



WP/20/282

IMF Working Paper

COVID-19 Pandemic: Asymmetric Effects across Gender and Age

by Francesca Caselli, Francesco Grigoli,
Damiano Sandri, Antonio Spilimbergo

I N T E R N A T I O N A L M O N E T A R Y F U N D

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

December 21, 2020

Abstract

Overall mobility declined during the COVID-19 pandemic because of government lockdowns and voluntary social distancing. Yet, aggregate data mask important heterogeneous effects across segments of the population. Using unique mobility indicators based on anonymized and aggregate data provided by Vodafone for Italy, Portugal, and Spain, we find that lockdowns had a larger impact on the mobility of women and younger cohorts. Younger people also experienced a sharper drop in mobility in response to rising COVID-19 infections. Our findings, which are consistent across estimation methods and robust to a variety of tests, warn about a possible widening of 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. We thank Pedro Rente Loureço and Vodafone's Big Data and Artificial Intelligence team for constructing the mobility indicators used in this study.

Contents

	Page
1 Introduction	4
2 The COVID-19 crises in Italy, Portugal, and Spain	6
3 A unique dataset	7
4 Heterogeneous effects on mobility across gender	9
4.1 Regression discontinuity	10
4.2 Local projections	13
5 Heterogeneous effects on mobility across age groups	18
5.1 Regression discontinuity	18
5.2 Local projections	19
6 Robustness	21
7 Conclusions	28
References	30
A Lockdown stringency dynamics	32

List of Figures

	Page
1 Mobility Levels from Apple, Google, and Vodafone	8
2 Mobility and Lockdown Stringency	9
3 Impact of Stay-at-Home Orders on Mobility, by Gender	11
4 Impact of School Closures and Stay-at-Home Orders in Northern Italy	12
5 Impact of a Full Lockdown on Mobility, by Gender	16
6 Impact of a Doubling of COVID-19 Cases on Mobility, by Gender	17
7 Impact of a Stay-at-Home Orders on Mobility, by Age	18
8 Impact of a Full Lockdown on Mobility, by Age Group	20
9 Impact of a Doubling of COVID-19 Cases on Mobility, by Age Group	20
10 Impact of Lockdowns and COVID-19 Cases using Full Sample	24
11 Impact of Lockdowns and COVID-19 Cases using Day-of-the-Week Fixed Effects	25
12 Impact of Lockdowns and COVID-19 Cases Excluding Portugal	26
13 Impact of Lockdowns and COVID-19 Cases Excluding Spain	27
A.1 Lockdown Stringency Dynamics	32

List of Tables

	Page
1 RD Estimate of the Gender Gap by Age Group	14
2 Gender Gap at the Trough of the Estimated Response	16
3 Robustness of the RD Results by Gender	21
4 Robustness of the RD Results by Gender across Provinces	22
5 Robustness of the RD Results by Age	22
6 Robustness of the RD results by Age across Provinces	23

1 Introduction

The COVID-19 pandemic dramatically reduced people’s mobility. This was due in part to the lockdown measures that governments adopted to reduce personal contacts, 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 contracting 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 (Chetty et al., 2020; Glaeser et al., 2020; Goolsbee and Syverson, 2020; Maloney and Taskin, 2020).

This paper innovates relatively to the existing studies by showing that mobility patterns 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. These data offer the unique advantage of disaggregating mobility information across gender and age groups, making it possible to uncover important heterogeneous reactions to the pandemic and lockdown measures. This paper makes several contributions.

First, the analysis contributes to the growing evidence about the disproportionate impact of the COVID-19 crisis on women. This literature finds that some home production that in normal times can be outsourced had to be performed within the household during the pandemic. And this burden fell disproportionately on women. Hupkau and Petrangolo (2020), for example, find that in UK households women took on a larger share of increased childcare needs, even though fathers became the primary childcare providers in an important share of households. Alon et al. (2020) show that, contrary to past recessions, the current crisis led to a stronger increase in female 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). In this paper we document that the disproportional impact of the crisis on women is also visible in mobility patterns and examine the causes of this differential.

In this regard, this paper also contributes to the literature on the determinants of labor force participation. School closures and other lockdown measures reduce mobility in different ways across gender and thus have a differential effect on labor supply. Previous studies have found that exogenous changes in the length of school schedule impact female labor force participation (Berthelon et al., 2015). We provide high-frequency, complementary evidence.

Second, this paper documents considerable heterogeneity in the impact of lockdowns across age groups. These findings are quite relevant for the ongoing debate on the distribution of costs and benefits across generations (Belot et al., 2020; Glover et al., 2020a). While lockdowns protect mainly older people that are more likely to develop serious medical conditions from COVID-19, they impose economic costs especially on working age people that rely on labor income to support consumption.

Third, the heterogeneous impact of lockdowns across age can shed light on the scarring effects from the crisis. For instance, a given reduction in aggregate mobility has different long-term effects if it is concentrated on the young. Oreopoulos et al. (2012) find that individuals entering the labor force during a recession suffer from substantial and persistent loss of earning; they also document that graduates who can switch quickly to better firms suffer less. So a forced reduction in mobility has a particularly adverse and persistent effect if it reduces the mobility of young.

Fourth, heterogeneous effects across gender and age raise important methodological considerations. Lumping all groups together in estimating the effects of containment measures as done in most studies could lead to aggregation bias and mismeasurement problems.

Fifth, in presence of repeated infection waves in several countries (as we write, many countries are experiencing a strong resurgence) and in the context of an already weak economic activity, authorities need to consider more nuanced containment measures. To design these measures, authorities should consider the heterogeneous effects on different 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 on mobility is a difficult task because the adoption of lockdowns is an endogenous policy decision. For example, governments are more likely to impose measures when health risks are more acute. At that time people tend to also reduce mobility voluntarily because they fear being infected or infecting others. This raises the risk of detecting a spurious correlation between lockdowns and mobility. To address the endogeneity concern, we use a regression discontinuity (RD) design that focuses on high-frequency changes in mobility around specific lockdown measures, thereby reducing the risk that other factors may affect mobility at the same time. In addition, we estimate local projections in which we control for pre-existing mobility trends and for the severity of the country’s epidemic. We further strengthen the identification by excluding those provinces that experienced early or particularly severe outbreaks. In this case, the analysis looks at the mobility patterns in provinces less affected by COVID-19 for which the adoption of national lockdowns was largely an exogenous event.

The results from the RD design show that the adoption of stay-at-home orders in Italy, Portugal, and Spain led to a sharper contraction in women’s mobility relative to men’s. Stay-at-home orders coincided in most provinces with school closures. To isolate the impact of school closures on mobility, we narrow the analysis to provinces in Northern Italy that imposed school closures before stay-at-home orders. We find that school closures led already to a considerable gender differential in mobility, highlighting the uneven role of women in caring for children.

The findings based on local projections corroborate the results of the RD analysis, showing that lockdowns had a disproportionate effect on the mobility of women. The differential impact is statistically significant and quantitatively relevant. Among people aged 24 to 45, a full lockdown—including all measures used by governments during the pandemic—reduces the number of women leaving home by almost 26 percent, against an impact on men of about 21 percent.

The local projections also allow to examine the effects of rising COVID-19 cases on mobility for a given level of 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 very important because much of the public debate on the need for lockdowns centers 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.

Regarding the differential effects across age, RD designs show that stay-at-home orders disproportionately reduce the mobility of working age people, especially those below 45 years of age. Local projections provide additional evidence that lockdowns have a larger impact on younger cohorts. 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 based evidence confirms that younger people have been more likely to suffer an income loss during the pandemic (Belot et al., 2020; Montenovolo 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 inter-generational fairness (Glover et al., 2020b).

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 could be expected that these people are more likely to isolate themselves when infections rise. 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. Section 2 presents background information on the COVID-19 crises in Italy, Portugal, and Spain, and the relative containment measures. Section 3 describes the data provided by Vodafone. 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 evolution of the COVID-19 crises in Italy, Portugal, and Spain and the lockdown measures adopted at the local and national level.

Italy. Italy was among the first countries to be hit by COVID-19 after China. On January 31, the Italian government declared a state of emergency and stopped flights 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 were 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. 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, 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. Broader diffusion began by mid-February and all 50 provinces had confirmed cases by mid-March. 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 until June 1 which 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 municipalities in Catalunya with 70,000 people affected. On March 14, the entire country entered in the state of emergency and many nonessential activities were closed, such as restaurants and 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 anonymized and aggregated data on mobility provided by Vodafone through a confidential agreement. By analysing connections of mobile phones 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.¹ The age groups include four categories: people aged between 18 and 24, between 25 and 44, between 45 and 64, and 65 and above.² For Spain and Portugal, however, some daily observations are missing, especially for the oldest and youngest age groups. In these cases, we linearly interpolate the series to have a balanced panel throughout the sample period.

