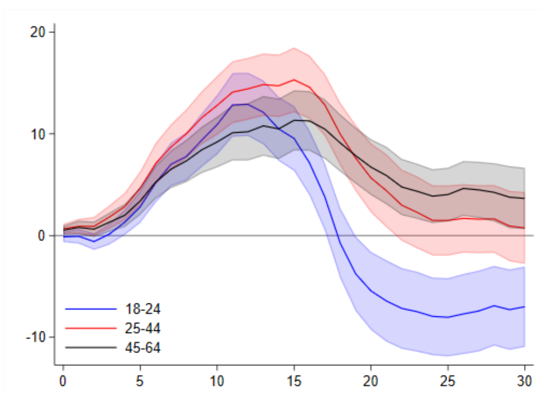

(d) Impact of a doubling of COVID-19 cases, differential with respect to 65+

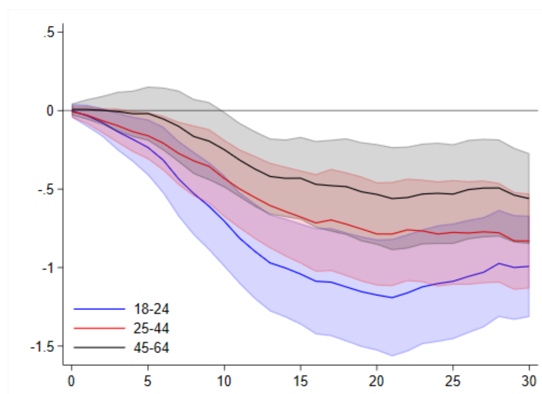

Notes: The x-axes denote the number of days, the lines denote the point estimates, and the shaded areas correspond to 90 percent confidence intervals computed with standard errors clustered at the province level.

7 Conclusions

In this paper, we have used unique mobility indicators provided by Vodafone which differentiate by gender and age to shed light on several key themes that have emerged during the COVID-19 pandemic.

First, the analysis provides novel evidence about the disproportionate impact of the crisis on women. Lockdown measures reduce the mobility of women more than men's. This seems largely due to women carrying an uneven burden in caring for children when schools are closed. Stay-at-home orders have indeed a disproportionate impact on women's mobility especially for people aged between 24 and 45 that are more likely to have young children. Furthermore, evidence from a few regions in Northern Italy that closed schools before adopting stay-at-home orders shows that the gender gap in mobility opens already at the time of school closures. In this respect, this paper contributes to the broader literature on the determinants of labor force participation. Previous studies have found that (exogenous changes in the) length of school schedule impact female labor force participation (Berthelon et al., 2015.) Our study provide complementary and high frequency evidence.

These findings warn about a possible widening of inequality across gender, as women may compromise their employment opportunities if they have to stay home to care for children. These concerns are further heightened by the fact that women tend to be employed in contact-intensive sectors—such hospitality, personal care and retail—that have been more severely impacted by the pandemic. Targeted policy intervention is required to support women during the pandemic, for example by offering parental leave to both men and women to encourage equal burden sharing in caring for children when schools are closed.

Second, the analysis contributes to the debate about the uneven effects of the crisis across age groups. By containing the spread of the virus, lockdowns benefit especially people above 65 years of age because they face much greater health risks from COVID-19. The economic costs of lockdowns fall instead disproportionately on working age people. The analysis shows that lockdowns lead to a stronger reduction in the mobility of younger people, for example preventing them from reaching their work places and bringing children to school to free up time for work.

Interestingly, the mobility of younger people responds more strongly also to rising infections, for a given level of the stringency of lockdowns. This could be because younger people are more concerned about the virus in line with survey evidence despite being less likely to develop severe health conditions. Or it may capture that rising infections reduce business activities in contact-intensive sectors, such as bars and restaurants, leaving many young people that work in those sectors unemployed.

The disproportionate impact of lockdowns on the mobility of the young is particularly concerning because younger workers depend on labor income to sustain consumption while older people have access to larger personal saving and often receive stable retirement income. Younger workers also have generally less stable job contracts that are more likely to be terminated during a crisis. These considerations highlight the need for a social path across generations to at least partially compensate younger workers for the economic losses they face because of lockdowns. This is essential not only from a fairness standpoint but also to ensure a sufficiently strong public support to deploy lockdown measures when needed.

Third, the results on the differential effects on age groups also provide insights on the possible long-term effects of the lockdown. The fact that younger generations reduced mobility more the older generations during the lockdown suggests that the scarring effects could be long term. This effect would be compounding the known effect that generations entering the labor force during a recession suffer a long-term scarring effect. This is a preliminary insight that should be investigated further in the future

Fourth, the fact that different demographic groups reacted differently to stay-at-home orders, school closure, and COVID-19 cases calls into question the assumption that population can be treated as homogeneous; future studies should investigate further the extend of a possible aggregation bias.

Finally, especially in the context of a resurgent epidemics in several regions of the world, policies should take into account this heterogeneity.

References

Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D. Whinston (2020). "Optimal Targeted Lockdowns in a Multi-Group SIR Model". NBER Working Paper No. 27102.

Adams, Abigail, Teodora Boneva, Marta Golin, and Christopher Rauh (2020). "Inequality in the Impact of the Coronavirus Shock: Evidence from Real Time Surveys". CEPR Discussion Paper 14665.

