DISCUSSION PAPER SERIES

DP14814

CLOSED FOR BUSINESS

Dion Bongaerts, Francesco Mazzola and Wolf Wagner

PUBLIC ECONOMICS

ISSN 0265-8003

CLOSED FOR BUSINESS

Dion Bongaerts, Francesco Mazzola and Wolf Wagner

Discussion Paper DP14814
Published 25 May 2020
Submitted 22 May 2020

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

CLOSED FOR BUSINESS

Abstract

We investigate the effectiveness of business shutdowns to contain the Covid-19 disease. In March 2020, Italy shut down operations in a number of sectors. Using a difference-in-difference approach, we find that municipalities with higher exposure to closed sectors experience subsequently lower mortality rates. We estimate the resulting life savings to exceed 12.000 people over less than a month. Using estimates of remaining life-years, this translates into monetary benefits of 12 billion Euros We also show that business shutdowns exhibit rapidly diminishing returns and have effects outside the closed sectors and in other municipalities. This suggests that effective containment policies require central coordination.

JEL Classification: H12, I18

Keywords: Pandemic, COVID-19, business shutdown, coordination

Dion Bongaerts - dbongaerts@rsm.nl
Erasmus University

Francesco Mazzola - mazzola@rsm.nl
Erasmus University

Wolf Wagner - wagner@rsm.nl
Erasmus University and CEPR

Closed for business

DION BONGAERTS, FRANCESCO MAZZOLA and WOLF WAGNER*

May 22, 2020

Abstract

We investigate the effectiveness of business shutdowns to contain the Covid-19 disease. In March 2020, Italy shut down operations in a number of sectors. Using a difference-in-difference approach, we find that municipalities with higher exposure to closed sectors experience subsequently lower mortality rates. We estimate the resulting life savings to exceed 12.000 people over less than a month. Using estimates of remaining life-years, this translates into monetary benefits of 12 billion Euros We also show that business shutdowns exhibit rapidly diminishing returns and have effects outside the closed sectors and in other municipalities. This suggests that effective containment policies require central coordination.

Keywords: Pandemic; Covid-19; business shutdown; coordination;

*All authors are affiliated with Erasmus University, Rotterdam School of Management. Wolf Wagner is also affiliated with CEPR. Emails: dbongaerts@rsm.nl, mazzola@rsm.nl and wagner@rsm.nl.

1 Introduction

In attempts to contain virus outbreaks, such as Covid-19, policy makers trade off public health against economic costs. Yet, little is known about how effective containment policies are, and whether such measures are better organized in a centralized or decentralized way.[1] In the wake of economically harmful containment measures, this tradeoff is under scrutiny. Moreover, as measures are slowly reversed, a discussion has emerged about decentralizing containment policies, leaving decisions more to local governments,[2] and, possibly, even to businesses themselves. Clearly, a localized approach requires absence of strong spillovers from containment policies. The risk of a potential second wave underlines the urgency of understanding and quantifying the effects of containment measures.

The empirical identification of containment effects is challenging, for several reasons. Policymakers have implemented containment measures in response to a rapidly evolving pandemic. There is thus no clear counterfactual to the policy – the spread of the virus will have changed also in the absence of the policy. Containment decisions are also often clustered as pandemics develop very quickly, making it difficult to isolate the impact of a specific policy. They tend to be introduced at a time when the general public becomes very aware of the dangers of the virus, and takes (self-imposed) measures to reduce the risk of contracting the virus (the spread of the virus may thus slow unrelated to the policy). These identification issues are compounded by the fact that there is an uncertain, and variable, lag between contracting the virus and the disease taking effects, and thus also between policy and health outcomes.

This paper exploits within-country variation in exposure to a nationwide policy to

[1] In some countries (e.g., the US and Germany) containment policies are largely locally organized, whereas in others (e.g., the Netherlands), policies are mainly organized at the national level. Containment policies are not necessarily orchestrated by public authorities. The first lockdowns in Brazil were imposed by drug gangs in favelas in Rio de Janeiro.

[2] In Germany, the state of Thuringia is the first to reverse Covid-19 policies, leaving decisions about containment policies to local municipalities.

study the impact of business shutdowns on mortality. During March 2020 Italy was the first country in Europe to completely shut down a selected number of sectors in its economy. Specifically, on March 11th sectors comprising 17% of the economy (in terms of employment) were closed; additionally, on March 25th sectors comprising 34% of total employment were closed. The (nationwide) closure policy affected Italian municipalities differently due to differences in sectoral exposures. Our empirical strategy is to study differential changes in mortality patterns across municipalities to this (common) policy. This allows to control for confounding factors taking place in Italy around the time of the policy. Our empirical strategy also accounts for heterogeneity across municipalities arising from the fact that they were at different stages in the pandemic at the time of the policy. We use a statistical approach to classify the begin of the pandemic in each municipality, and then examine variation among municipalities with similar (pre-policy) pandemics.

The results suggest that business shutdowns are effective in saving human lives: municipalities with greater exposures to either the first or the second shutdown see a decline in death rates relative to other municipalities. We undertake a counterfactual analysis that shows that the first shutdown (which provides cleaner identification than the second) saved more than 9000 Italian lives over a period of 23 days. We find the second shutdown to be less effective (per unit of economy closed down), consistent with the fact that at that time the virus was already more under control. Estimates of the value of a statistical life-year imply a large societal benefit from business closures, which for the first shutdown exceed nine billion Euros.

Our analysis also shows that business shutdowns have important spillovers. Business closures may affect the spread of the virus outside a municipality because of commuting (other forms of travel were fairly restricted during our sample period). Consistent with this we find that shutdown exposure in business centres affects death rates in neighboring municipalities. The beneficial effect is large, and comparable in size to the impact of a municipalities' own shutdown exposure. We also find that greater shutdown exposure has a strong effect on parts of the population that is very unlikely to work. This points to significant contagion effects that reach beyond the firms undertaking economic activities. The existence of different forms of spillovers points to benefits from centralizing or coordinating shutdown decisions at a central level.