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 Lourenco et al. (2020).

The mobility patterns detected by Vodafone are broadly in line with those according to Apple and Google data.³ Figure 1 shows that all indicators correlate fairly closely at the national level. Correlations

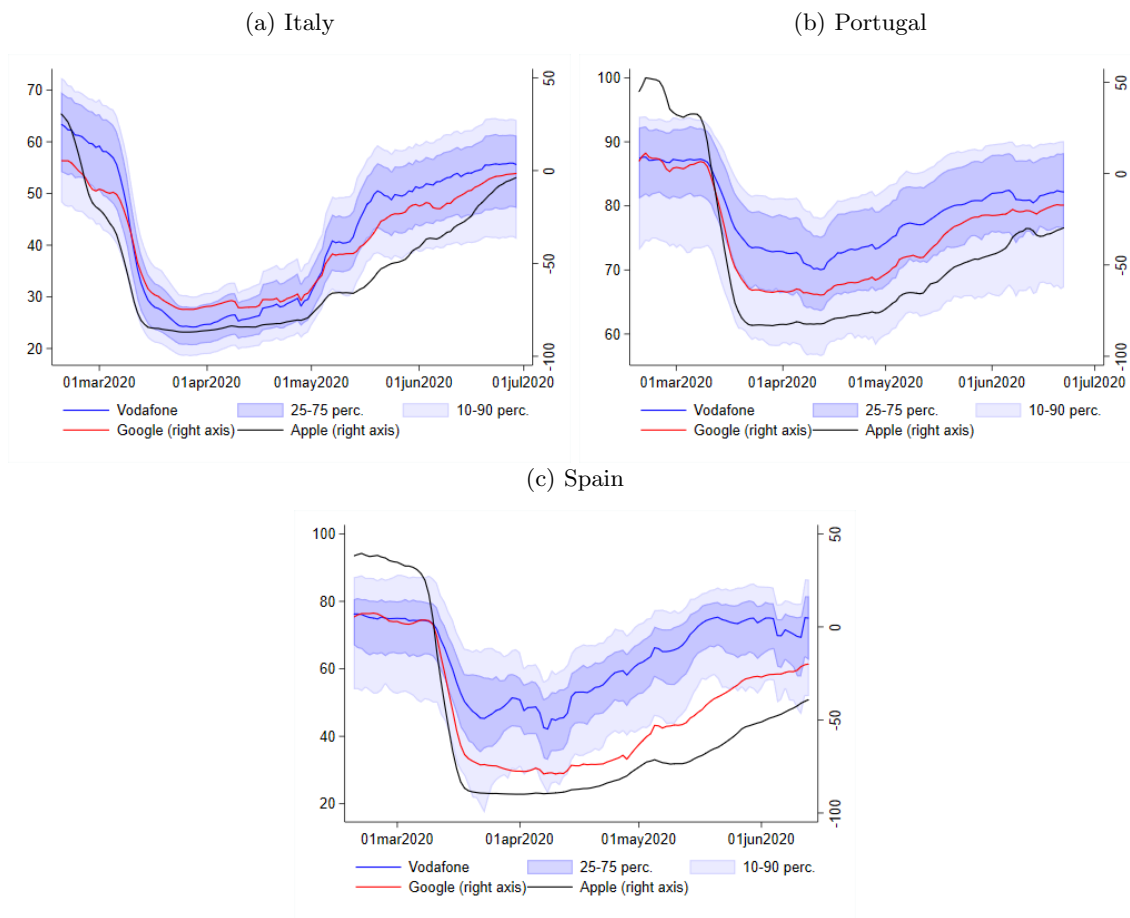
¹These indicators were prepared by Vodafone’s Big Data and Artificial Intelligence team. To protect the privacy of individuals and minority groups, the data have been provided in anonymized form, reporting the average mobility for a given gender and age group at the provincial (NUTS3) level when a minimum 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.

²In a few cases, the age information is inferred. For example, in Spain, the age group 18–24 is separated from family contracts based on several factors, including 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.

³Apple mobility data are available at <https://covid19.apple.com/mobility>, and Google mobility data are available at <https://www.google.com/covid19/mobility>.

between the Vodafone indicator and the Apple and Google indicators range between 93 and 99 percent for Italy and Portugal and are 72 and 88 percent for Spain. The geographical disaggregation of the 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.

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

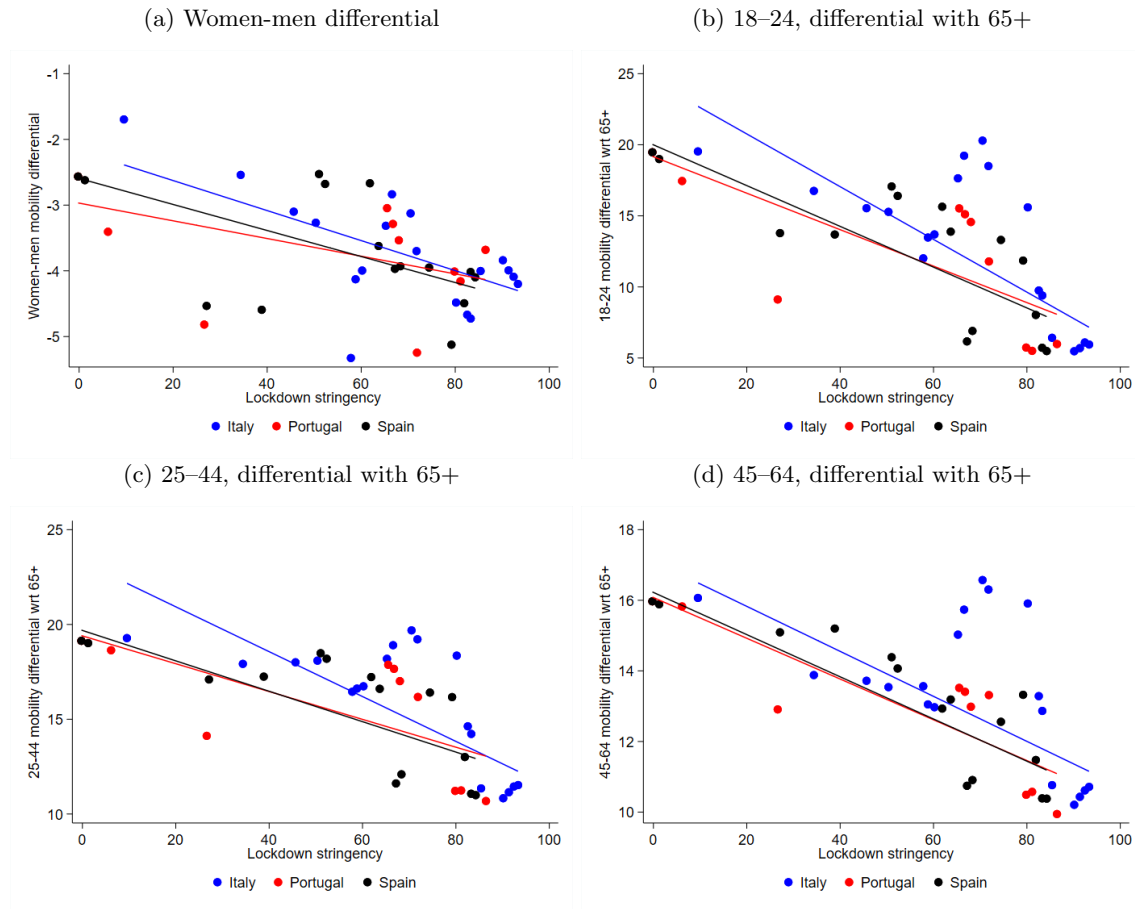


Notes: The lines denote the country-level mobility levels. In the case of Vodafone, the line corresponds to the cross-province population-weighted average of the percent of people moving, using 2018 population levels as weights; and the shaded areas denote the cross-province interquartile range (dark 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.

The key advantage provided by the Vodafone data is the ability to differentiate mobility across gender and age groups. This makes it possible to examine 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 the correlation between the stringency of lockdowns during the period of analysis and the average mobility differential between women and men. In all countries lockdowns have been associated with a larger drop in women’s mobility relative to men’s. The other three panels in Figure 2 show the correlation between the stringency of lockdowns and the mobility differential relative to the oldest age category of 65

and above. The charts suggest that lockdowns reduce more strongly the mobility of working age people relative to people aged 65 and above. The rest of the analysis will test more formally for these patterns using RD approaches and local projections.

