An empirical estimate of the infection fatality rate of COVID-19 from the first Italian outbreak

Gianluca Rinaldi^{☆a}, Matteo Paradisi^{☆b}

^a*Harvard Economics Department*

^b*PhD. NBER and Einaudi Institute for Economics and Finance (EIEF)*

Abstract

Background: The coronavirus 2019 (COVID-19) pandemic has been spreading globally for months, yet the infection fatality ratio of the disease is still uncertain. This is partly because of inconsistencies in testing and death reporting standards across countries. We provide estimates which don't rely on official cases and deaths data but only on population level statistics.

Methods: We collected demographic and death records data from the Italian Institute of Statistics. We focus on the area in Italy that experienced the initial outbreak of COVID-19 and estimated a Bayesian model fitting age-stratified mortality data from 2020 and previous years. We also assessed the sensitivity of our results to alternative assumptions on the proportion of population infected. Based on our estimates we finally studied the heterogeneity in overall lethality across countries.

Findings: We estimate an overall infection fatality rate of 1.31% (95% credible interval [CrI] 0.94 – 1.89), as well as large differences by age, with a low infection fatality rate of 0.05% for under 60 year old (CrI 0 – 0.17) and a substantially higher 4.16% (CrI 3.05 – 5.80) for people above 60 years of age. In our sensitivity analysis, we found that even under extreme assumptions, our method delivered useful information. For instance, even if only 10% of the population were infected, the infection fatality rate would not rise above 0.2% for people under 60. Finally, using data on demographics we show large expected heterogeneity in overall IFR across countries.

Interpretation: Our empirical estimates show a sharp difference in fatality rates between young and old people and rule out overall fatality ratios below 0.5% in populations with more than 30% over 60 years old.

[☆]Contributed Equally. The authors can be reached at rinaldi@g.harvard.edu and paradisi@nber.org.

May 12, 2020

1. Introduction

Estimating the severity of the coronavirus disease 2019 (COVID-19), caused by the novel severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2) across demographic groups is of medical interest, but will also be crucial in the design of health and economic policies to deal with the next phases of the pandemic.

Several approaches to measure lethality have been proposed since the start of the outbreak^{1,2,3,4}, with most of the relevant population infected. Given the mounting evidence that a substantial proportion of infected people are either asymptomatic or display very mild symptoms^{5,6}, as well as the difficulties many countries are encountering in ramping up testing, it has been difficult to obtain precise estimates of the total number of infected. Testing mainly symptomatic cases on the basis of clinical studies on the symptoms of COVID-19^{7,8,9,10,11}, as announced for instance by the Italian government, might also have led to underestimating the total number of infected. While measurement of deaths is reliable, statistically significant deviations in total deaths relative to previous years have been observed in the most affected areas, leading to concerns that official COVID-19 death counts might also be underestimated in some cases¹². Together, those hurdles make estimating the true infection fatality rate challenging.

To sidestep some of those issues, in this article we used an empirical approach employing publicly available aggregate deaths and demographic data to obtain infection fatality ratio estimates without relying on official data on COVID-19 positive cases and deaths. The key observation is that, assuming an accurate measurement of fatalities in a population, infection fatality ratio estimates are less strongly dependent on accurate measurement of total cases when the share of population infected is larger: keeping the number of deaths fixed, the estimate changes much less when the fraction of the population infected varies from 40% to 60% than when it changes from 2% to 3%, simply because ratios are nonlinear.

We were therefore able to obtain precise fatality estimates by age range focusing on one of the hardest hit areas in Lombardy, which was placed under lock-down order already on February 21st 2020. This area includes ten towns and has a population of around 50 thousands people. The first recorded patient infected through community spread in Italy was admitted

to intensive care on February 20th and the area already had 36 confirmed cases the day after¹³. While widespread randomized antibody testing hasn't yet been conducted in this area, 30% of a sample of blood donors from all ten municipalities tested positive to antibody¹⁴, and a smaller sample of 60 asymptomatic blood donors in one of the towns under lock-down showed 40 (66%) positive cases¹⁵. Although these samples might not be fully representative of the population of these areas, this figure is highly suggestive of a widespread contagion.