Our analysis further points at rapidly declining benefits to scale from sectoral shutdowns. We compare the marginal effectiveness of shutdowns across municipalities that differ with respect to what proportion of their economy was affected by the first shutdown. We find that the marginal effectiveness in municipalities with the lowest sector exposure to be more than three times higher than the average marginal effectiveness across all municipalities. In addition, (marginal) effectiveness declines monotonically with the total share of sectors closed down. These results are consistent with a lower effectiveness of the second shutdown, and suggest benefits from targeted business closures, rather than widespread closures.

In the wake of the Covid-19 crisis a significant theoretical literature has emerged that examines optimal policies during a pandemic. This literature emphasizes production externalities as a rationale for public policies. Production externalities arise when the provision of goods and services result in the spread of the virus to individuals not directly involved in the business activities. As firms (and their workers) will not internalize the social cost of such contagion, they will make inefficient containment decisions, providing a rationale for government-imposed shutdowns. Our paper is the first to document empirical evidence consistent with such production externalities. Eichenbaum et al. (2020) study optimal taxation of business activities (which can be interpreted as shutdown intensity) in an environment where the severity of the production externalities vary with the spread of the virus. Our finding of lower policy effectiveness once the pandemic is more under control is consistent with the theoretical premises of their model. Krueger et al. (2020) show that production externalities are mitigated by individuals shifting activities to environments that pose less contagion risk. Our estimates – which are net of such mitigating behavior – suggest that production externalities remain significant. Calibrating an SIR-model to the US economy, Bethune and Korinek (2020) show that the social cost of infections exceeds the private cost by factor two, consisting with our results of large spillovers on individuals that are unlikely to work. Beck and Wagner (2020) analyze virus contagion across jurisdictions (countries in their model) and show that interjurisdictional externalities creates a need for coordination of containment policies. Our findings of strong geographical spillovers provides evidence for the existence of such externalities.

The various containment measures enacted during the Covid-19 crisis have led to a very rapidly growing literature that tries to understand their benefits and costs. Several papers have examined the impact on infections and mortality. Using either time-series or cross-country variation, these studies have generally concluded that these measures are effective (see Hartl et al. (2020) for Germany, Qiu et al. (2020) for China, and Ullah and Ajala (2020) for a cross-country study). Our study, using within-country variation in exposure to national business shutdowns in Italy, confirms the effectiveness of containment measures.

The next section describes our empirical approach and data. Section 3 contains the empirical analysis of death rates in Italian municipalities. Section 4 presents the results. Finally, section 5 concludes.

2 Methodology and Data

On March 11th the Italian Prime Minister announced a nationwide shutdown of all food, retail and personal-services activities, except for first-necessity goods. Businesses like supermarkets, small grocery shops, pharmacies, and newsstand kiosks were allowed to remained open. Next to business shutdowns, the decree also limited personal mobility.

Two weeks later, on March 25th, an extension of the duration of the shutdown was announced. In addition, the list of sectors included in the shutdown was enlarged.

For reasons of identification, our study primarily focuses on the first shutdown. We study the impact of this shutdown using a difference-in-difference (DD) approach. Given the lag between a virus infection and (possible) subsequent a death, the "treatment" date does not coincide with the day the policy was enacted. In particular, it is impossible for the policy to have any effect on deaths already from the first day after the policy onwards. Estimates suggest a median of five days between the exposure to the virus and the occurrence of first symptoms ("the incubation" period, Lauer et al., 2020) and about eight days between first symptoms and death.[3] Thus, the average time to death across individuals is about 13 days. As the policy should already start to show effects prior to the average, we take the treatment date to be day 10.

Our empirical model takes the following form:

$$y_{m,t} = y_{m,t-1} + \gamma d11_t + \phi(d11_t \times PolicyExposure11_m) + \tag{1}$$

$$+\eta d25_t + \psi(d25_m \times PolicyExposure25_m) + FE_m + FE_t + \varepsilon_{m,t},$$

where y is the number of (Covid-related) deaths in municipality m on day t. We include the lagged value, $y_{m,t-1}$, as a determinant since epidemiological models (such as the SIR model, (Kermack and McKendrick, 1927; Allen, 2017)) show that new infections condition highly on the prevalent share of infected people in the population. We include two dummies, $d11_t$ and $d25_t$, to indicate the treatment date for the first and second shutdown (10 days after the respective announcement dates). The variables $PolicyExposure11_m$ and $PolicyExposure25_m$ measure the exposure of a municipality to sectors that were (newly) shut down at the first and second shutdown. We include time and municipality fixed effects to absorb any day-specific effects and (time-invariant) municipality effects,

[3]See for example the CS 22/2020 report by the Health Superior Institute at `https://www.iss.it/coronavirus/-/asset_publisher/1SRKHcCJJQ7E/content/id/5304852`

respectively. The inclusion of time fixed effects in particular allows to account for any (common) effects arising around the shutdowns, for example, due to reduced personal mobility. Our variable of interest is the interaction coefficient ϕ, which estimates whether municipalities with a higher shutdown exposure experience fewer (daily) deaths as a consequence of the policy. If the policy shutdown is effective at reducing the spread of the virus and thus, ultimately, reduces deaths, the prediction is that the coefficient ϕ will enter with a negative sign.

We next describe the calculation of the variables. We construct measures of shutdown exposure using data on employment and establishments of Italian firms made available by Italian Statistical Agency (ISTAT). The data provides sectoral data at the municipality level from the year 2017, including information on the number of employees and employers, revenues and number of establishments.We construct a (continuous) municipality-level shutdown exposure $Shutdown11_m$ by dividing the total of employees and employers in sectors shutdown on March 11th by the total number of employees and employers in the municipality.[4] We construct an equivalent measure of the second shutdown, $Shutdown25_m$, using the employment ratio of sectors that were (newly) shutdown on the 25th to total employment. Figure 1 depicts sectoral shutdown exposures across municipalities, ordered by a municipality's exposure to the first shutdown. We can see that municipalities differ significantly with respect to the first shutdown exposure (blue portion).