Figure 2: Mobility and Lockdown Stringency



Notes: Panel (a) presents a binned scatter plot showing the mobility differential between women and men over the stringency of lockdowns during the period of analysis. Panel (b) to (d) present similar scatter plots but considering the differential in percentage points between people in different age groups relative to people aged 65 and above. Each dot denotes the cross-province average at any given time. The mobility indicator is residualized with respect to days of the week fixed effects. The lines denote the linear fit.

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 because 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. If voluntary social distancing is not controlled for, the empirical analysis would thus

overplay the impact of lockdowns on 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 RD 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 provinces 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 is also helpful to examine how people voluntarily reduce mobility in response to rising infections.

4.1 Regression discontinuity

To test whether lockdowns have unequal effects across gender, we first use an 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 a particular temporal threshold, making this approach akin to an event study exercise. As in more standard RD, 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.⁴

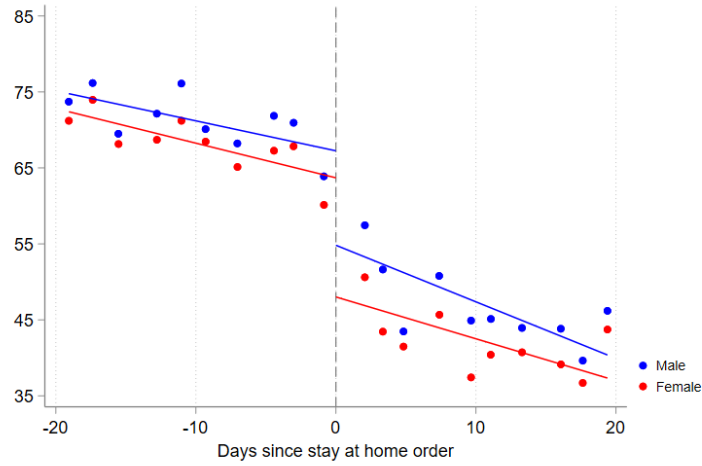
Figure 3 uses a bin scatter plot to present preliminary evidence that lockdowns are associated with a discontinuity in the mobility of women relative to men. Each dot represents the average mobility levels of men and women calculated using 20 equally sized bins around the introduction of national stay-at-home orders. We start by considering people aged between 25–44 for whom the differential effect of lockdowns across gender is the greatest, as we will later show. 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 points. Yet, the impact on women was stronger, as their mobility declined by about 3 percentage points more than for men.

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.⁵ Using mobility data from provinces in Northern Italy, Figure 4a presents an RD exercise with two discontinuities: the first is set on February 23, the day when local schools closed, and the second one on March 9, when the national lockdown was implemented. The divergence in mobility between men and women started already at the time of school closures. Men’s mobility declined very marginally when schools closed, while women’s mobility saw a clear discontinuity. This corroborates the hypothesis that women carry uneven

⁴Of course, the number of contagions surged in the weeks leading to the stay-at-home orders, which were imposed 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. To test if this assumption is valid, we did an RD on cases at the time of lockdowns. There is no clear discontinuity.

⁵Schools in Northern Italy closed on February 23. On March 4, the Italian government closed all schools and universities nationwide. The national stay-at-home order was announced with a Presidential Decree on March 9.

Figure 3: Impact of Stay-at-Home Orders on Mobility, by Gender
(Age group 25–44, percentage of people moving)



Notes: The chart shows the percentage of men and women moving divided into 20 equal-sized bins. The series are residualized with respect to province and day-of-the-week fixed effects.

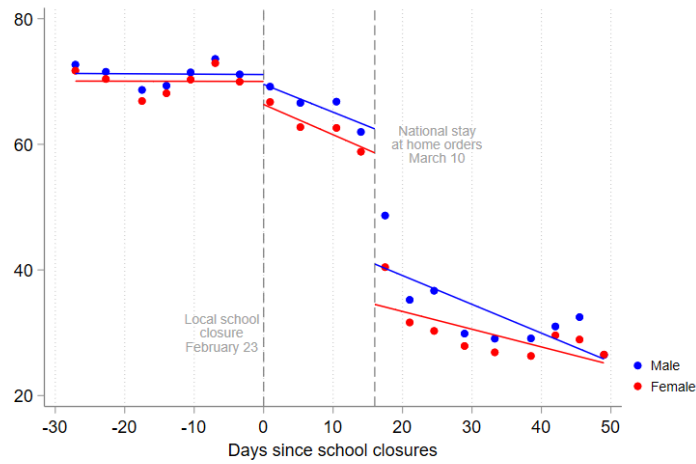
responsibilities in looking after the children when schools are closed.

Figure 4b provides additional evidence about the importance 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 a decline in women’s mobility relative to men’s. The mobility of men and women was similar before February 23. When schools closed, the heat map shows darker colors across all provinces, reflecting a disproportionate reduction in women’s mobility. The adoption of stay-at-home orders led to a further widening of the mobility gap.⁶

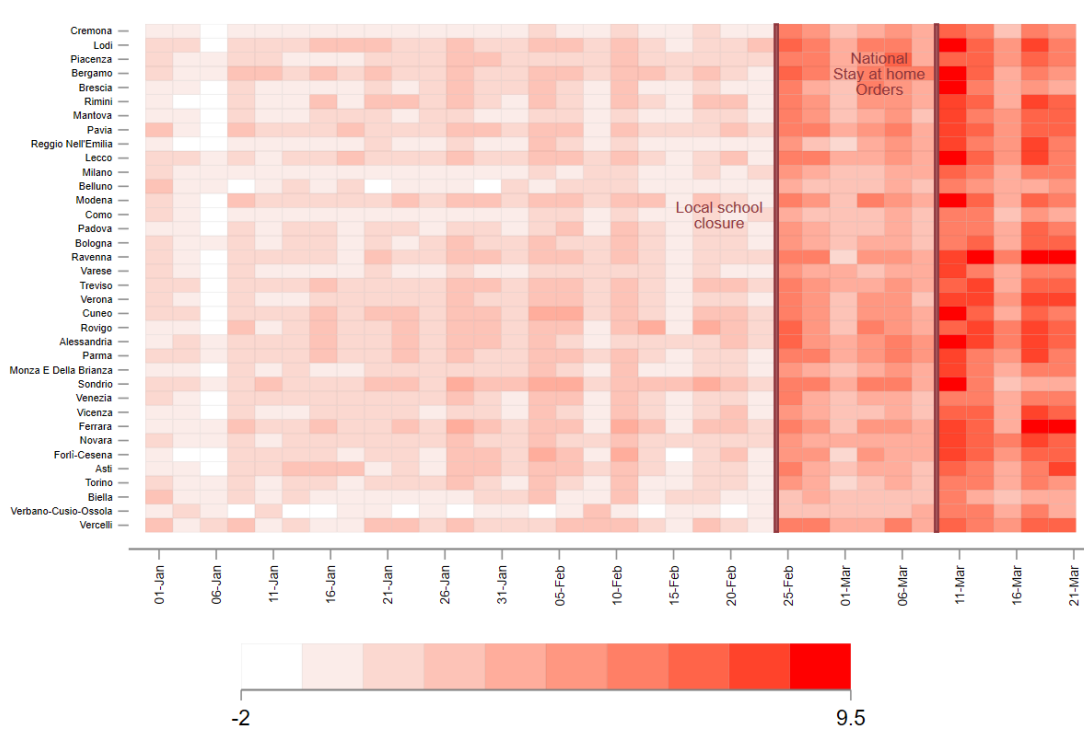
⁶Figure 4b provides also an additional insight. The provinces in figure 4b are listed in order of decreasing frequency of COVID-19 cases as a share of the province population on March 9. The absence of a clear vertical pattern indicates that the effects of school closures and lockdowns on the difference in mobility across gender 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, percentage of people moving)

(a) Average impact across provinces



(b) Impact across each provinces



Notes: Panel (a) reports a binned scatter plot where the percentage 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. Panel (b) reports the difference between men and women mobility in Friuli-Venezia-Giulia, Emilia-Romagna, Lombardia, Piemonte, and Veneto. Local school closures were introduced on February 23rd, and national stay-at-home orders on March 9th.

The differences in drop in mobility across age groups further confirm that school closures impacted disproportionately women’s mobility. The analysis presented so far has focused on people aged between 25 and 44. These cohorts are more likely to have young kids that require home supervision when schools are closed.⁷ 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 at the time of the lockdown, we follow Anderson (2014) and estimate this local linear regression:

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

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 $stay_{j,t} \times date_{j,t}$ 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 the 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 drop 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 when stay-at-home orders were imposed which coincided in most provinces with school closures. Column (1) and (3) present the results of the same specification for the age groups 18–24 and 45–64: the gap between women and men mobility is still present, statistically significant, but smaller, equal to 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+. These results show that the mobility gender gap is the largest for people aged 25–44 that are more likely to have young children. Therefore, they suggest that childcare needs largely explain the disproportionate impact of lockdowns on women’s mobility.