2. Methodology

2.1. Data

We focused on ten Italian municipalities in Lombardy that experienced the initial outbreak of COVID-19. Data on deaths has been collected from the Italian Institute of Statistics (ISTAT)¹⁶. We built estimates of total death counts based on daily data on recorded deaths for 2020 and previous years for the period 2015-2019 that ISTAT collects from the *Anagrafe Nazionale della Popolazione Residente* (National Census of Resident Population). The data contains information about gender and age group (in 5-years bins) for each recorded death until April 15th 2020. ISTAT has released data for only eight of the ten municipalities that experienced the initial outbreak until April 15th, while information on nine municipalities is available until April 4th. Hence, we focused our baseline analysis on the latter subsample, while using the smaller sample until April 15th to test the robustness of our results. The excluded town in the baseline sample is much smaller than the others: its total population in 2019 was 1,118, while the remaining nine municipalities had a total population of 49,431. For one of the municipalities (Casalpusterlengo) we used the most recent data from¹⁶ available until May 31st, while we drew information for the period April 1st to April 4th from a previous data release¹⁷.

Death counts data has been complemented with information provided by ISTAT on the demographics of each municipality¹⁸. For every city we collected total population by age range in each year from 2015 to 2019. We used year 2019 information for year 2020 since data on 2020 has not been released yet. The municipality of Castelgerundo was created in 2018 by the union of the municipalities of Camairago and Cavacurta. We therefore used total population counts from the latter two for the years 2015-2018.

Finally, data on demographics for OECD countries was collected¹⁹. Unfortunately, OECD age groups do not perfectly overlap with the age bins in the Italian death data. Indeed, there is a one year difference: OECD age groups are, for instance, 20 – 24, while the Italian deaths data uses bins such as 21 – 25. Since we aggregate across multiple bins this is not a major concern, for instance the 21 – 40 deaths data bin corresponds to 20 – 39 in the OECD data.

2.2. Comparing 2020 deaths to previous years

In order to motivate the use of administrative death counts for our infection fatality rate estimation, we begin by analysing patterns in overall mortality in 2020 and previous years. In particular, on every day between February 21st (the beginning of the outbreak) and April 4th the difference between the number of 2020 deaths and the 2015-2019 average has been computed for total population and for different age groups. We compared total deaths to average number and fluctuations in deaths in past years to assess the signal to noise ratio of this measure.

2.3. Bayesian Estimation of COVID-19 infection fatality rate

We employ a Bayesian framework to estimate the infection fatality rate of COVID-19 by adapting a standard binomial mortality model to our setting²⁰. The likelihood function of the model is obtained by assuming deaths in the period between February 28th and April 4th of each year are binomially distributed according to:

$$D_{i,a,y} \sim \text{Binomial}(\delta_a, N_{i,a,y}) \quad \text{for } y \in \{2015, \dots, 2019\} \quad (1)$$

$$D_{i,a,2020} \sim \text{Binomial}(\delta_a + \delta_a^{\text{Covid}} \cdot \theta_i, N_{i,a,2019}) \quad (2)$$

where i denotes the municipality, y the year, and a the age range. We used seven age ranges: 0 – 20, 21 – 40, 41 – 50, 51 – 60, 61 – 70, 71 – 80, 81+. $D_{i,a,y}$ and $N_{i,a,y}$ are the total deaths and population in town i , year y and age range a , respectively. For population in 2020, we use 2019 data since 2020 data is not yet available. The baseline lethality rates δ_a are heterogeneous across age ranges, but were assumed to be constant across municipalities and years.

δ_a^{Covid} is the infection fatality rate for age range a and we assumed $\delta_{a,t}^{\text{Covid}} = 0$ in every year before 2020 when COVID-19 was not present. We also assumed

that infection rates θ_i are heterogeneous across municipalities but constant across age groups.

We assumed the following priors for the parameters of interest:

$$\delta_a \sim \text{Uniform}[0, 0.1] \quad (3)$$

$$\delta_a^{\text{Covid}} \sim \text{Uniform}[0, 0.3] \quad (4)$$

$$\theta_i \sim \text{Beta}(3, 2). \quad (5)$$

Priors on baseline and COVID-19 death rates were chosen to be uninformative, while we chose the prior on infection rates to reflect the results of the antibody testing in one of the municipalities¹⁵, while at the same time maintaining a weakly informative prior.