[Figure 1 here]

In particular, it is interesting to note that the second shutdown exposure exhibits a

[4]Shutdown sectors correspond to the following European classification of the economic activities (NACE) codes: "451", "452", "473", "474", "477", "478" for the retail industry; "561", "563" for the food and beverages industry; "96" for the personal-services industry. We exclude employment in schooling and sports (NACE codes "85" and "931") from the denominator since these sectors were already shut down weeks before.

strong correlation with the first exposure (the correlation coefficient is -0.53).[5]

We next describe our measure of policy effectiveness, which is based on fatalities per 100,000 people. ISTAT recently released death registry data on a sample of 6,866 municipalities covering all Italian regions, corresponding to approximately 90% of the entire population (Italy has 7,904 municipalities in total). The dataset contains the number of deaths, together with the residence location of the deceased, gender and age bracket, from January 2020 to mid April 2020. Using death rates offers several advantages. First, alternative measures based on infections or hospital admissions suffer from biases.[6] For example, a higher number of sampling tests will inevitably show higher infection rates. In addition, in regions with better healthcare conditions (and close proximity of hospitals), usage of hospitals will be higher. Many deceased people who had shown no or mild symptoms (asymptomatics) were simply not accounted in the official figures because they are not hospitalized (e.g. due to limited capacity of hospitals). Deaths, by contrast, may be argued to be the ultimate variable of interest. Second, the collection process for death registry records minimizes reporting lags and subjectivity in recording information (e.g. residence at time of death). Third, deaths allow to capture also mortality cases that are indirectly attributable to the Covid-19 disease. For example, recent evidence shows that the number of deaths from heart attacks more than tripled during the pandemic in Italy (De Rosa et al., 2020), likely because of hospitals congestion or unavailability of ambulances.

A disadvantage of the death registry data is that it does not allow us to observe the cause of death. We thus use a statistical method to infer Covid-19-related deaths, based on deviations from the historical pattern. Specifically, we calculate excess deaths attributable to Covid-19 by deducting from a municipality's (daily) death rate the average death rate over the previous five years in the municipality using a rolling window of 7

[5] This negative correlation is not just mechanical. It persists (albeit with smaller scale) if we calculate the second shutdown exposure relative to employment in sectors that were not shut down the first time.

[6] Ciminelli and Garcia-Mandicó (2020) show that there is significant underreporting of Covid-19 deaths in Italy.

days. We scale excess deaths by the population to arrive at the following measure:

$$ExcDeathRate_{m,t} = \frac{Deaths_{m,t,2020} - avgDeaths_{m,7d,2019-2015}}{Population_m} \times 100,000 \qquad (2)$$

There is an important source of heterogeneity across municipalities: the virus hit places in Italy at different points in time. Failing to address this heterogeneity is likely to lead to an inappropriate econometric specification. In particular, a municipality that was hit early by the virus may likely display lower growth in contagion (as the curve has already levelled off) compared to a region with low virus intensity. Due to the highly non-linear dynamics of a pandemic, municipality-level fixed effects may not appropriately account for such heterogeneity. We thus condition in our empirical analysis on the "time of arrival" of the virus to compare municipalities that are at the same stage of the pandemic. We classify the time of arrival in a municipality based on two criteria: "anomaly" and "persistence". The former is measured by the day in which the cumulative excess deaths in a municipality surpasses one standard deviation of its distribution over the period January-mid April 2020. For the latter we require that the cumulative death rate among residents of a municipality m reaches a threshold of 10 deaths per 100,000 inhabitants at any point in time during the sample period. We classify the time of arrival of a virus as the day where the first criteria is met for a municipality that fulfills the second criteria (which is time-invariant). Note that according to our definition, not all municipalities are subjected to the virus (about 500 municipalities in total). We have visually inspected our classifications for a number of municipalities, and have found them to be reasonable. Studying the econometric properties, we find that on average the arrival of the virus lies in between the first and second structural break of a municipality's cumulative death rate time series.

To be included in our final dataset, we require a municipality to have been hit by the virus at some point during the first four months of 2020 as defined by our excess death measure. We also require a municipality to have at least 2500 inhabitants, to limit the influence of noise in the excess-death rate (which is scaled by population). This leaves us with 3,025 municipalities, spanning 105 provinces and 20 regions. Our sample spans the period from February 22nd to April 12th.

[Table 1 here]

Table 1 provides the summary statistics. The mean of the $ExcDeathRate_{m,t}$ is 4, that is, there are four Covid-related deaths a day per 100,000 inhabitants. The average exposure to the first shutdown, $Shutdown11_m$, is about 17.2%, whereas the average exposure to the second policy is larger (33.7%). As explained, we will focus mostly on the first shutdown; the impact of the second shutdown maybe partly confounded by the first one and hence offers a less clean setting. The table also contains information on the breakdown of sectors closed on the 11th. We can see that the food and beverage sector and the retail sector on average are about 7%, whereas the personal services sector is smaller (less than 3%).

The table also lists several control variables. The $Hospitaliz_p$ variable is the hospital capacity in a province p, measured as the sum of the number of all beds available in hospital, as a fraction of the total population (source: Ministry of Health). To study spillovers, we also include information on the shutdown exposures of the largest business centre of the provence where the municipality is located, $Shutdown11_n$ and $Shutdown25_n$. We identify the business centre as the municipality with the highest share of closed sectors, searching the municipalities with at least 15.000 inhabitants. Next, the variable $WinterTourists_p$ measures tourist intensity in a province. It is calculated as (foreign) tourists visits during January and February, scaled by population (either tourist and residents information is obtained from ISTAT). The top-10 provinces according to our tourist proxy contains skiing provinces (e.g. Trento, Bolzano, Sondrio) and historical cities (e.g. Firenze, Venice, Rome). The $WeekArrival_m$ variable is the number of weeks that lapsed between the arrival of the virus (calculated as described above) and the effective date of the first policy (March 21). We can see that on average a municipality starts experiencing the virus for the first time one month before the first-policy effective date.

3 Empirical analysis

We start with a graphical analysis of death rates across municipalities over the sample period. Figure 2 shows excess deaths, comparing municipalities with above and below the median shutdown exposure.

[Figure 2 here]

As explained before, an important heterogeneity in terms of dynamics among municipalities is the time of arrival of the virus. To take this into account, for the construction of the graph we group municipalities according to the week of the virus arrival. Within each group we calculate the average excess death rates for municipalities above and below the median shutdown exposure of this group, and average across groups.

