COVID-19: R_0 is lower where outbreak is larger

Abstract

We use daily data from Lombardy, the Italian region most affected by the COVID-19 outbreak, to calibrate a SIR model on each municipality. Municipalities with a higher initial number of cases feature a lower rate of diffusion, not attributable to herd immunity: there is a robust and strongly significant negative correlation between the estimated basic reproduction number (R_0) and the initial outbreak size. This represents novel evidence of the prevalence-response elasticity in a cross-sectional setting, characterized by a same health system and homogeneous social distancing regulations. By ruling out alternative explanations, we conclude that a higher number of cases causes changes of behavior, such as a more strict adoption of social distancing measures among the population, that reduce the spread. This finding calls for the distribution of detailed epidemiological data to populations affected by COVID-19 outbreaks.

Keywords: COVID-19, prevalence-response elasticity, basic reproduction number, social distancing, containment.
JEL classification: I12, I18, C53, C22.

1 Introduction

An infectious disease's basic reproduction number, or R_0, represents the average number of secondary cases produced by a single infected case in an otherwise susceptible population: it is typically used as a reference value to assess the transmissibility of an infectious disease in a given population. Given a number of individuals susceptible to infection, a disease with higher R_0 will infect a larger number of individuals. There is hence an obvious positive relationship between the R_0 and the resulting size of an outbreak (Tildesley and Keeling, 2009).

However, the spread of an outbreak does not only depend on ex-ante features of a virus or a population, but potentially also on the response of both population and authorities to the outbreak. A stream of literature has described the possible negative relationship between the prevalence of a disease in a given population and its rate of spread. Indeed, if a population's risk taking behavior is endogenous, then the larger the outbreak size, the more precautions individuals and authorities will take: this is the essence of the prevalence-response elasticity (Fenichel, 2013; Philipson, 2000; Laxminarayan and Malani, 2011). This is particularly important in the context of the COVID-19 pandemic, to which most countries in the world have reacted with some form of social distancing measures, or lockdown. In absence of a vaccine or effective drugs, these measures are the best weapon to reduce the number of deaths, as well as the number of intensive care unit beds required (Flaxman et al., 2020; Ferguson et al., 2020; Greenstone and Nigam, 2020) and have been shown to obtain a substantial reduction in the speed of contagion (Kucharski et al., 2020; Wang et al., 2020; Chudik et al., 2020). However, their effectiveness crucially relies on the willingness of the population to follow the recommendations.

Several studies have found evidence of the prevalence-response elasticity by employing data from surveys of behavior (Oster, 2012; Mullahy, 1999; Ahituv et al., 1996; Philipson, 1996), observing for instance an increase in reported condom use in areas with a higher AIDS prevalence. Epidemiological data could allow to overcome the possible problem of response bias, but the task is made non-trivial by the fact that public policies also react to the prevalence of a diseases, and by spillovers across regions due to human mobility. The present study relies on epidemiological panel data on COVID-19 in Lombardy, the region of Italy most heavily affected by the pandemic (Cereda et al., 2020 provide an accurate description of the early phase of the

outbreak in such region). Specifically, we employ daily data on the number of individuals positive to COVID-19 at the municipality level, focusing on a period in which the entire country was subject to a lockdown. All municipalities under analysis share the same public health system and, in the period considered, were subject to the same social distancing regulation. However, at the start of the period, they were characterized by a strong heterogeneity in the number of cases, both in absolute and in per capita terms.

We study a period beginning on March 25, 2020, that is, more than two weeks after the lockdown regulation was put in place, and covering three weeks, during which such regulations still held: this means that movements across municipalities were severely restricted, requiring any travelers to present a valid (typically work or health related) justification for their journey. Hence, such municipalities were living in an exceptional state of isolation.

We fit a Susceptible-Infected-Recovered (SIR) model (Kermack and McKendrick, 1927) on data from each municipality and find that the estimated R_0 is negatively correlated with the prevalence in the municipality at the beginning of our period. This result holds both when considering the absolute and per capita number of cases, and is robust to different specifications and sample disaggregations.

We present and compare different complementary explanations for this finding. Early and widespread testing can increase the reported number of cases and simultaneously allow the authorities to slow the spread of the pandemic by isolating known cases. At the same time, where the number of cases is higher, the population might comply more strictly with the lockdown measures, thus reducing the rate of spread: we show in Section 4 why this latter mechanism is most likely to drive our results.