4.2 Local projections

The RD 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

⁷In Italy, as in other countries, grandparents often play an important 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 of school closures on parents’ mobility.

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*** (0.59)	-12.95*** (0.46)	-12.52*** (0.41)	-11.76*** (0.46)
Women × stay-at-home	-2.31*** (0.50)	-3.24*** (0.35)	-1.74*** (0.34)	1.55*** (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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

the epidemic is worsening and when 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 disaggregation of the Vodafone mobility data at the NUTS-3 level. Italy, Portugal, and Spain imposed lockdowns on a *national* scale in reaction to *localized* outbreaks. For example, in Italy 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. Leveraging on this observation, we exclude from the regression sample of each country the provinces that (i) registered the first 100 (cumulative) cases, (ii) had the highest 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.⁸ The regression thus examines the mobility response in the 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:

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

where variable $mob_{i,g,a,t+h}$ denotes the percentage of people moving in province i , of gender g and age a , at time $t+h$, with $h = \{1, \dots, 20\}$ being the horizon; $ln\Delta cases_{i,t-p}$ is the log of daily COVID-19 cases, which

⁸These 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 as shown in the robustness section.

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 persistence.⁹ The specification also features lags of the dependent variable to control for pre-existing trends; province and gender fixed effects to capture 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 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—among which travel restrictions, school and business closures, and stay-at-home orders—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 when the lockdown is introduced, reaching the trough after seven days. Mobility gradually resumes afterwards as the lockdown stringency impulse dissipates, as shown in Figure A.1 of Appendix A.

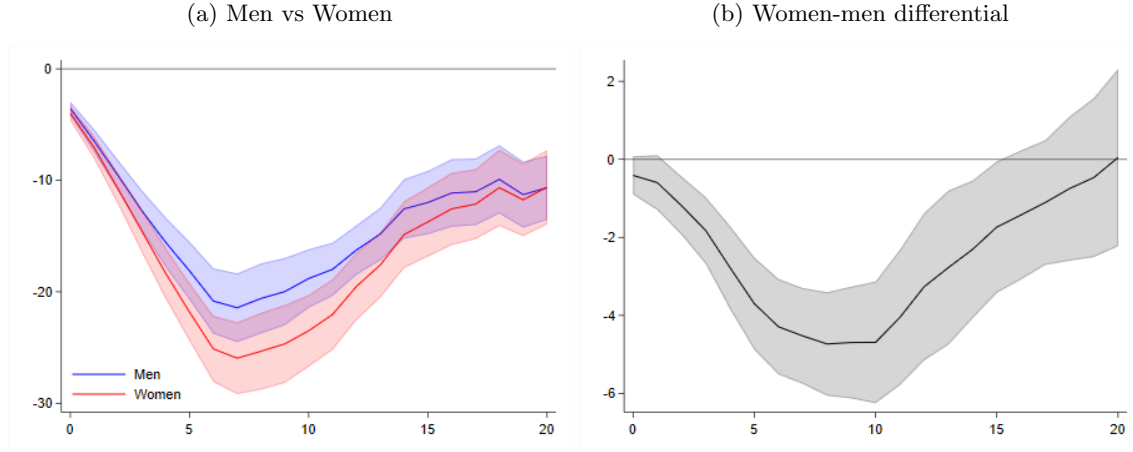
Most importantly, panel 5a reveals that lockdowns have an uneven effect on mobility across gender, impacting women more strongly. Women’s mobility falls by 26 percentage points seven days after the introduction of lockdowns, while that of men declines by about 21 percentage points. Panel 5b shows that the differential between the mobility of women and that of men is statistically significant until the lockdown impulse weakens 15 days after the initial tightening.

These results corroborate the findings of the RD analysis that lockdowns tend to impact women disproportionately. The RD analysis also showed 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 supervised 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 largest mobility gap between women and men in response to a full lockdown over the 20-day horizon of the local projections. We confirm that women’s mobility falls the most relative to men’s for people aged 25–44, with a differential of 4.7 percentage points. The gender gap declines to 3.6 percentage points for people aged 18–24 and to 3.7 percentage points 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. These results are thus closely aligned with those of the RD analysis.

Besides capturing the impact of lockdowns, the local projection in equation (3) also measures how mobility responds to an increase in COVID-19 infections holding constant the stringency of lockdowns. This is an important issue because during the pandemic people 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 was tightly correlated with the spread of COVID-19

⁹We 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, both at the subnational and national level. Since we want to measure the impact of *actual* restrictions at the national level, we re-construct the index excluding public information campaigns (as they aim to promote voluntary social distancing) and considering only measures that were adopted at the national level.

Figure 5: Impact of a Full Lockdown on Mobility, by Gender
(Age group 25–44, percentage of people moving)



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 2: Gender Gap at the Trough of the Estimated Response

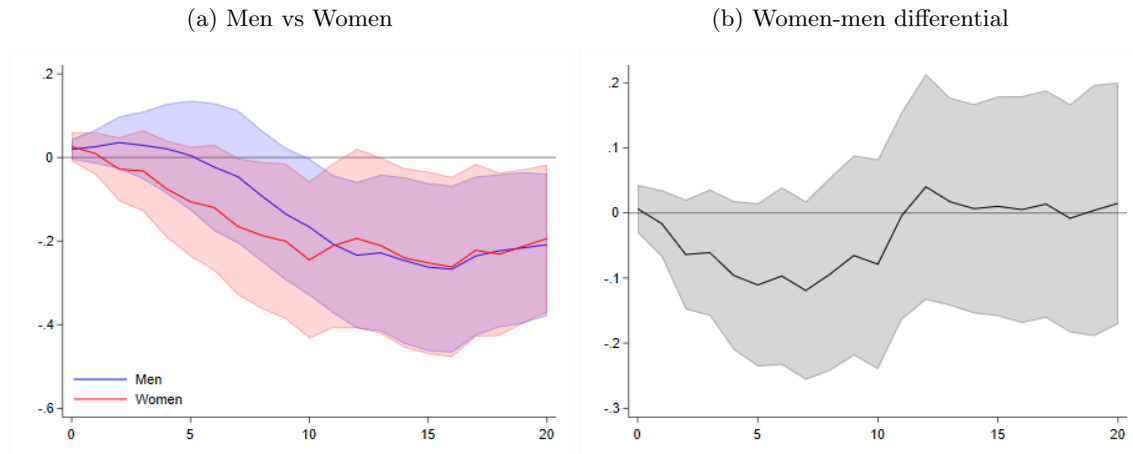
	18–24 (1)	25–44 (2)	45–64 (3)	65+ (4)
Lockdown stringency	-27.07*** (1.80)	-20.60*** (1.60)	-19.75*** (1.57)	-15.19*** (1.77)
Women × lockdown stringency	-3.58*** (0.95)	-4.73*** (0.67)	-3.69*** (0.73)	2.33 (1.46)
Days after the shock	7	7	7	7
Observations	18,798	18,798	18,830	17,872
Provinces	163	163	163	157
R-squared	0.94	0.94	0.94	0.91

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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

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.¹⁰ The interaction term between daily COVID-19 infections and the gender dummy 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 the mobility of both men and women. A doubling of daily COVID-19 cases leads to a contraction in mobility by about 0.3 percentage points 20 days after the introduction of the lockdowns. The effect is similar across men and women. Panel 6b shows indeed that there is no statistically difference across gender in how mobility responds to rising infections.

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



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.