There are three takeaways from Figure 2. First, there is no visible difference among the two groups prior to (first) treatment date, both in terms of trends but also in terms of their level. This is confirmed in Table 2, which tests for "parallel-trends" using a balance variables test, showing that mortality rates do not statistically differ - both in levels and in changes - among the two groups during the pre-treatment period. This strengthens the premises of our diff-in-diff analysis.

[Table 2 here]

Second the policy seems to become effective around date 10, consistent with our priors. Third, following the effective date, excess deaths decline in high exposure municipalities relative to low exposure municipalities.

Table 3 contains the OLS estimates for our main empirical model (equation 1). The regression includes only the municipalities that were hit by the virus (according to our statistical identification) during the sample period.

[Table 3 here]

The first column reports the results including municipality fixed effects, showing that the lagged value of mortality rates positively predict next day values. This is consistent with epidemiological models. The coefficient on the post-treatment time dummy $d11_t$ obtains a positive sign. This is explained by the fact that the policy was initiated in response to information about a rapidly spreading virus, thus around the time where contagion rates where peaking. As discussed previously, this points to the endogeneity of (national) containment policies, and reinforces the need to use within-country variation for identification. The interaction term of the treatment dummy and the shutdown variable obtains a coefficient of -0.0263 that is significant. This indicates that the shutdown was effective, as municipalities that had a higher share of sectors that were shutdown saw their death rates declining relative to other municipalities. In column 2, we add dummy variables for each weekday (from Monday to Saturday) to control for day-of-the-week effects possibly contaminating coefficient estimates. The interaction effect remains significant, with a similar coefficient.

Column 3 includes exposure to the second shutdown. The interaction term with the first shutdown increases in (absolute) size, to -0.0363. The dummy for the second shutdown obtains a negative and significant value, consistent with the second shutdown taking place at a time of a (nationwide) decline in death rates. The interaction effect with the second shutdown exposure, $Shutdown25_m$, obtains a negative value of -0.019 that is significant. This indicates that the second shutdown was effective as well in reducing death rates. It is interesting to compare the coefficients for the interaction effects on the first and second shutdown. The first shutdown obtains a coefficient that is about twice as large as the second one, indicating that the first shutdown was more effective per unit of employment. There are several interpretations to this. One is that there are declining returns to shutting down sectors, an issue to which we will return below. The final column of the table adds time-fixed effects to the regressions (the weekday dummies are redundant then and hence dropped). The results are broadly similar.

Our regression results suggest that the business shutdowns implemented in Italy reduced mortality arising from Covid-19. The size of the coefficients also shows that the effect is substantial in economic terms. We can obtain an estimate of the total effect of the (first) shutdown as follows. Given an average shutdown exposure across municipalities of 17.2% and a coefficient estimate of -0.04 (last column of Table 3), the first shutdown reduced daily deaths by 0.68 per 100,000 inhabitants. Given a total population of Italy of 60.36 millions this totals to $9,432$ lives saved over our 23-day sample period. Using estimates for the "Value of Statistical Life", we can express this into monetary terms. A common estimate for the value of one year of life in Europe is 80,000 Euros (e.g., Stadhouders et al., 2019). Considering 12 years of average remaining life of Covid-19 victims (Hanlon et al., 2020), we can calculate the monetary benefit of the policy to be $9,432 \times 80,000 \times 12yrs = 9$ Billion Euros.

The total implied benefit of the second shutdown is lower, mostly on account of the smaller event window used. Given a coefficient estimate of -0.0196, a mean shutdown exposure of 33.7% and a period of only 9 days until the end of the sample period, we obtain live savings of 3,358, amounting to 3.2 billion Euros.

3.1 Robustness

Table 4 shows that our results are robust to various modifications in the empirical model, focusing on the baseline specification of the last column of Table 3.

[Table 4 here]

The first column considers a change in the calculation of the virus arrival time. In particular, we only classify a municipality to have been hit by the virus if cumulative death rates surpass two standard deviations of its distribution (in the baseline we used one standard deviation). This leads to a more restrictive classification. Column (1) shows that the coefficient on both interaction terms are very similar, if compared to the baseline estimates.[7] In the second column, we shorten the post-treatment period to April 4th, that is, the treatment date for the second policy. This avoids any confounding effect stemming from the second policy. The coefficient now increases to -0.0435, which is consistent with Figure 2 showing that the policy is more effective earlier on. Column (3) uses logarithmic death rates, to take into account that epidemiological models predict non-linear relationships. Specifically, we replace excess death rates $y_{m,t}$ by $log(y_{m,t}+1)$. The coefficients on both policy interactions remain significant and negative. In column 4 we drop the municipalities in the most hit region (Lombardy) to investigate whether a single region is driving our results. The interaction term for the first policy remains negative and significant with a similar coefficient as in the baseline, while the interaction term for the second policy becomes insignificant. This latter is consistent with the first policy providing a clearer identification setting. In column (5) we exclude tourist regions. We expect closures of businesses (such as restaurants) to be less effective in tourist regions over the sample period, simply because Italy had already shut its borders

[7]The results are also robust to changing the "anomaly" threshold to the first decile of each municipality cumulative deaths distribution.

to tourists earlier on. Consistent with this we find that if we exclude tourist regions the coefficients on the first and second policy interactions increase in absolute terms. Columns (6) and (7) consider falsification tests, where we move both treatment dates to 20 days earlier and consider municipalities with no virus in circulation according to our methodology (totalling to 160 municipalities). Both policy interaction terms shrink substantially in size and become now insignificant. In column (8) we re-estimate the baseline model, now including time fixed-effects for each *municipality cohort* (in terms of the week of arrival of the virus). The interaction term for the first policy is significant and with very similar coefficient to the baseline, while the interation terms for the second policy is insignificant.

3.2 Contagion channels

In Table 5 we explore further the mechanism behind our baseline results (column 4 of Table 1).

[Table 5 here]

In column (1) we consider death rates of people older than 65 years, that is, among a group that is unlikely to work. We thus exclude individuals directly connected with the business (as employees or employers). Column (1) shows that the results continues to hold.[8] This points to a contagion externality (e.g., Eichenbaum et al., 2020) from business activities. This is an important result from a policy perspective because if predominantly people within a firm were infected, standard economic theory would suggests less of a need for policy interventions as any utility loss due to contagion is

[8]Similar results are obtained when we restrict the age to above 80, in which case a direct involvement in business activities becomes very unlikely.

then more likely to be internalized (in particular, workers may require higher wages to keep working during the pandemic, or simply stop turning up at work).