2 Materials and methods

We employ count data of per-municipality (cumulated) recorded cases. These are updated daily and distributed by regional authorities. We do not rely on data on recovered and deceased individuals, as such data are not available with the required geographical disaggregation.

Data are available starting from March 25, 2020. This study focuses on a period of twenty-one days during which lockdown measures were always in place: in Appendix A.3, we perform a sensitivity analysis by employing later

data (until April 19).

We verify that only minimal deviations appear between regional data and the aggregation of municipal data (with the exception of data from March 24, which is indeeed removed from our sample). Out of 1507 municipalities in Lombardy, 960 had at least one recorded COVID-19 case as of March 25. Figure 1a displays the number of cases (size of the dots) and the cases per capita (color of the dots) as of March 25 for each of these 960 municipalities. Similarly, Figure 1b displays the number of new cases (size of the dots) and the number of new cases per capita (color of the dots) recorded in each municipality during the three weeks under analysis.

It should be noted that official data concerning the COVID-19 outbreak in Italy has been found to be strongly incomplete, both in terms of positive individuals and of casualties: several researchers have estimated an outbreak size much higher than that suggested by official numbers (Flaxman et al., 2020) and others have corroborated this with an analysis of anomalies in death rates (Bartoszek et al., 2020). Moreover, local testing strategies are known to have deviated from WHO guidelines and to have changed over time, also depending on available resources: towards the end of our period of interest, more subjects with mild symptoms were tested. For this reason, some researchers have put forwards adaptations of the SIR model that account for a threshold in the capacity of the health system (Ichino et al., 2020). Such problems are not specific to Italy, as official data from a number of countries have been questioned. More in general, the difficulty in obtaining reliable data on the number of infected, deceased and recovered individuals calls for refinements of traditional epidemiological models (Atkeson et al., 2020; Riccardo et al., 2020).

Given our research question, these caveats are of limited importance. Indeed, the focus of the present work is to document *differences* in response across municipalities, rather than to precisely estimate the epidemiological parameters or expected duration of the COVID-19 outbreak in Lombardy. Despite the presence of noise in the data, we show in what follows that the differential patterns identified can be distinguished from such noise.

Data on population size is obtained from the Italian National Istitute of Statistics (ISTAT).

Figure 1: Distribution of cases and cases increase

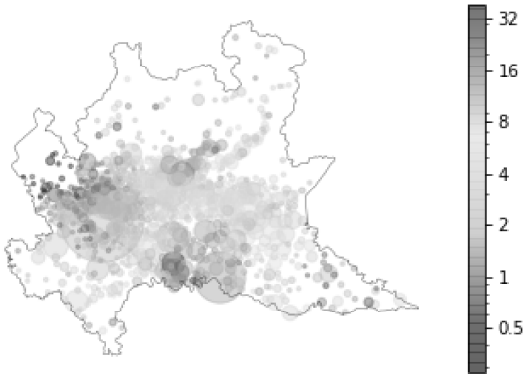

(a) Cases on March 25

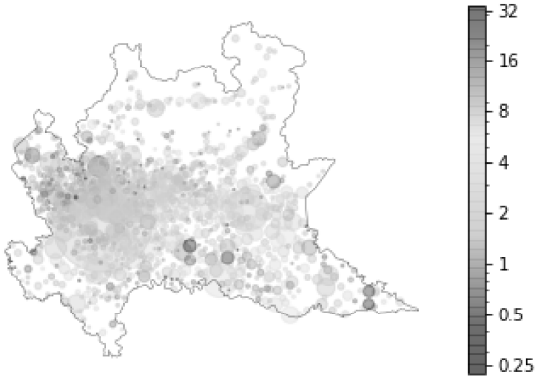

(b) New cases between March 25 and April 14

Note: dot size represent absolute numbers, colors represent cases per one thousand inhabitants.

2.1 Methods

We fit a SIR model on each municipality in the period of three weeks beginning with March 25. Given the short time span considered, we employ a simplified SIR model which does not account for natural rate of mortality. Hence, the model is entirely defined by setting few parameters: β, which determines the rate at which susceptible (S) individuals become infected (I); γ, which determines the rate at which infected individuals become recovered (R); the initial number of infected and recovered individuals, and the population size N ($= S + I + R$).