¹⁰Besides 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 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

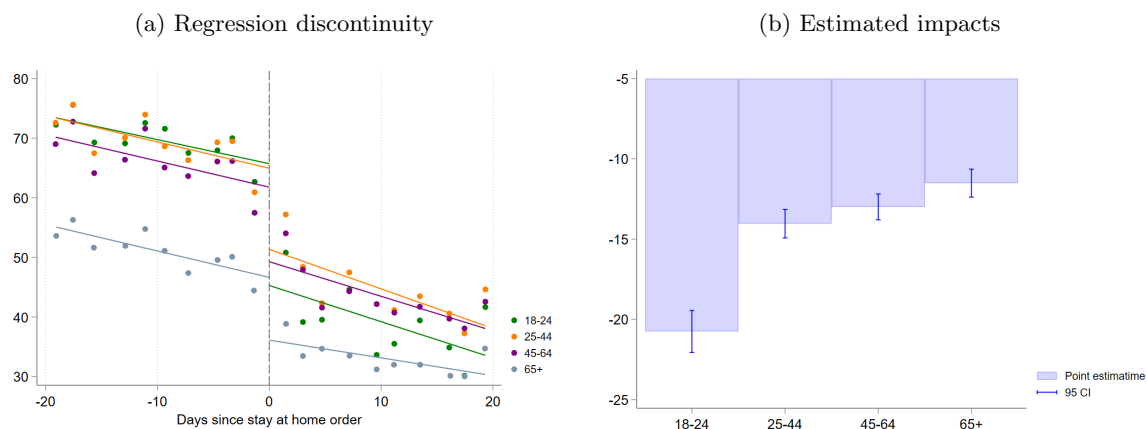
In this section, we examine if lockdowns have a different impact on mobility depending on people’s age. In line with the analysis on the effects across gender, we first examine the data using an RD approach. We then revisit the evidence using local projections which also allow us to examine if people 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 more clearly illustrated in panel 7b which shows the estimated mobility contraction for each age group using an RD specification akin to 1, with the associated 90 percent confidence intervals. The percentage of people below 25 years of age leaving home in a day declined by more than 20 points when stay-at-home orders were imposed. The mobility drop becomes progressively smaller for older people, being equal to only 11 percentage points for people aged 65 and above.

Figure 7: Impact of a Stay-at-Home Orders on Mobility, by Age
(Percentage of people moving)



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 5.

These results show that lockdowns tend to disproportionately impact the mobility of younger cohorts. This is not surprising if we consider 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 reach their work places and bring children to schools. Therefore, they are much more affected by lockdowns that, by impeding movements, have more adverse

effects on their employment opportunities.

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, and are estimated using data from provinces that did not experience early and/or major outbreaks and thus for whom the adoption of national lockdowns was largely exogenous. Formally, we estimate the following specification:

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

The specification features interaction terms between the lockdown stringency index and age group dummies $agegroup_i^a$, with $a = \{1 = [18, 24]; 2 = [25, 44]; 3 = [45, 64]\}$, 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 declines sharply across all age categories during the first ten days after the introduction of a lockdown. Yet, the younger cohorts experience a considerably larger drop in mobility, reaching a trough of 30 percentage points nine days after the introduction of lockdowns for people aged 18–24.¹¹

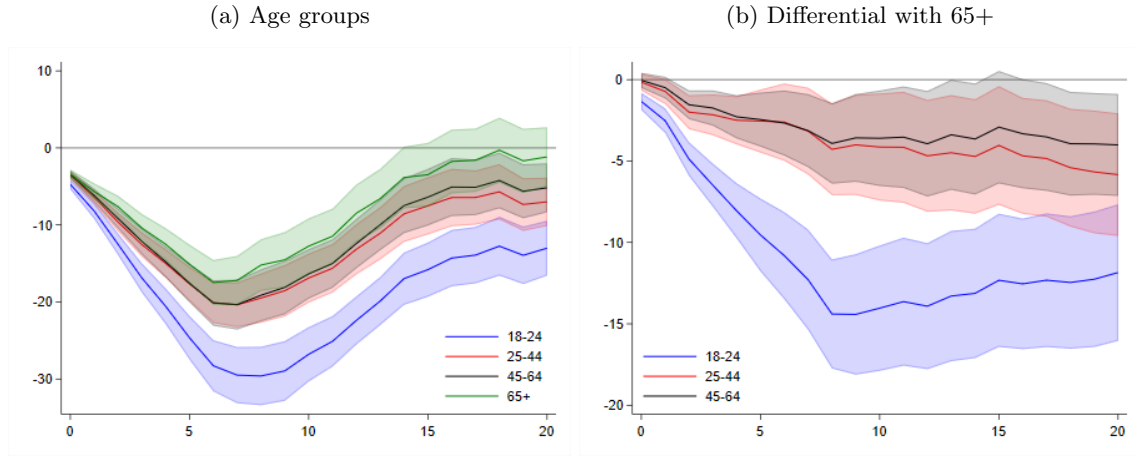
To illustrate the differences across age groups more clearly, panel 8b shows the mobility differential of each age group relative to people aged 65+. The mobility differential of those aged 18–24 compared to people 65+ is the largest and becomes statistically significant since the beginning of the projection horizon. The drop in mobility for the age groups 25–44 and 45–64 becomes statistically larger relative to people 65+ a few days after the lockdown stringency impulse. These findings corroborate the results from the RD analysis in the previous section, showing that lockdowns impact the mobility of younger cohorts disproportionately.

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 react most forcefully to a rise in infections. A doubling of COVID-19 cases leads to a fall in mobility by about one percentage point after 20 days for people aged 18–24. A rise in infections leads also to reduction in mobility for people aged 25–44 and 45–64, even though the effect is more modest. Mobility remains instead broadly unchanged for people aged 65+, actually increasing marginally towards the end of the projection horizon. Panel 9b confirms that mobility declines more for all working-age groups relative to people 65+.

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

¹¹The mobility dynamics reflect the underlying impulse to the lockdown stringency that dissipates after two weeks as illustrated in Figure A.1a of Appendix A.

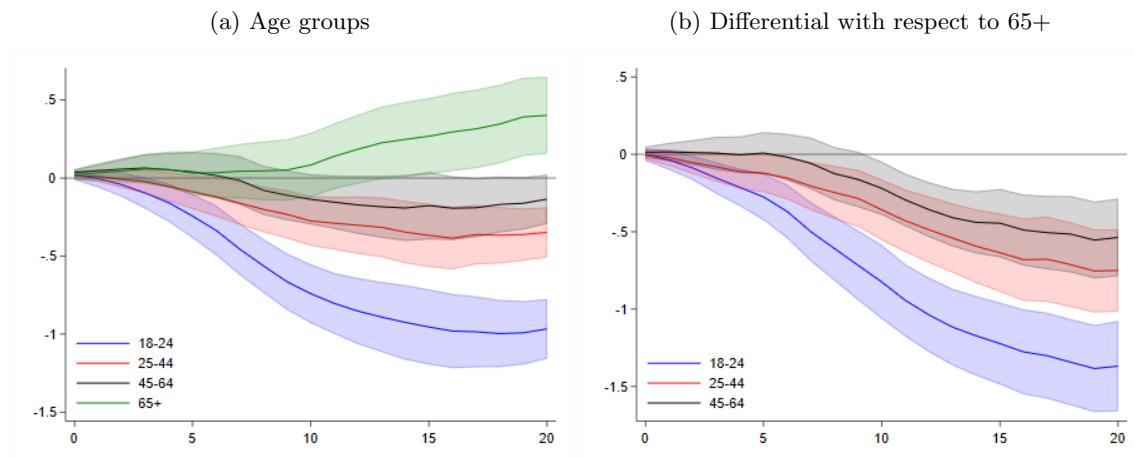
Figure 8: Impact of a Full Lockdown on Mobility, by Age Group
(Percentage of people moving)



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.

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 larger response in the mobility of the young may reflect their stronger concerns about the virus. This is consistent with the evidence presented in Bordalo et al. (2020). Based on 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
(Percentage of people moving)



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 impact of lockdowns on mobility across gender for people aged 25–44 based on the RD analysis. Column (1) reports the differential impact between women and men from the baseline specification in section 4.1. 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) and (5) we exclude, one at a time, data for Portugal and Spain for which the mobility series are interpolated. The differential impact on women’s mobility is confirmed on these different samples. Finally, in column (6) we 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. Finally, since the mobility indicators is bounded between 0 and 100 percent, we verify in column (7) that the results are robust to applying a logistic transformation.

Table 3: Robustness of the RD Results by Gender

	Baseline (1)	10-day window (2)	Excl. North. Italy (3)	Excluding Portugal (4)	Excluding Spain (5)	Controlling for cases (6)	Logistic transformation (7)
Stay at home	-12.95*** (0.459)	-8.99*** (0.424)	-12.82*** (0.548)	-14.42*** (0.440)	-13.04*** (0.454)	-11.69*** (0.469)	-0.62*** (0.024)
Women × stay-at-home	-3.24*** (0.353)	-3.07*** (0.431)	-3.07*** (0.442)	-3.58*** (0.403)	-3.79*** (0.149)	-3.25*** (0.352)	-0.13*** (0.018)
Observations	14,102	7,228	11,150	12,052	10,742	14,098	14,291
R-squared	0.86	0.85	0.85	0.84	0.92	0.86	0.87

Notes: The table reports the coefficients of an interaction term between the gender dummy 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 (6) considers a 20-day window and controls for the moving average of daily COVID-19 cases. Standard errors are clustered at the province level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

We also verify if the results are robust across provinces with different characteristics. In Table 4 we re-run the RD analysis splitting the sample depending on whether provinces are above or below the median level of GDP per capita, population density, and population size. Furthermore, we split the sample based on whether provinces are considered rural or not by the Eurostat. Across all these samples, lockdowns continue to have a disproportionate effect on the mobility of women, thus confirming that the results are not driven by provinces with peculiar characteristics.