We implemented a Bayesian procedure to derive point estimates and credible intervals for the infection fatality rates. The model was estimated using Markov Chain Monte Carlo (MCMC). We calculated the median and 95% credible interval using to the quantiles of the posterior distribution for all parameters. To check the sensitivity of our estimates we also calculated point estimates as the mode of the posterior distribution and confidence intervals as 95% highest posterior density interval²¹.

We fitted our model using R version 3.6.2. We drew 100,000 samples from the joint posterior distribution and used 10 independent chains, discarding the first 1000 samples for each chain. Trace plots of the Markov Chain Monte Carlo as well as posterior distributions for each variable are reported in Appendix A. All analyses are fully reproducible with the code available online.

Because our Bayesian procedures relies on modelling assumptions to derive the infected portion of the population, we also implemented a more agnostic approach, showing how infection fatality rates vary by contagion rate. In this exercise we computed the infection fatality rates from a simplified version of the model above in which we set a degenerate prior for each of the θ_i to be equal to a constant in the interval $[0.1, 1]$ and estimated the model for each choice.

2.4. Projections of IFR Across Countries

We used the estimates of our Bayesian model to draw projections on the overall lethality in OECD countries. We built these projections assuming homogeneous contagion across age groups and using each country's demographic composition to re-weight our IFR estimates through linear combinations. Point estimates in this exercise were produced combining the medians

of the posteriors of different age groups, while 95% credible intervals were obtained with linear combinations of the extremes of the 95% credible intervals of the age-specific IFRs.

3. Results

3.1. Raw Deaths Counts

We documented a substantial increase in total deaths at the beginning of the outbreak. Figure 1 shows the total daily deaths counts in the nine municipalities for 2020 and 2015-2019 (average) in the period between Jan 1st and April 4th. The year 2020 and previous years are very similar preceding the last week of February, when the first COVID-19 cases have been discovered. Starting that week, we observed a spike in the number of deaths. This spike is clearly related to COVID-19 as it starts at the beginning of the outbreak and it significantly overcomes average fluctuations in deaths observed in previous years. In total, deaths over the period from February 21st to April 4th in 2020 were almost five times the average in previous years over the same period of time (371 vs 73).

3.2. Infection Fatality Ratio Estimates

Using our Bayesian model we estimated an overall infection fatality ratio of 1.31% (95% credible interval [CrI] 0.94–1.89). We also uncovered substantial heterogeneity by age. For under 60 years old the infection fatality rate was 0.05% (CrI 0 – 0.17, Table 2), while for over 60 years old it was 4.16% (CrI 3.05 – 5.80, Table 2).

Figure 2 shows the estimated infection fatality ratios by age group together with 95% credible intervals and interquartile range. As expected, infection fatality ratios were found much larger for older age groups. Point estimates are 4.54 and 8.50 for 71-80 and 81+, respectively. We cannot exclude, however, that the infection fatality ratio for over 80 years old is as high as 11.57% or as low as 6.39%. Interestingly, we found an infection fatality rate close to zero for under 50 years old, and around 0.1% for the 51-60 group. For robustness, we have also recalculated all point estimates using the mode and confidence intervals as the highest posterior density interval. Results remained virtually identical (data not shown). Moreover, we checked that results remained unchanged when extending the period of analysis until April 15th and looking at the eight municipalities for which we have data until this date (data not shown).

Estimated infections rates were also heterogeneous by town, ranging between 24.3% and 84.7% (Table 2). Interestingly, Castiglione d’Adda, where antibody tests conducted on a sample of individuals detected a 66.6% infection rate, resulted as the municipality with the largest share of the population infected (84.7%). We estimate a population weighted overall infection rate for the nine towns of 45.58% (CrI 30.56% – 62.88%). This is broadly consistent with a recent study on blood donors for the entire area¹⁴ has found a 30% overall infection rate.

We then performed an exercise to assess how sensitive our infection fatality rate estimates are to different levels of contagion (Figure 3). Focusing on a large range of potential infection rates, we found that even in the conservative assumption that only 15% of the population was infected, under 60 years old still experienced an infection fatality rate significantly below 1%, with under 40 being around 0.1%. Obviously, estimates spiked for older age groups as we set the infection rate very low. These results confirm the view that the infection has low lethality rates for younger individuals, but large rates for the elders. Moreover, this exercise showed that the overall infection fatality rate was significantly above zero and around 0.5% even in the most conservative assumption of 100% contagion.