It is worth mentioning that while β and γ might not be identifiable by analyzing daily *active* cases (for sufficiently low prevalence, I approximately follows an exponential growth of base $1 + \beta - \gamma$), this concern does not apply to $I + R$ (see Appendix A.1).

We consider, for each municipality, a discretized version of the continuous SIR model – each period corresponding to a day – and automatically explore the parameter space for β, γ, and the initial value for I and R (we take population size from official statistics), looking for the combination that provides the best fit. Specifically, the goodness of fit is maximized by minimizing the sum of square residuals between the cases count and the sum of the I and R pools sizes. The initial values for the free parameters are set to those calibrated on the entire Lombardy region. The optimization algorithm is described in Appendix A.2.

Given that the SIR model assumes a non-null initial population of infected individuals, we only consider the 960 municipalities satisfying this condition. We further drop 47 municipalities which had new cases recorded on only one or two dates, hence reducing to 913 municipalities: the fitting procedure may become unreliable if provided too few updating points. See Figure 4 in Appendix A.2 for a comparison of fitted SIR models with actual cases count in three representative examples.

Once we find the best SIR parameters for each municipality, we regress the estimated R_0 (the ratio of the estimated β and γ) on the outbreak size within the municipality as of March 25. We focus on the *per capita* number of cases, as we expect any effect to be related to the *prevalence* of the outbreak – a same number of cases will be perceived in a very different way in Milan or in a small municipality.

Figure 2: Distribution of R_0

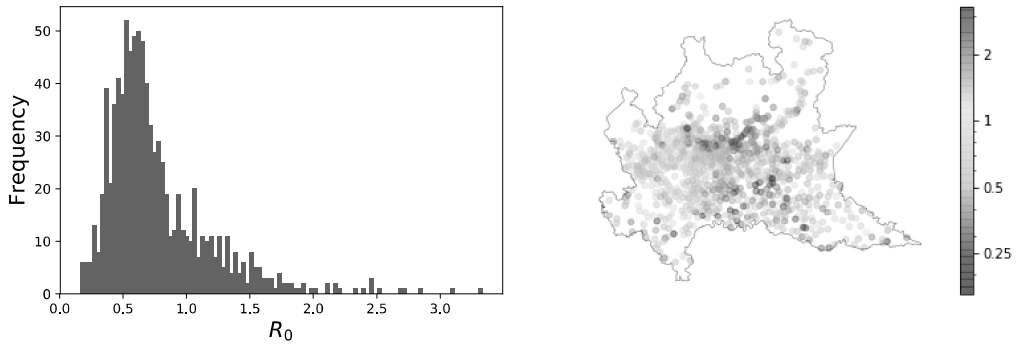

Note: distribution of estimated per-municipality R_0, computed as $\frac{\beta}{\gamma}$. Not shown: five municipalities with $R_0 < 0.16$, five with $R_0 > 3.33$.

3 Results

Among municipalities with at least one recorded case as of March 25, we find a strong (-0.38) and strongly significant (p=0.000) negative correlation between the initial number of cases per one thousand inhabitants and the percentage increase of cases in the period under analysis.

A strong heterogeneity (partly attributed to statistical noise – several municipalities count only a few cases each) can be observed in the distribution of estimated R_0 across municipalities: in what follows, unless differently specified, we trim data by dropping 0.5% of outliers on each side of the distribution of R_0, hence reducing to 902 municipalities.

Figure 2 shows the distribution of the estimated values of R_0: the mean value is 0.77 (0.79 before trimming) and becomes 0.82 when weighted on population; the median is 0.65. In only few municipalities (195) the value of R_0 appears to be larger than the critical threshold of 1: in the remaining 707 municipalities, the outbreak is expected to spontaneously extinguish without requiring herd immunity.