Following a similar set of robustness tests, Table 5 and Table 6 corroborate the RD findings that lockdowns have a stronger impact on the mobility of younger cohorts.

We also perform various robustness tests for findings of the local projections. As discussed in section 4.2, the local projections are estimated 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 affect our estimates if lockdowns or voluntary social distancing had different effects on mobility in regions more

Table 4: Robustness of the RD Results by Gender across Provinces

	GDP pc below median (1)	GDP pc above median (2)	Pop. density below median (3)	Pop. density above median (4)	Population below median (5)	Population above median (6)	Rural (7)	Not rural (8)
Stay-at-home	-12.166*** (0.716)	-14.141*** (0.512)	-11.588*** (0.682)	-14.617*** (0.538)	-10.985*** (0.641)	-14.826*** (0.588)	-15.930*** (0.744)	-11.902*** (0.534)
Women \times stay at home	-2.655*** (0.555)	-3.859*** (0.415)	-2.849*** (0.644)	-3.648*** (0.253)	-3.502*** (0.372)	-2.974*** (0.595)	-3.379*** (0.389)	-3.185*** (0.462)
Observations	7,203	6,899	7,293	6,809	6,992	7,110	3,695	10,407
R-squared	0.861	0.880	0.817	0.915	0.876	0.838	0.881	0.852
NUTS3 FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Days of week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

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 (1) and (2) are estimated splitting the sample above and below the median of GDP per capita in a given province for 2018. Column (3) and (4) splits the sample below and above the median of population density as measured by the population per squared km in 2018. Column (5) and (6) reports results splitting the sample below and above the median of population in 2018. Column (7) and (8) reports the coefficients estimated splitting the sample according to whether the province is considered rural or not according to Eurostat. Standard errors are clustered at the province level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: Robustness of the RD Results by Age

	Baseline (1)	10-day window (2)	Excl. North. Italy (3)	Excluding Portugal (4)	Excluding Spain (5)	Controlling for cases (6)	Logistic transformation (7)
Stay-at-home	-11.515*** (0.441)	-9.107*** (0.429)	-11.015*** (0.537)	-12.346*** (0.459)	-12.707*** (0.354)	-10.486*** (0.425)	-0.515*** (0.020)
18–24 \times stay at home	-9.238*** (0.505)	-6.992*** (0.461)	-8.932*** (0.630)	-10.854*** (0.468)	-9.333*** (0.464)	-9.124*** (0.502)	-0.412*** (0.020)
25–44 \times stay at home	-2.526*** (0.445)	-0.752* (0.447)	-3.031*** (0.550)	-3.127*** (0.485)	-1.354*** (0.244)	-2.453*** (0.435)	-0.122*** (0.021)
45–64 \times stay at home	-1.482*** (0.489)	0.244 (0.504)	-2.361*** (0.597)	-1.957*** (0.542)	0.708*** (0.196)	-1.421*** (0.478)	-0.063*** (0.023)
Observations	55,264	28,258	43,456	47,672	42,360	55,248	55,264
R-squared	0.852	0.847	0.838	0.834	0.919	0.857	0.864

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 (6) considers a 20-day window and controls for the moving average of daily COVID-19 cases. Standard errors are clustered at the province level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Robustness of the RD results by Age across Provinces

	GDP pc below median (1)	GDP pc above median (2)	Pop. density below median (3)	Pop. density above median (4)	Population below median (5)	Population above median (6)	Rural (7)	Not rural (8)
Stay at home	-9.866*** (0.649)	-13.328*** (0.534)	-9.876*** (0.725)	-13.186*** (0.437)	-11.278*** (0.656)	-11.730*** (0.603)	-13.559*** (0.661)	-13.559*** (0.661)
18-24 × stay at home	-8.307*** (0.779)	-10.188*** (0.621)	-8.359*** (0.829)	-10.327*** (0.508)	-8.069*** (0.668)	-10.405*** (0.707)	-10.039*** (0.693)	-10.039*** (0.693)
25-44 × stay at home	-3.463*** (0.748)	-1.666*** (0.434)	-2.937*** (0.799)	-2.293*** (0.352)	-0.950* (0.549)	-4.074*** (0.648)	-3.415*** (0.541)	-3.415*** (0.541)
45-64 × stay at home	-2.970*** (0.778)	-0.007 (0.542)	-2.978*** (0.891)	-0.087 (0.325)	-0.178 (0.608)	-2.764*** (0.737)	-1.550*** (0.538)	-1.550*** (0.538)
Observations	27,914	27,350	28,115	27,149	27,089	28,175	14,659	14,659
R-squared	0.854	0.870	0.810	0.908	0.865	0.840	0.876	0.876
NUTS3 FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Days of week FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Gender FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The table reports the coefficients of an interaction term between the age dummy and and stay-at-home variable. All regressions include the age 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, gender and day-of-the-week fixed effects. Column (1) and (2) are estimated splitting the sample above and below the median of GDP per capita in a given province for 2018. Column (3) and (4) splits the sample below and above the median of population density as measured by the population per squared km in 2018. Column (5) and (6) reports results splitting the sample below and above the median of population in 2018. Column (7) and (8) reports the coefficients estimated splitting the sample according to whether the province is considered rural or not according to Eurostat. Standard errors are clustered at the province level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

impacted by the virus. Thus, we test if our findings are robust to the inclusion of those provinces. Figure 10 shows the coefficients on the interaction terms that capture the differential impact of lockdowns and COVID-19 cases on 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 took a larger toll on younger cohorts. In terms of magnitudes, the estimated effects are virtually identical to those in the baseline.

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.¹² 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 11a and 11b of Figure 11 confirm the disproportional impact of lockdowns on women’s mobility compared to men’s, as well as that rising infection do not generally have a statistically significant different effect across gender. If anything, the women-men mobility differential in response to lockdowns appear more persistent than with time fixed effects. Similarly, without time fixed effects, panel 11c shows that lockdowns are still found to impact more strongly the mobility of younger people. Finally, panel 11d 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.

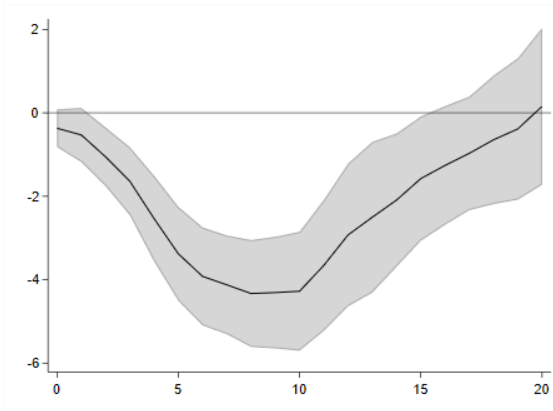
Another possible concern with the local projections is that the results could be affected by the linear interpolation of the missing observations for Portugal and Spain. Thus, we re-estimate the local projections excluding one of these two countries at a time.¹³ Figure 12 presents the results of the impact of lockdowns and voluntary social distancing when data for Portugal are excluded. Figure 13 repeats the exercise excluding

¹²Such movements could be caused by public announcements by the government, public health officials, and international organizations, or by news about celebrities being infected.

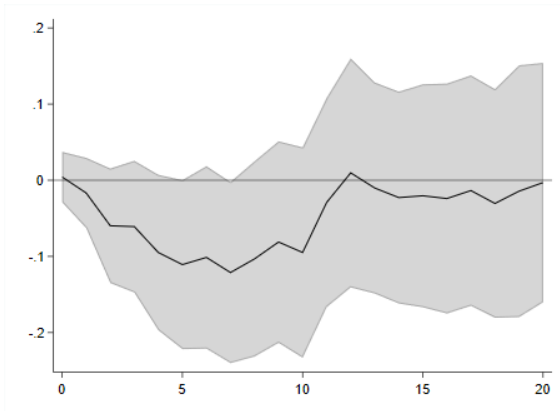
¹³Running regression with data for a single country would prevent the possibility of including time fixed effects.