Finally, we used our estimates to build overall IFRs for OECD countries. We re-weighted the IFR of each age group by its respective share of the total population of each country. This procedure relies on the assumption that contagion was homogeneous across age groups. Figure 4 shows the results ranking countries based on their projected IFR. Given their high shares of over 60 years old, Japan and Italy are the countries with the highest IFR, with point estimates of 1.44% and 1.19%, respectively. The Europe-28 IFR is slightly above 1%. Interestingly, the U.S. are projected to have a IFR below 1% and equal to 0.8% if contagion spread homogeneously across age groups. Very few developed countries are expected to have a IFR below 0.5% in the scenario of homogeneous infection across ages.

4. Discussion

In this paper we estimated the infection fatality rate of COVID-19 from administrative death counts on nine Italian municipalities that experienced the first outbreak of this disease in late February 2020. We found an overall infection fatality rate of 1.31% (CrI 0.94 – 1.89). We uncovered significant heterogeneity across age groups. Under 60 years old have infection fatal-

ity ratios around 0.05% (CrI 0.00 – 0.17). On the contrary, older people are at significantly larger risk: over-60 infection fatality ratio is 4.16% (CrI 3.05 – 5.80), and for over-80 it is estimated at 8.50% (CrI 6.4 – 11.57). We excluded very low fatality rates even under very conservative assumptions on the population infection rate: if the entire population had been infected, the overall infection fatality ratio would still be around 0.5%. Finally, we showed projections of overall IFRs across countries in the scenario of equal contagion across age groups. Japan and Italy are the countries expected to suffer the most with IFRs of 1.44% and 1.19%, while the US have a smaller IFR equal to about 0.8%.

Our result for overall infection fatality ratio is larger than estimates in²² based on travellers' data. It is, however, remarkably close to estimates for two case studies where the entire population was tested: Diamond Princess Cruise (1.3%)²³, and Vo' Euganeo (1%) - a 3,000 inhabitants municipality in the Italian Veneto region. Despite the similarity there are two main reasons to expect differences between our estimates and those in the two mentioned case studies. First, we estimate a large infection fatality ratio for over 80 years old who are likely underrepresented in the Diamond Princess Cruise. Second, total population in both Diamond Princess Cruise and Vo' Euganeo is limited and for this reason those estimates might be affected by under-sampling of the number of infected in the tails. These small differences notwithstanding, all these figures are substantially lower than estimated case fatality rates (CFRs) computed with official contagion data and for previous influenza pandemics²⁴. The most plausible reason is that COVID-19 case counts are subject to downward biases due to limited testing capacity and testing strategies that prioritize symptomatic cases despite a large number of asymptomatic patients.

Our estimates could suffer from the fact that deaths data is missing for the period after April 4th 2020. This would be an issue if the contagion had not stopped by that date and therefore useful information on deaths could not be used in our model. However, as the trend in number of deaths suggests, the contagion likely stopped by April 4th in these nine municipalities. Indeed, the number of deaths in the last days of our sample went back to the average number in the previous five years. Moreover, using data available until April 15th for eight of the nine municipalities, we checked that the number of deaths remained aligned to the average in previous years between April 4th and April 15th. For this reason, there are reasons to believe that data for the following weeks would not be very informative on the lethality of the first

wave of contagion.

One limitation of our estimates is that they should not be taken at face value in the analysis of contexts where the hospital system is under stress or close to capacity. This is because our exercise was performed on a case study at the beginning of the Italian outbreak when the hospital system was still fully functioning. Moreover, the quarantine measures implemented in the nine municipalities on February 21st likely reduced contagion, potentially affecting infection fatality ratios and making them hard to extrapolate to context where similar measures were not undertaken.

Another limitation is that our model assumed a constant baseline lethality rate in absence of COVID-19. This implies that the COVID-19 outbreak did not change the baseline death rate in the population. Plausibly, the lockdown policy decreased deaths from, among others, violence and traffic, while at the same time the outbreak could have increased other fatalities due to lower availability of healthcare resources for other diseases. Fluctuations due to these causes are, however, likely to be quantitatively small compared to the large spike in deaths that we observed in 2020, where total deaths were five times the average in previous years over the same period of time (371 vs 73). For this reason, confounding factors in baseline deaths levels should not have substantially affected the signal to noise ratio of our death data.