Table 1 presents the results of the regression analysis. We follow the recommendation of Lewis and Linzer (2005) who show that, unless an overwhelming share of the residual in the final model is due to measurement error in the dependent variable (rather than to actual noise in the relation of in-

Table 1: Main results

	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	0.888***	0.887***	0.876***	0.786***	0.818***	0.831***
	(0.021)	(0.022)	(0.022)	(0.015)	(0.019)	(0.019)
cases‰	-0.024***	-0.024***	-0.029***			
	(0.004)	(0.004)	(0.005)			
cases			-0.000	-0.001***		-0.001***
			(0.000)	(0.000)		(0.000)
new_cases‰					-0.160***	-0.150***
					(0.034)	(0.033)
population		0.000		0.000***	0.000	0.000***
		(0.000)		(0.000)	(0.000)	(0.000)
population^{-1}			121.738**			
			(58.473)			
Observations	902	902	902	902	902	902
R2	0.047	0.047	0.053	0.019	0.026	0.041

Note: dependent variable: estimated R_0. White standard errors reported in parentheses.
*p<0.1; **p<0.05; ***p<0.01

terest), an OLS with heteroskedastic robust standard errors provides better results than Weighted Least Squares (WLS) – we employ the latter approach as a robustness exercise in Appendix A.3. We see a negative and strongly significant relationship between the initial number of cases per one thousand inhabitants and the estimated R_0 (column (1)); this relationship is robust to controlling for population size (column (2)), and to both the absolute number of cases and the inverse of population size (column (3)), i.e., a full interaction model where the per capita count represents the interaction term (Kronmal, 1993). The coefficient for the per capita number of cases can be interpreted as the reduction in R_0 resulting from an increase of one case per one thousand individuals in the prevalence of the outbreak. The value of -0.024 observed in column (2), which we consider as our baseline specification, indicates a sizeable effect: for reference, given that the prevalence in Milan as of March 25 was around 1.7‰, the above mentioned result suggests that had it been 2.7‰, the average R_0 would have been around 0.909 instead of the observed 0.884. The same negative and strongly significant effect is observed if we consider as explanatory variable the absolute number of cases, controlling for population size (column (4)). This holds also when we consider as explanatory variable the number of *new* cases per capita observed on March 26 as compared to March 25, always controlling for population size (column (5)). While the number of new cases is more noisy than the total number of cases, it might be considered a more up-to-date information on the hazard posed at the local level by the epidemic. Hence, both pieces of information can in principle be relevant for individual behavior: indeed, column (6) shows that they are both strongly significant even when considered simultaneously.

It should be noted that any intrinsic characteristic of municipalities – such as demography, location, structure of the economy – which might explain a larger outbreak size should also favor a larger R_0 (Dowd et al., 2020). Thus, our minimal model will yield a lower bound for the coefficient of interest in presence of any unobservable variable which influences contagion.

4 Discussion

There are a few reasons that might explain why a larger COVID-19 outbreak should result in a *subsequent* lower R_0.

The first is related to herd immunity, by which areas where the outbreak is initially more present have less scope for further spread because a large

share of individuals have already caught, and possibly developed immunity to the virus. This is in principle not a problem for our approach, as the SIR model accounts for this effect and should estimate an R_0 net of it – in other terms, R_0 describes the evolution of the outbreak in an hypothetical situation in which the pool of susceptible individuals is never reduced. However, the problem might still arise if the count data employed severely underestimate the actual spread of the virus: the number of positive cases could actually be much larger than the detected one, leading to an estimated R_0 lower than the real one because of the undetected effect of herd immunity in reducing the rate of contagion.

The underestimation of infected population might also suggest an alternative explanation of the result, related to *test capacity*: to the extent that a lower detected prevalence reflects a lower ability of health authorities to identify infected individuals, it should then correlate with a lower ability to isolate, hospitalize and cure them, and hence to a faster outbreak growth.

A third, social, explanation is related to the prevalence-response elasticity (Philipson, 1996): if the local population is aware of a larger prevalence of the disease, it reacts by changing its behavior towards a stricter application of social distancing rules, thus leading to a lower R_0. This might in fact depend from other motives than minimizing the individual risk of infection. Subjects may want to reduce the risk for their fellow citizens, and in particular relatives and family members; likewise, they might want to avoid the risk that a larger outbreak results in postponing the end of the lockdown. In what follows, we provide evidence in favor of this social mechanism.

We start by analyzing the first possible explanation, related to herd immunity: several sources have argued that the actual size of the infected population might lie between four and ten times the official reported numbers. In the most affected municipalities in our sample during the period analyzed, 57 infections per one thousand inhabitants were recorded, and according to the most pessimistic estimates this would mean that up to 57% of the population was infected. While most municipalities have a number of recorded cases per one thousand inhabitants which is orders of magnitude lower, to avoid the possibility that an even partial herd immunity effect might be driving the results, we re-estimate our main model on subsamples of municipalities, based on their initial number of cases per capita. Specifically, we split the sample according to quartiles of cases per capita on March 25. The results, presented in columns (1) to (4) of Table 2, show that our findings are not driven by herd immunity, as the coefficient for *cases‰* is negative in each

quartile. The coefficient is significant, and its absolute value particularly large, for municipalities with a relatively low prevalence (first three quartiles), including municipalities with 2 cases per one thousand inhabitants or less (first quartile).