Column (2) explores next the impact of hospital capacity. We would expect policies to be more effective in reducing death rates in areas with congested hospitals (as virus contagion is then more likely to be fatal). Consistent with this we find that the interaction effect of hospital capacity and the first policy effect to be positive.[9] That is, in provinces where hospital capacity is less likely to be an issue, the policy effect is weakened. The interaction effect with the second policy is insignificant. This is consistent with the fact that the virus spread was already more under control at the time the second policy was enacted, and hence hospital congestion less likely to be an issue.

In column 3 of Table 5 we examine whether policy effectiveness exhibits decreasing or increasing returns. We sort municipalities into quintiles according to their share of closed down sectors, and then run a regression interacting our variable of interest with each quintile dummy. Comparing the coefficients on the shutdown exposure across the different quintiles, we see that they are monotonically declining (in absolute terms) as we move from low to high shutdown exposures. In other words, one unit of shutdown matters less in municipalities with higher shutdown exposure. This means that we face declining returns to shutting down businesses. This is fully consistent with the non-linear nature of epidemiological dynamics. In particular, once the virus is sufficiently contained, the marginal effect of reducing the reproduction rate further declines.

So far we have examined the impact of shutdowns on the municipality itself. However, larger cities also attract workers from other municipalities. We may thus expect that in this case a business shutdown has also spillovers on other municipalities. To investigate, in column (4) we include in our baseline model also an interaction effect of the largest exposed municipality in a province. In line with the idea that cities with a high sector concentration attract people from neighboring municipalities, the coefficient on the in-

[9] Though we carefully include all pair interaction terms combinations, they are omitted from the output table.

teraction between $Shutdown11_n$ and the first shutdown is negative and significant. It is similar in magnitude as the main within-municipality effect, albeit a bit smaller. This points to important spillovers from shutdowns across jurisdictions.[10]

The last column of the table decomposes the first shutdown exposure into sectors. We create exposure variables for all three sectors (food, retail and personal services) following the same approach as for the total exposure variable. The results show that all individual exposures obtain a negative coefficient (the coefficient on personal services is insignificant though, possible owing to the fact that there is little variation in this variable across municipalities, as shown in Table 1). The retail sector obtains the highest coefficient (although not statistically different from the other two coefficients). This is good news for policy makers since (brick-and-mortar) retail has a close substitute (online shopping). Thus the most effective policy may also be one that has a low cost in terms of lost consumer welfare.

4 Conclusion

This paper has examined the impact of national business shutdowns in Italy during the Covid-19 crisis. Employing a diff-in-diff setting we have found that municipalities more exposed to shutdowns experience subsequently lower mortality rates. This suggest that business shutdowns are effective in containing the spread of a virus and, ultimately, save lives. The effects are economically large and point to significant benefits from shutdowns. Our analysis suggests that shutdowns may be particularly desirable in the retail sector, both in terms of benefits and costs. Furthermore, our estimates show rapidly declining (marginal) benefits from shutdowns. Finally, our analysis points to business decisions having important (contagion) spillovers, on individuals outside the business as well as individuals in other localities. Our findings provide valuable information to policy-

[10]Beck and Wagner (2020) show that such cross-jurisdictional spillovers make uncoordinated lockdown decisions inefficient, and derive implications for international cooperation of lockdown policies.

makers involved in navigating the exit from the Covid-19 crisis, as well as for managing future pandemics.

References

Allen, L. J. (2017). A primer on stochastic epidemic models: Formulation, numerical simulation, and analysis. *Infectious Disease Modelling*, 2(2):128–142.

Beck, T. and Wagner, W. (2020). National containment policies and international cooperation. *Covid Economics: Vetted and Real-Time Papers*, 8:120–134.

Bethune, Z. A. and Korinek, A. (2020). Covid-19 infection externalities: Trading off lives vs. livelihoods. Technical report, National Bureau of Economic Research.

Ciminelli, G. and Garcia-Mandicó, S. (2020). Covid-19 in italy: an analysis of death registry data. *mimeo*.

De Rosa, S., Spaccarotella, C., Calabrò, M. P., Curcio, A., Perrone Filardi, P., Mancone, M., Mercuro, G., Muscoli, S., Nodari, S., Pedrinelli, R., Sinagra, G., and Indolfi, C. (2020). Reduction of hospitalizations for myocardial infarction in italy in the covid-19 era. *European Heart Journal*.

Eichenbaum, M., Rebelo, S., and Trabandt, M. (2020). The macroeconomics of epidemics. *NBER Working Paper No. 26882*.

Hanlon, P., Chadwick, F., Shah, A., Wood, R., Minton, J., McCartney, G., Fischbacher, C., Mair, F. S., Husmeier, D., Matthiopoulos, J., et al. (2020). Covid-19–exploring the implications of long-term condition type and extent of multimorbidity on years of life lost: a modelling study. *Wellcome Open Research*, 5.

Hartl, T., Wälde, K., and Weber, E. (2020). Measuring the impact of the german public shutdown on the spread of covid19. *Covid economics, Vetted and real-time papers, CEPR press*, 1:25–32.

Kermack, W. O. and McKendrick, A. G. (1927). A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721.

Krueger, D., Uhlig, H., and Xie, T. (2020). Macroeconomic dynamics and reallocation in an epidemic. Technical report, National Bureau of Economic Research.

Lauer, S. A., Grantz, K. H., Bi, Q., Jones, F. K., Zheng, Q., Meredith, H. R., Azman, A. S., Reich, N. G., and Lessler, J. (2020). The incubation period of coronavirus disease 2019 (covid-19) from publicly reported confirmed cases: estimation and application. *Annals of internal medicine*.

Qiu, Y., Chen, X., and Shi, W. (2020). Impacts of social and economic factors on the transmission of coronavirus disease 2019 (covid-19) in china. *Journal of Population Economics*, page 1.

Stadhouders, N., Koolman, X., van Dijk, C., Jeurissen, P., and Adang, E. (2019). The marginal benefits of healthcare spending in the netherlands: Estimating cost-effectiveness thresholds using a translog production function. *Health economics*, 28(11):1331–1344.