In conclusion, our results support the need for isolating policies especially for the elder part of the population given the extremely high fatality rates from COVID-19 that we estimated.

Contributors

GR conceived the study with input from MP. MP collected demographic and deaths data and cross checked the accuracy of the data. GR developed the code for the Bayesian model and estimated it with input from MP. MP produced the first draft of the manuscript. All authors contributed to the final draft.

Declaration of interests

GR and MP declare no competing interests.

Data Sharing

All data is publicly available and can be found at the cited sources. Code for replication is available at https://github.com/gianlucaRinaldi/covid_IFR_Lombardy.

References

- [1] Verity R, Okell LC, Dorigatti I, Winskill P, Whittaker C, Imai N, et al. Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet Infectious Diseases*. 2020;.
- [2] Baud D, Qi X, Nielsen-Saines K, Musso D, Pomar L, Favre G. Real estimates of mortality following COVID-19 infection. *The Lancet infectious diseases*. 2020;.
- [3] Kim DH, Choe YJ, Jeong JY. Understanding and Interpretation of Case Fatality Rate of Coronavirus Disease 2019. *J Korean Med Sci*. 2020 Mar;35(12). Publisher: The Korean Academy of Medical Sciences. Available from: <http://synapse.koreamed.org/DOIx.php?id=10.3346%2Fjkms.2020.35.e137>.
- [4] Villa M, Myers JF, Turkheimer F. COVID-19: Recovering estimates of the infected fatality rate during an ongoing pandemic through partial data. *medRxiv*. 2020;Available from: <https://www.medrxiv.org/content/early/2020/04/14/2020.04.10.20060764>.
- [5] Zhou X, Li Y, Li T, Zhang W. Follow-up of asymptomatic patients with SARS-CoV-2 infection. *Clinical Microbiology and Infection*. 2020;Available from: <http://www.sciencedirect.com/science/article/pii/S1198743X20301695>.
- [6] Quilty BJ, Clifford S, group2 Cnw, Flasche S, Eggo RM. Effectiveness of airport screening at detecting travellers infected with novel coronavirus (2019-nCoV). *Eurosurveillance*. 2020;25(5). Available from: <https://www.eurosurveillance.org/content/10.2807/1560-7917.ES.2020.25.5.2000080>.
- [7] Chan JFW, et al. A familial cluster of pneumonia associated with the 2019 novel coronavirus indicating person-to-person transmission: a study of a family cluster. *Lancet*. 2020;395(10223):514–523.

- [8] Chen N, et al. Epidemiological and clinical characteristics of 99 cases of 2019 novel coronavirus pneumonia in Wuhan, China: a descriptive study. *Lancet*. 2020;395(10223):507–513.
- [9] Guan WJ, et al. Clinical characteristics of coronavirus disease 2019 in China. *N Engl J Med*. 2020;.
- [10] Huang C, et al. Clinical features of patients infected with 2019 novel coronavirus in Wuhan, China. *Lancet*. 2020;395(10223):497–506.
- [11] Li Q, et al. Early transmission dynamics in Wuhan, China, of novel coronavirus-infected pneumonia. *N Engl J Med*. 2020;.
- [12] Galeotti A, Surico P. Deaths from Covid-19 - lessons from Italy and other nations. *londonedu*. 2020 04;Available from: <https://www.london.edu/news/the-true-death-toll-of-covid-19-1771>.
- [13] Grasselli G, Pesenti A, Cecconi M. Critical Care Utilization for the COVID-19 Outbreak in Lombardy, Italy: Early Experience and Forecast During an Emergency Response. *JAMA*. 2020 03;Available from: <https://doi.org/10.1001/jama.2020.4031>.
- [14] Il Giorno Pavia, editor. Coronavirus, Baldanti: "Test sierologico? Non parlerei di patente di immunità"; 2020. <https://www.ilgiorno.it/> (accessed on 04/18/2020).
- [15] La Stampa, editor. Castiglione d'Adda è un caso di studio: 'Il 70% dei donatori di sangue è positivo'; 2020. <https://www.lastampa.it/> (accessed on 04/15/2020).
- [16] ISTAT. Dataset analitico con i decessi giornalieri in ogni singolo comune di residenza (accessed on 05/6/2020). 2020 04;Available from: <https://www.istat.it/it/archivio/240401>.
- [17] ISTAT. Dataset Analitico con i Decessi Giornalieri (accessed on 04/14/2020). 2020 04;Available from: <https://www.istat.it/it/archivio/240401>.
- [18] ISTAT. Demografia in Cifre. <http://demoistatit/> (accessed on 04/13/2020). 2020 04;.