Adda, Jérôme (2016). "Economic Activity and the Spread of Viral Diseases: Evidence from High Frequency Data". *Quarterly Journal of Economics* 131(2), pp. 891–941.

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". CEPR Discussion Paper 15149.

Anderson, Michael L. (2014). "Subways, Strikes, and Slowdowns: The Impacts of Public Transit on Traffic Congestion". *American Economic Review* 104(9), pp. 2763–96.

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". NBER Working Paper No. 27264.

Belot, Michèle, Syngjoo Choi, Egon Tripodi, Eline van den Broek-Altenburg, Julian C. Jamison, and Nicholas W. Papageorge (2020). "Unequal Consequences of COVID-19 across Age and Income: Representative Evidence from Six Countries". IZA Discussion Paper No. 13366.

Berthelon, Matias, Diana Kruger, and Melanie Oyarzún (2015). "The Effects of Longer School Days on Mothers' Labor Force Participation". IZA Discussion Paper No. 9212.

Bordalo, Pedro, Katherine B. Coffman, Nicola Gennaioli, and Andrei Shleifer (2020). "Older People are Less Pessimistic about the Health Risks of Covid-19". NBER Working Paper No. 27494.

Chetty, Raj, John N. Friedman, Nathaniel Hendren, and Michael Stepner (2020). "How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data". NBER Working Paper No. 27431.

Davis, Lucas W. (2008). "The Effect of Driving Restrictions on Air Quality in Mexico City". *Journal of Political Economy* 116(1), pp. 38–81.

Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull (2020). "Health versus Wealth: On the Distributional Effects of Controlling a Pandemic". NBER Working Paper No. 27046.

Goolsbee, Austan and Chad Syverson (2020). "Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline". NBER Working Paper No. 27432.

Hupkau, Claudia and Barbara Petrangolo (2020). "Work, Care and Gender during the COVID-19 Crisis". CEPR Discussion Paper 15358.

Jordà, Òscar (2005). "Estimation and Inference of Impulse Responses by Local Projections". *American Economic Review* 95(1), pp. 161–182.

Kariotis, Timothy, Mad Price Ball, Bastian Greshake Tzovaras, Simon Dennis, Tony Sahama, Carolyn Johnston, Helen Almond, and Ann Borda (2020). "Emerging Health Data Platforms: From Individual Control to Collective Data Governance". *Data & Policy* 2, e13.

Maloney, William and Temel Taskin (2020). "Determinants of Social Distancing and Economic Activity During COVID-19: A Global View". The World Bank, Policy Research Working Paper 9242.

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". *NBER Working Paper 27132*.

Sevilla, Almudena and Sarah Smith (2020). "Baby Steps: The Gender Division of Childcare during the COVID19 Pandemic". Covid Economics 23.

Appendix A. Lockdown stringency dynamics

To better understand the dynamics uncovered by the local projections regarding how lockdowns affect mobility, it is helpful to examine how the stringency of lockdowns evolves over the local projection horizon. Panel A.1a shows that a lockdown tightening tends to gradually decline during the first two weeks to then experience a temporary renewed increase in the third week.

These estimated dynamics reflect the way in which Italy, Portugal, and Spain have adjusted their lockdown stringency during the sample of analysis. As illustrated in panel A.1b, both tightening and loosening phases have occurred in steps, giving rise to the estimated pattern. Countries have indeed progressively tightened lockdown restrictions between February and March and then progressively eased in May and June.

Figure A.1: Lockdown Stringency Dynamics

(a) Impact of a full lockdown on lockdown stringency (Percent)

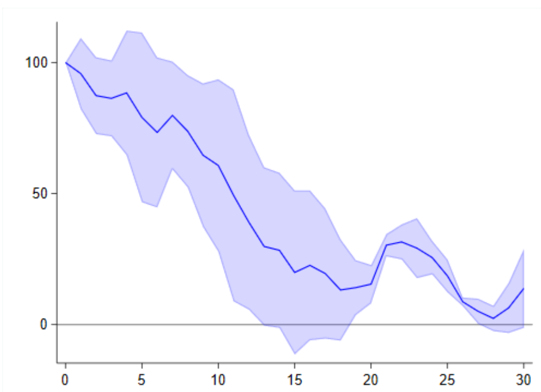

(b) Lockdown stringency (Index)

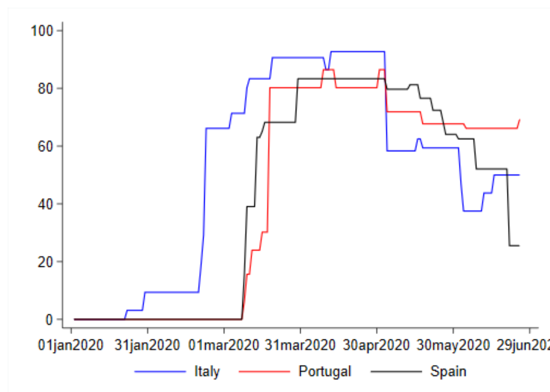

Notes: In panel A.1a, the x-axis denotes the number of days, the line denotes the point estimates, and the shaded area corresponds to the 90 percent confidence interval computed with standard errors clustered at the province level.