Figure 10: Impact of Lockdowns and COVID-19 Cases using Full Sample
(Percentage of people moving)

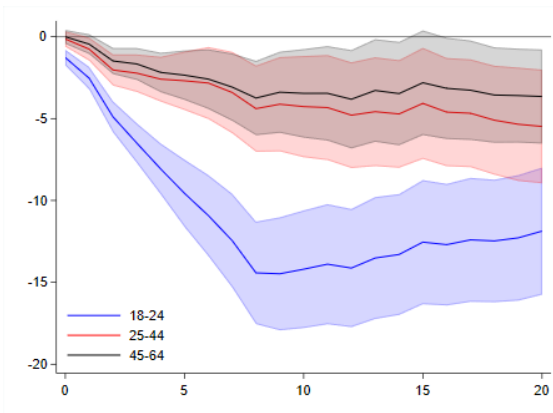
(a) Impact of a full lockdown,
women-men differential



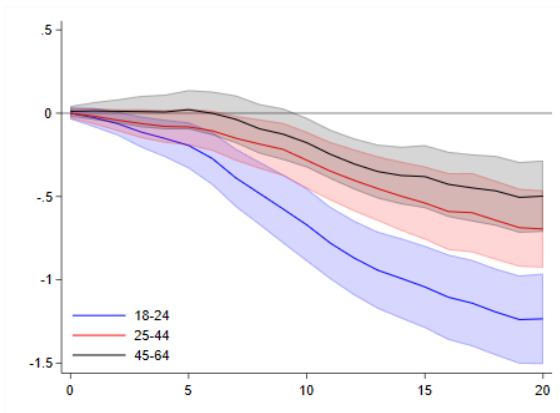
(b) Impact of a doubling of COVID-19 cases,
women-men differential



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

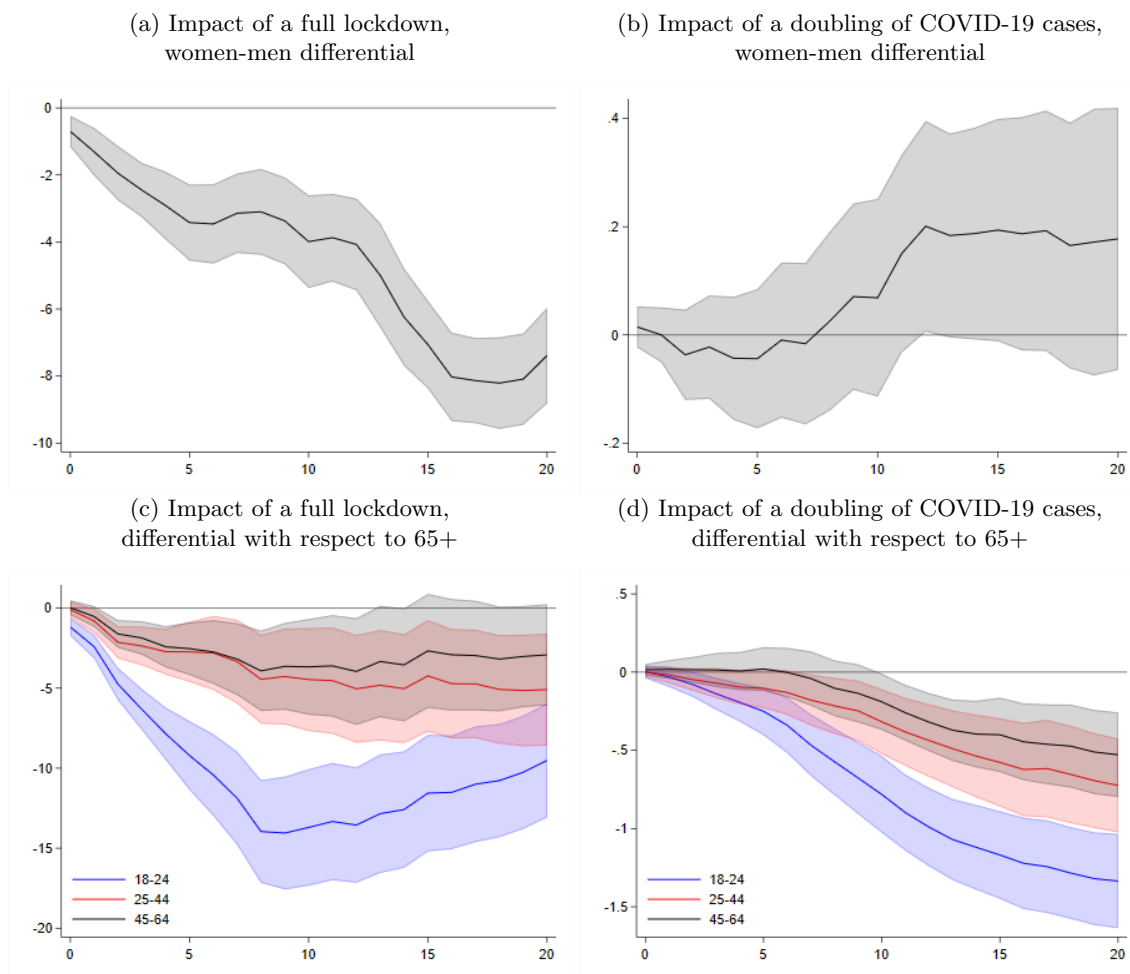


(d) Impact of a doubling of COVID-19 cases,
differential with respect to 65+



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.

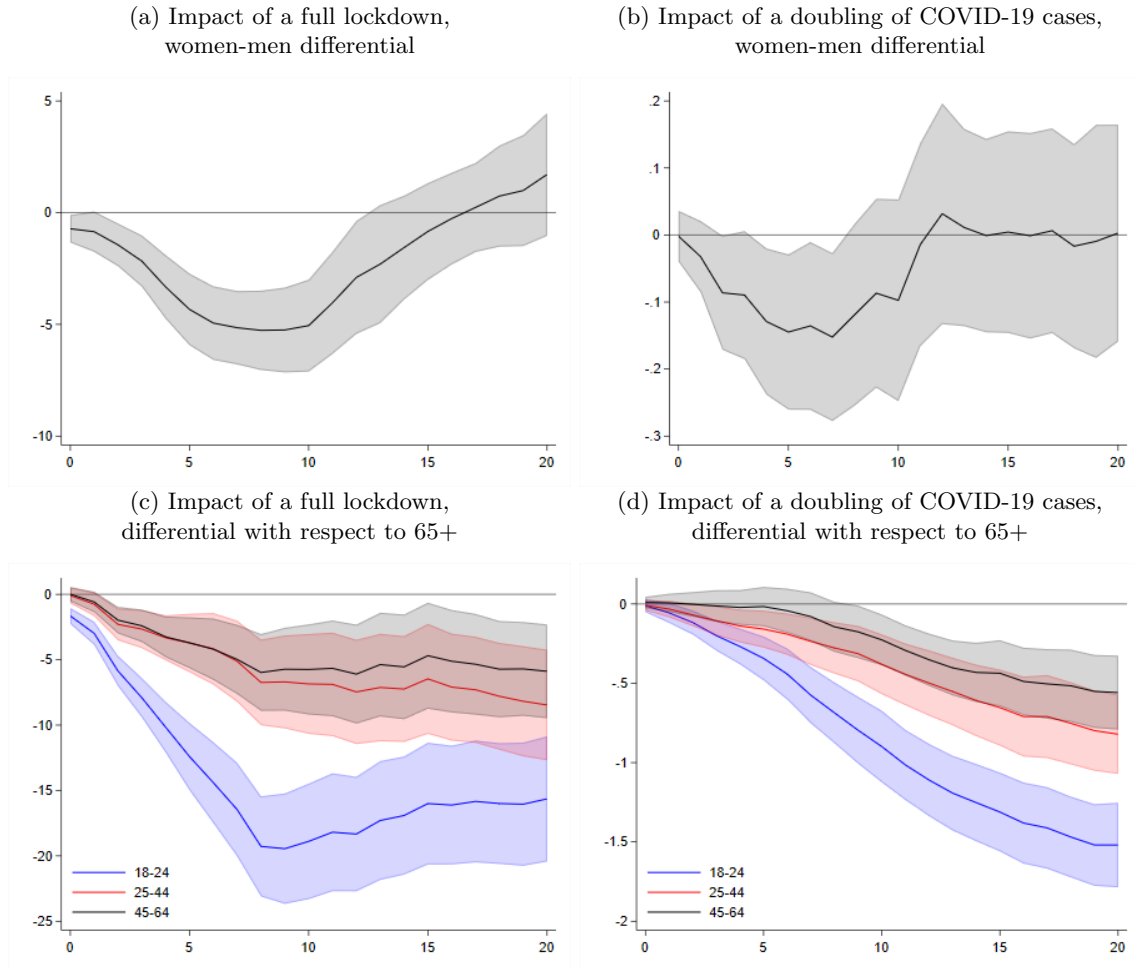
Figure 11: Impact of Lockdowns and COVID-19 Cases using Day-of-the-Week Fixed Effects
(Percentage of people moving)



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.