- [19] OECD. Historical population (accessed on 05/10/2020); Available from: https://stats.oecd.org/OECDStat_Metadata/ShowMetadata.ashx?Dataset=HISTPOP&Lang=en.
- [20] Manton KG, Woodbury MA, Stallard E, Riggan WB, Creason JP, Pelom AC. Empirical Bayes Procedures for Stabilizing Maps of U.S. Cancer Mortality Rates. *Journal of the American Statistical Association*. 1989;84(407):637–650. PMID: 12155378. Available from: <https://doi.org/10.1080/01621459.1989.10478816>.
- [21] McElreath R. Statistical rethinking, texts in statistical science. Boca Raton: CRC Press; 2015.
- [22] Niehus R, De Salazar PM, Taylor AR, Lipsitch M. Using observational data to quantify bias of traveller-derived COVID-19 prevalence estimates in Wuhan, China. *The Lancet Infectious Diseases*. 2020;.
- [23] Russell T, Hellewell J, Jarvis C, van Zandvoort K, Abbott S, Ratnayake R, et al. Estimating the infection and case fatality ratio for coronavirus disease (COVID-19) using age-adjusted data from the outbreak on the Diamond Princess cruise ship, February 2020. *Eurosurveillance*. 2020;25:123–456.
- [24] Wilson MB Nih. The emerging influenza pandemic: estimating the case fatality ratio. *Eurosurveillance*. 2009;25(5).

List of Tables

Table 1: Total population and deaths by age range

Age Range	Total Population					Total Deaths					
	2015	2016	2017	2018	2019	2015	2016	2017	2018	2019	2020
0 – 20	8452	8447	8420	8438	8460	0	0	0	0	0	1
21 – 40	10777	10616	10374	10233	10217	0	0	1	0	0	0
41 – 50	8407	8239	8053	7884	7630	1	2	0	1	2	2
51 – 60	7258	7480	7646	7872	7992	3	2	5	2	5	6
61 – 70	6037	6024	6022	5984	5984	5	5	4	9	7	31
71 – 80	4943	4968	5054	5090	5144	13	10	17	18	10	120
81+	3556	3607	3732	3822	4004	57	36	43	48	59	211
Overall	49430	49381	49301	49323	49431	79	55	70	78	83	371

Notes: This table reports descriptive statistics by age range on total population for 2015-2019, and deaths for 2015-2020. Data refers to the period from January 1st to April 4th.

Table 2: Model Estimates

	Estimate	CrI
Infection Fatality Ratios		
$\delta_{\text{Overall}}^{\text{Covid}}$	1.3068	[0.9355; 1.8914]
$\delta_{\leq 60}^{\text{Covid}}$	0.0479	[0.0040; 0.1658]
$\delta_{61+}^{\text{Covid}}$	4.1602	[3.0469; 5.8029]
$\delta_{0-20}^{\text{Covid}}$	0.0406	[0.0038; 0.1479]
$\delta_{21-40}^{\text{Covid}}$	0.0153	[0.0005; 0.0803]
$\delta_{41-50}^{\text{Covid}}$	0.0406	[0.0019; 0.1731]
$\delta_{51-60}^{\text{Covid}}$	0.1044	[0.0105; 0.2872]
$\delta_{61-70}^{\text{Covid}}$	0.9294	[0.5531; 1.4820]
$\delta_{71-80}^{\text{Covid}}$	4.5370	[3.3480; 6.3430]
$\delta_{81+}^{\text{Covid}}$	8.5040	[6.3870; 11.5700]
Baseline Death Rates		
δ_{0-20}	0.0015	[0.0001; 0.0080]
δ_{21-40}	0.0027	[0.0004; 0.0087]
δ_{41-50}	0.0163	[0.0070; 0.0314]
δ_{51-60}	0.0443	[0.0267; 0.0676]
δ_{61-70}	0.1011	[0.0700; 0.1405]
δ_{71-80}	0.2721	[0.2133; 0.3398]
δ_{81+}	1.3040	[1.1480; 1.4740]
Infection Rates		
θ_{Overall}	45.58	[30.56; 62.88]
$\theta_{\text{Casalpuusterlengo}}$	30.38	[20.14; 43.71]
$\theta_{\text{Castelgerundo}}$	24.28	[7.53; 54.62]
$\theta_{\text{Castiglione d'Adda}}$	84.68	[62.89; 97.18]
θ_{Codogno}	47.78	[33.81; 65.51]
θ_{Fombio}	60.08	[33.86; 88.78]
θ_{Maleo}	75.34	[50.69; 94.77]
$\theta_{\text{San Fiorano}}$	45.39	[21.45; 77.46]
θ_{Somaglia}	25.5	[11.22; 46.34]
$\theta_{\text{Terranova dei Passerini}}$	33.61	[9.397; 72.83]