We then consider the second possible explanation: that a lower *detected* prevalence signals a lower detection ability, and that this mechanically correlates with a lower ability to track and isolate infected subjects, hence raising the subsequent rate of diffusion. In order to disentangle this *test capacity* explanation from the third, social, one, we sketch two simple models of how these would be expected to affect R_0.

Let us represent with u_t the unknown real number of infected subjects per one thousand inhabitants at time t in a given municipality, and with i_t the corresponding known number. We are interested in the extent to which unidentified infected subjects (which are $u_t - i_t$ cases per one thousand inhabitants) will raise the R_0 for the municipality in the subsequent period. More specifically, we can assume that identified and unidentified patients form two different pools of infected subjects and that the latter has a much higher β – probability of infecting susceptible individuals – that leads to a corresponding higher R_0. Since β enters linearly in R_0 – and assuming for simplicity that γ is constant – the relationship between $u_t - i_t$ and R_0 would be expected to be *linear*. Moreover, not only identified patients are subject to a stronger form of isolation, but also close contacts of such patient (some of which are not infected) are recommended to self-quarantine: this does not happen in municipalities with a larger number of undetected cases, which implies that the effect of each unidentified patient should be *more* than linear in increasing the R_0. This would imply a linear or concave relationship between *cases‰* and R_0.

Vice-versa, any social explanation (Sabat et al., 2020; Ichino et al., 2020) is based on the assumption that inhabitants react to the news of the cases in their municipality. A same increase in per capita cases is likely to be more salient if the initial number of cases is lower. That is, we can expect inhabitants of two towns with respectively 1 and 11 known cases per one thousand inhabitants to differ in their compliance with social distancing prescriptions more than inhabitants of two towns with respectively 20 and 30 known cases per one thousands inhabitants: a same difference of one percentage point in prevalence will have a weaker effect on people behavior where prevalence is higher. This alternative explanation predicts a *convex* relationship (given the negative sign) between *cases‰* and R_0.

Table 2: Additional specifications

	Q 1	Q 2	Q 3	Q 4	Full
	(1)	(2)	(3)	(4)	(5)
Intercept	1.130***	1.129***	1.027***	0.713***	0.952***
	(0.072)	(0.133)	(0.211)	(0.068)	(0.030)
cases‰	-0.142***	-0.094**	-0.067*	-0.004	-0.047***
	(0.049)	(0.041)	(0.041)	(0.006)	(0.009)
cases‰^{2}					0.001***
					(0.000)
population	0.000	-0.000**	0.000	-0.000**	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	226	225	225	226	902
R2	0.036	0.022	0.014	0.008	0.061

Note: dependent variable: estimated R_0. White standard errors reported in parentheses. Columns (1) to (4): model restricted to municipalities with a number of cases per thousand inhabitants in the interval (0.278, 2.173], (2.173, 4.11], (4.11, 6.265] and (6.265, 38.743], respectively. Column (5): full sample. *p<0.1; **p<0.05; ***p<0.01

To disentangle between the test capacity and the social explanation, we enrich our basic model by introducing a quadratic term in *cases‰*. This is done in column (5) of Table 2. We see that the quadratic term has a positive sign and is strongly significant, while the sign of the linear term is still negative and has increased in absolute terms. Hence, while this does not allow us to exclude that the two other explanations might play a role, we can conclude that the *social* explanation is the main driver of the negative relationship between *cases‰* and R_0. See Appendix A.3 for additional evidence and sensitivity tests.

Jones et al. (2020) describes two possible opposite reactions to the COVID-19 outbreak: a precautionary attitude that leads to a stricter adherence to guidelines, and a "fatalism effect" according to which an individual who is more likely to be infected in the future "*reduces her incentives to be careful today*". Our results provide strong evidence in favor of the first mechanism.