Spain. In both cases, the results closely mimic those in the baseline analysis. The only difference is that an increase in COVID-19 cases tends to have a larger impact on women’s mobility relative to men’s, especially when Spain is excluded in panel 13b. Yet, the point estimates are quantitatively very small.

Figure 12: Impact of Lockdowns and COVID-19 Cases Excluding Portugal
(Percentage of people moving)

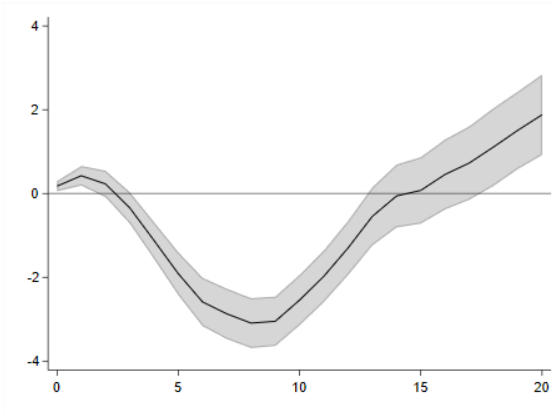


Notes: The x-axes denote the number of days, the lines denote the point estimates, and the shaded areas correspond to 90 percent confidence intervals computed with standard errors clustered at the province level.

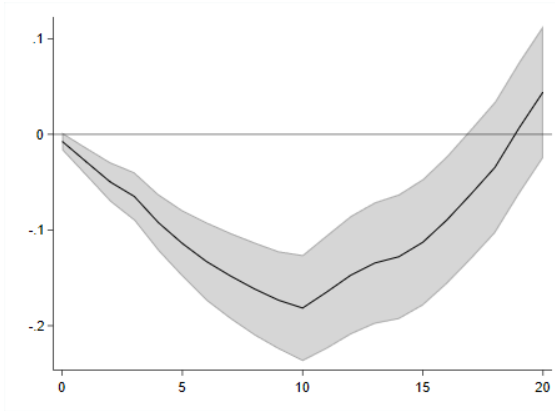
As in the case of the RD analysis, we also verified that the results of the local projections are robust to applying a logistic transformation to the mobility indicator and to splitting the sample depending on province characteristics. These results are available upon request.

Figure 13: Impact of Lockdowns and COVID-19 Cases Excluding Spain
(Percentage of people moving)

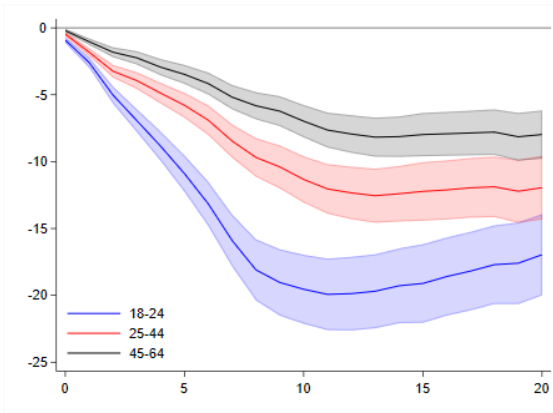
(a) Impact of a full lockdown,
women-men differential



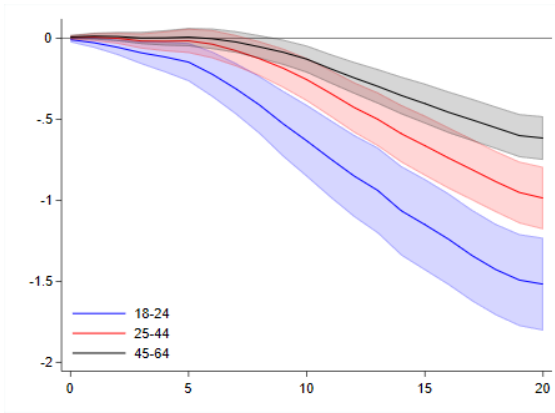
(b) Impact of a doubling of COVID-19 cases,
women-men differential



(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

We have used unique mobility indicators provided by Vodafone which differentiate by gender and age characteristics 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 larger impact on women's mobility especially for those aged between 25 and 44 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 mobility gender gap opened 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 schedules impact female labor force participation (Berthelon et al., 2015.) Our study provides complementary high-frequency evidence.

These findings warn about a possible widening of gender inequality, as women may compromise their employment opportunities if they have to stay at 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 by preventing them from reaching their work places and bringing children to school.

Interestingly, the mobility of younger people responds more strongly also to rising infections, for a given level of the stringency of lockdowns. Consistent with survey evidence, this could be because younger people are more concerned about the virus 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 young 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 tend to have less stable job contracts that are more likely to be terminated during a crisis. These considerations highlight the need for a social pact 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 enough public support to deploy lockdown measures when needed.

Third, the results on the differential effects on age groups 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 lasting. 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 react differently to stay-at-home orders, school closures, and COVID-19 cases can provide important inputs for the formulation of targeted policies in the context of resurgent epidemics in several regions of the world.

References

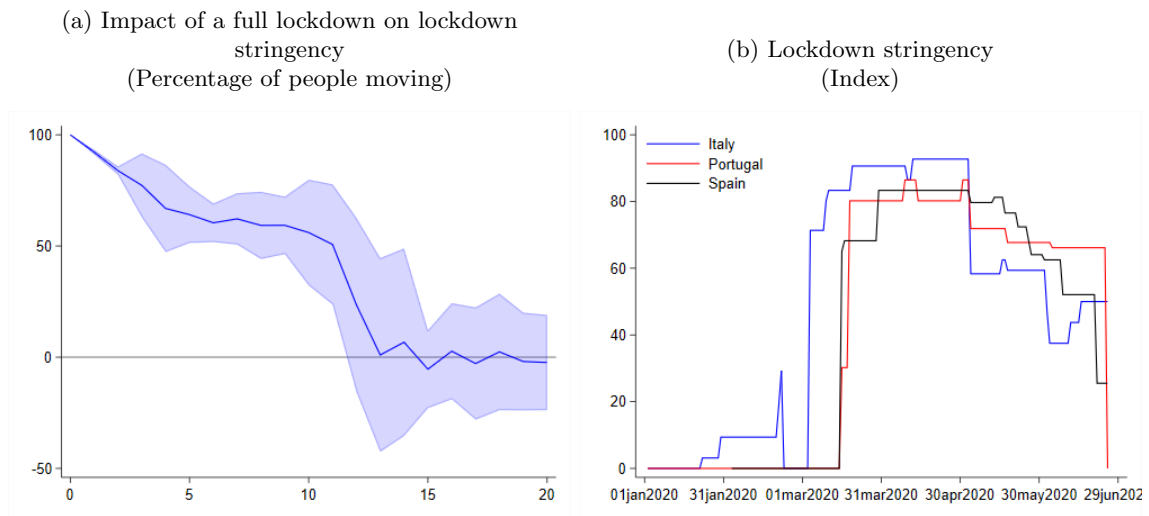
- Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D. Whinston (2020). “Optimal Targeted Lockdowns in a Multi-Group SIR Model”. NBER Working Paper No. 27102.
- Adams, Abigail, Teodora Boneva, Marta Golin, and Christopher Rauh (2020). “Inequality in the Impact of the Coronavirus Shock: 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.
- Glaeser, Edward L, Caitlin S Gorbach, and Stephen J Redding (2020). “How Much Does COVID-19 Increase with Mobility? Evidence from New York and Four Other us Cities”. NBER, Working Paper No. 27519.
- Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and Jose-Victor Rios-Rull (2020a). “Health versus Wealth: On the Distributional Effects of Controlling a Pandemic”. CEPR DP 14606.
- Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Victor Ríos-Rull (2020b). “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.
- Lourenco, Pedro Rente, Gurjeet Kaur, Matthew Allison, and Terry Evetts (2020). “Data sharing and collaborations with Telco data during the COVID-19 pandemic: A Vodafone case study”. *Data & Policy*.
- 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.
- Oreopoulos, Philip, Till Von Wachter, and Andrew Heisz (2012). “The short-and long-term career effects of graduating in a recession”. *American Economic Journal: Applied Economics* 4(1), pp. 1–29.
- Sevilla, Almudena and Sarah Smith (2020). “Baby Steps: The Gender Division of Childcare during the COVID19 Pandemic”. CEPR Covid Economics Issue 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 and dissipate after about two weeks. 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, countries have adjusted the stringency of lockdowns rather frequently.

Figure A.1: Lockdown Stringency Dynamics



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.