Ullah, A. and Ajala, O. A. (2020). Lockdown and testing effects. *Covid economics, Vetted and real-time papers, CEPR press*, 13:138–156.

A Figures

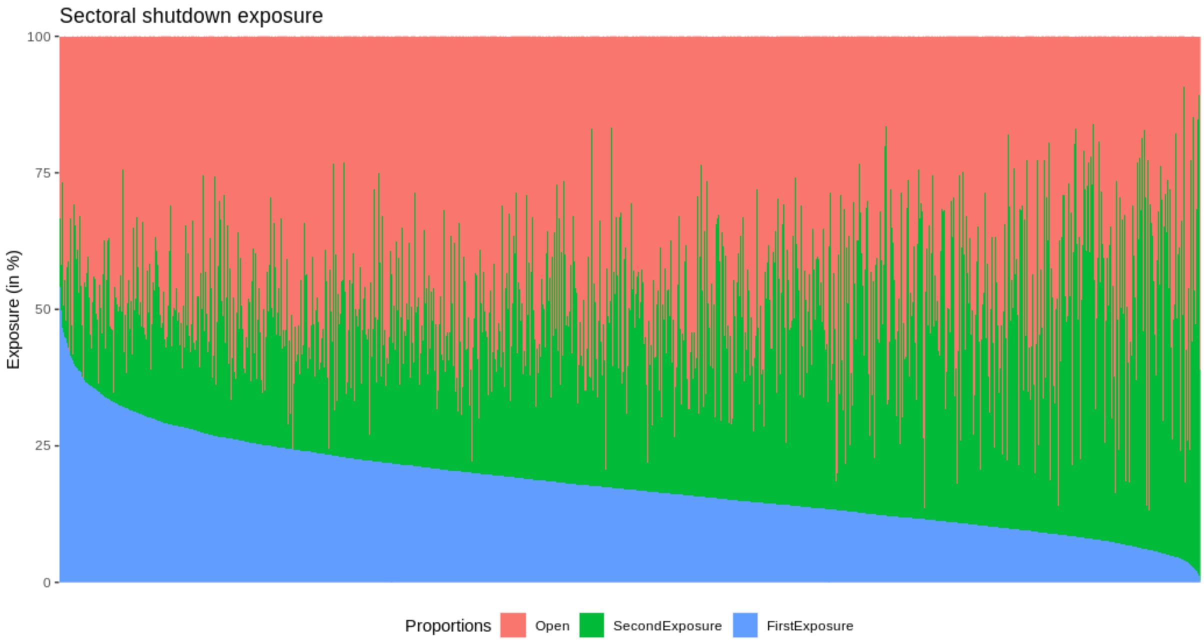

Figure 1: Sector breakdown within municipalities. The blue portion is the municipality exposure in terms of labor force to the March 11th shutdown sector list, the green one follows the March 25th decree, and the red part is the exposure to remaining (open) sectors.

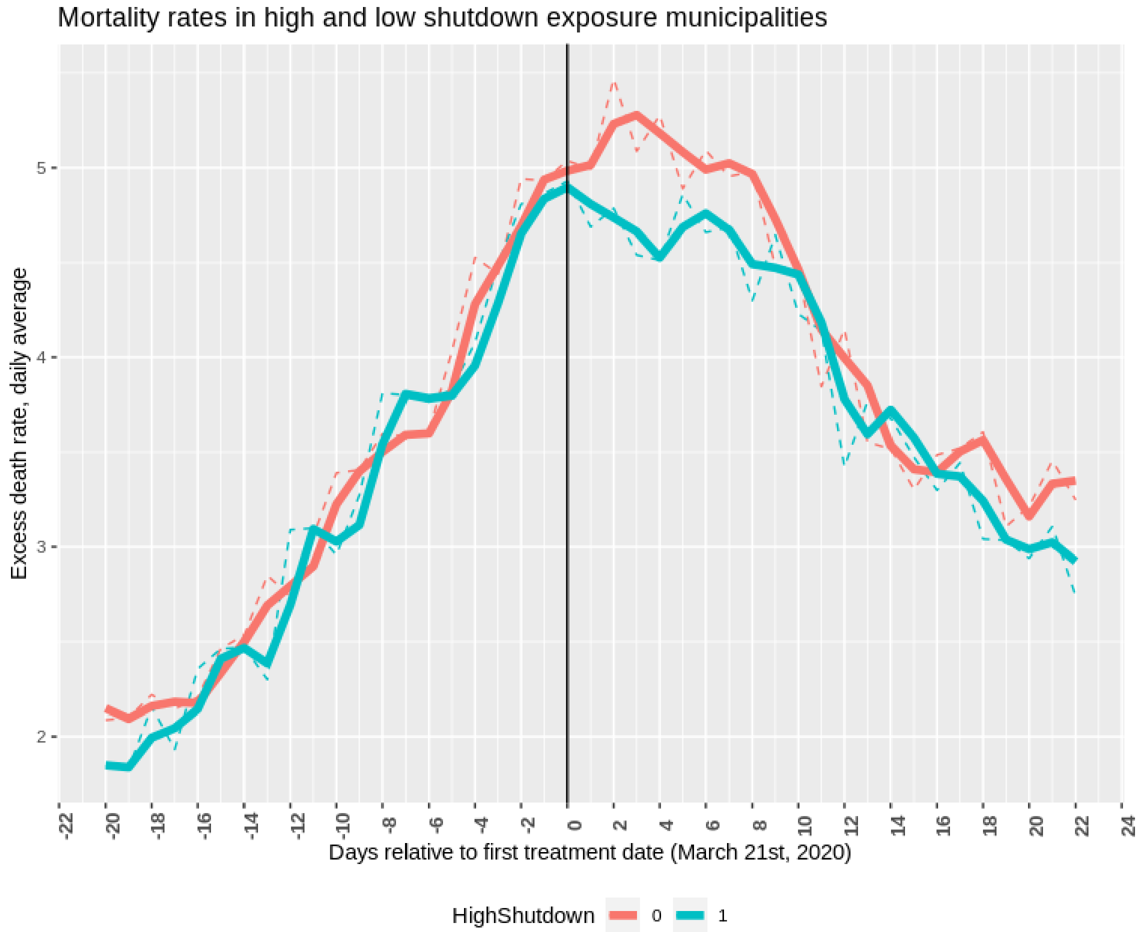

Figure 2: Average within-group excess death rates over time, around the first treatment date ($\tau = 0$ corresponds 03/21). A group excess death rate is calculated by averaging values across municipalities with above the median $Shutdown11_m$ (treated) and those below the median (control), conditional on pre-sorting municipalities on the virus arrival week.