Chudik et al. (2020) follow a different approach to model how changes in behavior affect the speed of diffusion in a SIR model. Rather than looking for variations in R_0 as originating from changes in β or γ, they consider such parameters as fixed, and assume instead that in regions with more active cases, more people will isolate (both because of mandatory and voluntary social distancing measures). When this happens, the sample S of susceptible individuals shrinks: this results in a reduction of the speed of contagion, despite an *increase* in the infection probability for susceptible individuals. Intuitively, the latter effect comes from the assumption that, while a number of sane individuals disconnect from contagion, the remaining ones keep their frequency of contacts unchanged, increasing their chance $\beta \frac{I_t}{P_E}$ (with P_E being the size of the non-isolated population) of contracting the disease. Our model and theirs are complementary, as social distancing can imply a generalized change of habits in the population but also, as a result, the isolation from contagion of a share of the population. In their paradigm, the results we obtain when restricting to quantiles of municipalities (Table 2, columns 1 to 4) and when introducing a quadratic term in the number of initial cases (Table 2, column 5) can be interpreted as a less-than linear increase in the number of isolated subjects in response to an increase in the number of recorded cases.

The endogeneity of risk taking with respect to outbreak size reveals one of the subtleties in predicting an outbreak's date of extinction (Geoffard and Philipson, 1996). As already observed by Philipson (1996), a lower prevalence might lead to less precautions been taken and hence to a longer outbreak

duration. To the extent that there is habit formation in the population, this might result in a larger outbreak becoming extinct before a smaller one, all else equal, and even in the case $R_0 < 1$.

5 Conclusions

We show that in the Italian region of Lombardy, during a lockdown, the basic reproduction number for COVID-19 reacts negatively to the initial size of an outbreak at the municipality level, an effect which cannot be explained by the population having reached herd immunity. Limited test capacity – and hence a limited ability by health authorities to isolate and treat affected individuals – appear to have at most a marginal role in explaining our result. Instead, our results provide support for an important role of the prevalence-response elasticity: population behavior is key to slowing down the contagion and, in particular, information about local outbreaks impacts on diffusion rates. This effect is consistent across all provinces and it is robust to the time period considered.

This is, to the best of our knowledge, the first time that a change of behavior in the population due to the prevalence-response elasticity is observed within a period of few days. The speed of the contagion and the reduced human mobility across interested municipalities reduce the relevance of geographic spillovers. This allows for a novel identification of the behavioral response directly from local epidemiological data, rather than having to rely on survey-based measures of health behavior. The fact that the behavioral response is particularly strong in municipalities characterized by smaller outbreaks suggests that individuals react more strongly to the first few cases. This aspect is confirmed by the convex relationship we find between the initial size of the outbreak and the R_0: the marginal effect on behavior of each new case seems to decrease in the number of cases.

Our results provide evidence in favor of a *precautionary* rather than *fatalistic* individual attitude towards the outbreak. This is a socially efficient reaction, as it reduces the likelihood of contagion precisely where it would have the highest potential to spread. The evidence provided hence calls for considering populations as an integral part of the decision making process (Sabat et al., 2020), and suggests that a timely and transparent provision of epidemiological data can increase the efficacy of a lockdown, hence reducing its negative social and economic impact, and yielding public health benefits

at no added cost.

References

Ahituv, A., V. J. Holz, and T. Philipson (1996). The Responsiveness of the Demand for Condoms to the Local Prevalence of AIDS. *Journal of Human Resources 31*(4), 869–897.

Atkeson, A. et al. (2020). How Deadly is COVID-19? Understanding the Difficulties with Estimation of its Fatality Rate. Technical report, Federal Reserve Bank of Minneapolis.

Bartoszek, K., E. Guidotti, S. M. Iacus, and M. Okrój (2020). Are official confirmed cases and fatalities counts good enough to study the COVID-19 pandemic dynamics? A critical assessment through the case of Italy. *arXiv preprint arXiv:2005.07271*.

Cereda, D., M. Tirani, F. Rovida, V. Demicheli, M. Ajelli, P. Poletti, F. Trentini, G. Guzzetta, V. Marziano, A. Barone, M. Magoni, S. Deandrea, G. Diurno, M. Lombardo, M. Faccini, A. Pan, R. Bruno, E. Pariani, G. Grasselli, A. Piatti, M. Gramegna, F. Baldanti, A. Melegaro, and S. Merler (2020). The early phase of the COVID-19 outbreak in Lombardy, Italy.