Notes: This table reports posterior median estimates from the model described in Section 2.3. For each parameter the last column reports the 95% credible interval.

List of Figures

Figure 1: Daily deaths in 2020 and previous years average

Notes: The black solid line reports the total deaths by day in 2020 in the seven municipalities, while the blue dashed line reports the average fatalities on the same day of the previous five years. The horizontal red line marks February 20th 2020, when the first patient was admitted to the ICU in this area.

Figure 2: Estimates of the infection fatality rate by age range

Notes: For each age range a on the horizontal axis, the figure reports information on the posterior distribution of the COVID-19 infection fatality rate δ_a^{Covid} from estimating the model described in Section 2.3. Boxes represent estimated interquartile ranges, with the central line reporting the median posterior estimate, vertical lines represent 1.5 times the interquartile range.

Figure 3: Estimates of infection fatality rate by age and proportion of population infected

Notes: For each proportion of the population affected between 10% and 80% on the horizontal axis, we estimate a restricted version of the model described in Section 2.3 in which the infection rate θ_i is set equal for each municipality at that value. Solid lines represent the median of each posterior while the surrounding bands report 95% credible intervals from the 2.5 to the 97.5 percentiles of the posterior distributions. The Overall estimate is obtained by weighting posteriors by population shares in 2019.

Figure 4: Estimates of infection fatality rate by country

Notes: For each country we weighted the IFRs for different age groups by the share of each age group over the total population. The procedure performs a linear combination of IFR estimates and age groups population shares. We assumed a homogeneous contagion across age groups in every country. Country codes are reported on the horizontal axis. Dots represent weighted (by population shares) medians of the posterior distributions of different age groups obtained in the model described in Section 2.3. The bands surrounding the point estimates report 95% credible intervals from linear combinations of the 2.5 and the 97.5 percentiles of the posterior distributions.

Appendix A. Markov Chain Monte Carlo Estimation diagnostics

Figure Appendix A.1: Trace and density plots for MCMC posteriors of δ_a^{Covid} for each age range

Notes: For each of the seven age ranges a , this figure reports diagnostic plots for the MCMC simulation of the model describe in Section 2.3 for the COVID-19 infection mortality rate parameter: δ_a^{Covid} . The left panels report trace plots of the last 5000 draws from the posterior to check convergence. The right panels report the corresponding posterior distribution estimate (black solid line) together with the prior distribution for that parameter (red solid line). The % overlap reported in red is the PPO (prior posterior overlap).

Figure Appendix A.2: Trace and density plots for MCMC posteriors of δ_a for each age range

Notes: For each of the seven age ranges a , this figure reports diagnostic plots for the MCMC simulation of the model describe in Section 2.3 for the baseline mortality rate parameter: δ_a . The left panels report trace plots of the last 5000 draws from the posterior to check convergence. The right panels report the corresponding posterior distribution estimate (black solid line) together with the prior distribution for that parameter (red solid line). The % overlap reported in red is the PPO (prior posterior overlap).

Figure Appendix A.3: Trace and density plots for MCMC posteriors of θ_i for each municipality

Notes: For each of the nine municipalities i , this figure reports diagnostic plots for the MCMC simulation of the model describe in Section 2.3 on the baseline mortality rate parameter: θ_i . The left panels report trace plots of the last 5000 draws from the posterior to check convergence. The right panels report the corresponding posterior distribution estimate (black solid line) together with the prior distribution for that parameter (red solid line). The % overlap reported in red is the PPO (prior posterior overlap).