B Tables

Table 1: Summary statistics

Variable	Definition	Mean	Standard Dev.	P5	Median	P95	Observat.
$ExcDeathRate_{m,t}$	Deaths in day t over 2020, minus 7-day average around t over previous five years, divided by municipality m residents;	3.947	9.43	0	0	24.066	130,231
$\Delta ExcDeathRate_{m,t,t-1}$	Daily change in $ExcDeathRate_{m,t}$;	-.076	12.11	-20.784	0	20.427	130,231
$log(1 + ExcDeathRate_{m,t})$	Natural logarithm of $ExcDeathRate_{m,t} + 1$;	.660	1.16	0	0	3.222	130,231
$log(Pop_m)$	Natural logarithm of municipality m residents;	8.846	.74	7.903	8.727	10.191	130,231
$d11_t$	Dummy variable, taking a value of one from ten days after first policy (03/21), and zero before;	.434	.49	0	0	1	130,231
$d25_t$	Dummy variable, taking a value of one from ten days after second policy (04/05), and zero before;	.133	.34	0	0	1	130,231
$Shutdown11_m$	number of employees and employers of sectors shutdown on March 11th divided by the total labor force in a municipality m	17.173	7.32	7.190	16.152	30.140	130,231
$Shutdown25_m$	number of employees and employers of sectors shutdown on March 25th divided by the total labor force in a municipality m	33.720	15.06	12.782	31.841	61.508	130,231
$Food11_m$	employees and employers in the food sector divided by the total labor force in a municipality m	7.265	4.54	2.383	6.113	16.220	130,231

(Continued)

Variable	Definition	Mean	Standard Dev.	P5	Median	P95	Observat.
$Retail11_m$	employees and employers in the retail sector divided by the total labor force in a municipality m	7.081	3.61	2.512	6.551	13.613	130,231
$Personal11_m$	employees and employers in the personal-services sector divided by the total labor force in a municipality m	2.842	1.79	.948	2.523	5.583	130,231
$Hospitaliz_p$	Number of beds available in all hospitals of province p, divided by residents in province p per 100,000 people	.686	.38	.227	.597	1.354	130,231
$Shutdown11_n$	$Shutdown11_m$ of the largest exposed municipality m, with at least 15,000 inhabitants, in the province	22.988	5.38	16.411	22.636	32.175	126,505
$Shutdown25_n$	$Shutdown25_m$ of the largest exposed municipality m, with at least 15,000 inhabitants, in the province	40.665	15.28	17.540	41.128	67.739	126,505
$WinterTourists_p$	Number of tourists in all types of accommodations in a province p, over January-February 2007 (ISTAT)	5.468	12.74	.275	2.434	25.889	126,028
$WeekArrival_m$	Time from the virus arrival date in a municipality m to March 21th, measured in number of weeks	5.713	2.78	1	6	10	130,129

Note: This table shows the definition, mean, standard deviation, the 5th, 50th (median) and 95th percentile, and number of observations for each variable adopted in the empirical analysis. The sample consists of ISTAT municipalities reporting death data among residents in the February 22nd - April 12th period. Municipalities with less than 2,500 inhabitants and those that, according to our definition, were not hit by the virus are excluded.

Table 2: Parallel trend (Balanced variables test)

Variables	Mean Low $Shut11_m$	Mean High $Shut11_m$	Difference	T-test
$ExcDeathRate_m$	3.978	3.570	.4077	1.29
$\Delta ExcDeathRate_m$	.160	-.468	.62	1.49
$ExcDeathRate_m$ growth	.01%	-.05%	.06%	1.45
$log(Pop_m)$	8.81	8.87	-.051	-1.87*
$WeekArrival_m$	5.57	5.56	.006	.058

Note: Average daily characteristics of municipalities with above median shutdown exposure to March 11th sector list compared to those with values below median of each group. Municipalities are assigned into groups according to their virus arrival week. Within-group average values are calculated through an 8-days period surrounding the 03/11 policy, that is from 03/07 to 03/15. Municipalities with less than 2500 residents and those not showing any "anomaly" nor "persistence" in cumulative excess death rates are excluded.

Table 3: Main analysis

LHS: DailyDeathRate	(1)	(2)	(3)	(4)
$y_{m,t-1}$	0.0399***	0.0398***	0.0363***	0.0335***
	(6.49)	(6.48)	(6.03)	(5.65)
$d11_t$	1.009***	1.087***	0.851***	
	(5.55)	(5.95)	(4.53)	
$d11_t \times Shutdown11_m$	-0.0263***	-0.0264***	-0.0363***	-0.0403***
	(-3.28)	(-3.29)	(-4.36)	(-4.83)
$d25_t$			-1.314***	
			(-5.41)	
$d25_t \times Shutdown25_m$			-0.0190***	-0.0196***
			(-2.76)	(-2.85)
Municipality FE	✓	✓	✓	✓
Time FE	x	x	x	✓
Weekdays	x	✓	✓	x
Obs.	130,231	130,231	130,231	130,231
R^2	0.151	0.151	0.154	0.157
Adj. R^2	0.131	0.131	0.134	0.137

This table presents difference-in-differences estimates of municipality-level panel regressions of daily excess death rates (left-hand side). The lagged dependent variable $y_{m,t-1}$ is included in the model. $d11_t$ and $d25_t$ are dummy variables that take a value of one in the days after the first and second policy becoming effective, respectively. $Shutdown11_m$ and $Shutdown25_m$ are employment exposures of municipality m to the shutdown policies of 03/11 and 03/25, respectively. The sample consists in ISTAT death registry data over the period 02/22/2020-04/12/2020. We exclude municipalities with less than 2,500 inhabitants and those that were not hit by the virus, leaving 3,025 municipalities in total. t statistics in parentheses. Standard Errors clustered at municipality-level. *, ** and *** represent statistical significance at the 10%, 5% and, 1% level, respectively.

Table 4: Robustness tests

	Arrival Time	Shorter Window	Log-Linear	Exclude Lombardy	Exclude Touristic	Placebo		Arrival Week
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
LHS:	$y_{m,t}$	$y_{m,t}$	$\log(y_{m,t}+1)$	$y_{m,t}$	$y_{m,t}$	$y_{m,t}$	$y_{m,t}$	$y_{m,t}$
$y_{m,t-1}$	-0.00682	0.0182***	0.0116***	-0.00495	0.0410***	-0.0356***	0.00471	0.146***
	(-1.22)	(3.28)	(3.55)	(-1.01)	(6.26)	(-6.97)	(0.28)	(2.83)
$\text{Shutd11}_m \times d11_t$	-0.0279***	-0.0435***	-0.00672***	-0.0233***	-0.0394***	-0.00785	-0.00215	-0.0397***
	(-3.27)	(-4.95)	(-6.35)	(-2.80)	(-4.11)	(-0.99)	(0.28)	(-4.69)
$\text{Shutd25}_m \times d25_t$	-0.0146***		-0.00168**	0.00557	-0.0264***	0.00812	0.0119	-0.00148
	(-2.28)		(-2.21)	(0.98)	(-3.41)	(1.49)	(1.56)	(-0.25)
Municipality FE	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓	
Time×ArrWeek FE								✓
Obs.	130,231	115,111	130,231	100,443	104,751	72,186	6,274	130,157
R^2	0.157	0.182	0.131	0.071	0.168	0.099	0.055	0.199
Adj. R^2	0.137	0.160	0.110	0.049	0.148	0.061	0.023	0.157

This table conducts robustness checks focusing on the fourth column of the baseline model shown in table 3. Including its lagged values, the dependent variable is the daily death rate in municipality m. $d11_t$ and $d25_t$ are dummy variables that take value of one in the days after the first and second policy becoming effective, respectively. $Shutdown11_m$ and $Shutdown25_m$ are municipality m exposure to the first and second policy, respectively. First column reports results OLS estimates changing definition of "anomaly" in arrival time to 2 standard deviation of cumulative excess death rate. "Shorter window" shrinks sample period to 04/04. Column 3 applies the natural logarithmic function to excess death rates. Columns 4 and 5 exclude Lombardy and touristic municipalities, respectively. Columns 6 and 7 ("Placebo") shifts both policy treatment dates twenty days earlier and considers municipalities with no virus in circulation, respectively. Column 8 adds time interacted with week-of-arrival-cohort fixed effects. The sample consists in ISTAT death registry data over the period 02/22/2020-04/12/2020. We exclude municipalities with less than 2,500 inhabitants throughout specifications, and those that were not hit by the virus (except for column 7). t statistics in parentheses. Standard Errors clustered at municipality-level. *, ** and *** represent statistical significance at the 10%, 5% and, 1% level, respectively.

Table 5: Contagion channels

LHS: DailyMortalityRates	Elderly (1)	Hospitaliz. (2)	Decreas.Eff. (3)	Spillovers (4)	Sector Decompos. (5)
$y_{m,t-1}$	0.0336***	0.0332***	0.0332***	0.0327***	0.0335***
	(5.51)	(5.62)	(5.59)	(5.47)	(5.65)
$\text{Shutd11}_m \times d11_t$	-0.0389***	-0.0831***		-0.0341***	
	(-4.81)	(-3.67)		(-3.20)	
$\text{Shutd25}_m \times d25_t$	-0.0160**	-0.0173	-0.0222***	-0.0218***	-0.0199***
	(-2.38)	(-1.08)	(-3.21)	(-2.68)	(-2.88)
$\text{Shutd11}_m \times d11_t \times Hospitaliz_p$		0.0580**			
		(2.39)			
$\text{Shutd25}_m \times d25_t \times Hospitaliz_p$		-0.0066			
		(-0.25)			
$\text{Shutd11}_m \times d11_t \times Q1_m$			-0.123***		
			(-4.57)		
$\text{Shutd11}_m \times d11_t \times Q2_m$			-0.121***		
			(-3.86)		
$\text{Shutd11}_m \times d11_t \times Q3_m$			-0.0850***		
			(-2.95)		
$\text{Shutd11}_m \times d11_t \times Q4_m$			-0.0818***		
			(-2.64)		
$\text{Shutd11}_m \times d11_t \times Q5_m$			-0.0316		
			(-1.28)		
$\text{Shutd11}_n \times d11_t$				-0.0269**	
				(-2.25)	
$\text{Shutd25}_n \times d25_t$				-0.000207	
				(-0.03)	
$\text{food11}_m \times d11_t$					-0.0311**
					(-2.25)
$\text{retail11}_m \times d11_t$					-0.0558***
					(-2.87)
$\text{personal11}_m \times d11_t$					-0.0341
					(-0.95)
Municipality FE	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓
Obs.	128,335	130,231	124,303	130,231	130,231
R^2	0.153	0.158	0.155	0.158	0.157
Adj. R^2	0.132	0.137	0.134	0.137	0.137

This table explores different contagion channels focusing on the fourth column of the baseline model shown in table 3. Including its lagged values, the dependent variable is the daily excess death rate in municipality m. $d11_t$ and $d25_t$ are a dummy variables that take value of one in the days after the first and second policy, respectively, become effective. $Shutdown11_m$ and $Shutdown25_m$ are municipality m exposure to the first and second policy, respectively. $Hospitaliz_p$ is the ratio of beds in all hospitals in a province, divided by total residents. $Q1-Q5$ are dummy variables for each quintile of $Shut11_m$. $Shut11_n$ and $Shut25_n$ are exposures to first and second policies, respectively, of the largest hit municipality in the province. $food11_m$, $retail11_m$ and $personal11_m$ are the exposure of a municipality m to the food, retail and personal-services industry, respectively. All pairwise interaction terms are included yet omitted in the table. You can refer to table 1 for variable definitions, sample period and data. t statistics in parentheses. Standard Errors clustered at municipality-level. *, ** and *** represent statistical significance at the 10%, 5% and, 1% level, respectively.