Chudik, A., M. H. Pesaran, and A. Rebucci (2020). Voluntary and Mandatory Social Distancing: Evidence on COVID-19 Exposure Rates from Chinese Provinces and Selected Countries. Technical report, National Bureau of Economic Research.

Dowd, J. B., L. Andriano, D. M. Brazel, V. Rotondi, P. Block, X. Ding, Y. Liu, and M. C. Mills (2020). Demographic science aids in understanding the spread and fatality rates of COVID-19. *Proceedings of the National Academy of Sciences 117*(18), 9696–9698.

Fenichel, E. P. (2013). Economic considerations for social distancing and behavioral based policies during an epidemic. *Journal of Health Economics 32*(2), 440–451.

Ferguson, N., D. Laydon, G. Nedjati-Gilani, N. Imai, K. Ainslie, M. Baguelin, S. Bhatia, A. Boonyasiri, Z. Cucunubá, G. Cuomo-Dannenburg, et al. (2020). Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Imperial College COVID-19 Response Team.

Flaxman, S., S. Mishra, A. Gandy, et al. (2020). Estimating the number of infections and the impact of nonpharmaceutical interventions on COVID-19 in 11 European countries. *Imperial College COVID-19 Response Team 30*.

Geoffard, P.-Y. and T. Philipson (1996). Rational epidemics and their public control. *International Economic Review*, 603–624.

Greenstone, M. and V. Nigam (2020). Does social distancing matter? *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2020-26).

Hornstein, A. S. and W. H. Greene (2012). Usage of an estimated coefficient as a dependent variable. *Economics Letters 116*(3), 316–318.

Ichino, A., C. A. Favero, and A. Rustichini (2020). Restarting the economy while saving lives under COVID-19. *CEPR Discussion Paper Series* (DP14664).

Jones, C. J., T. Philippon, and V. Venkateswaran (2020). Optimal mitigation policies in a pandemic: Social distancing and working from home. Technical report, National Bureau of Economic Research.

Kermack, W. O. and A. G. McKendrick (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.

Kronmal, R. A. (1993). Spurious correlation and the fallacy of the ratio standard revisited. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 156*(3), 379–392.

Kucharski, A. J., T. W. Russell, C. Diamond, Y. Liu, J. Edmunds, S. Funk, R. M. Eggo, F. Sun, M. Jit, J. D. Munday, et al. (2020). Early dynamics of transmission and control of COVID-19: a mathematical modelling study. *The Lancet Infectious Diseases*.

Laxminarayan, R. and A. Malani (2011). The economics of infectious disease. In S. Glied and P. Smith (Eds.), *The Oxford Handbook of Health Economics*. Oxford, UK: Oxford University Press.

Lewis, J. B. and D. A. Linzer (2005). Estimating regression models in which the dependent variable is based on estimates. *Political Analysis 13*(4), 345–364.

Mullahy, J. (1999). It'll only hurt a second? Microeconomic determinants of who gets flu shots. *Health Economics 8*(1), 9–24.

Oster, E. (2012). HIV and sexual behavior change: Why not Africa? *Journal of Health Economics 31*(1), 35–49.

Philipson, T. (1996). Private vaccination and public health: an empirical examination for US measles. *Journal of Human Resources*, 611–630.

Philipson, T. (2000). Economic epidemiolgy and infectious diseases. In A. Culyer and J. Newhouse (Eds.), *Handbook of Health Economics*, Volume 1. San Diego, CA.: Elsevier Science.

Riccardo, F., M. Ajelli, X. Andrianou, A. Bella, M. Del Manso, M. Fabiani, S. Bellino, S. Boros, A. Mateo Urdiales, V. Marziano, M. C. Rota, A. Filia, F. P. D'Ancona, A. Siddu, O. Punzo, F. Trentini, G. Guzzetta, P. Poletti, P. Stefanelli, M. R. Castrucci, A. Ciervo, C. Di Benedetto, M. Tallon, A. Piccioli, S. Brusaferro, G. Rezza, S. Merler, and P. Pezzotti (2020). Epidemiological characteristics of COVID-19 cases in Italy and estimates of the reproductive numbers one month into the epidemic. *medRxiv*.

Sabat, I., S. Neuman-Böhme, N. E. Varghese, P. P. Barros, W. Brouwer, J. van Exel, J. Schreyögg, and T. Stargardt (2020). United but divided: policy responses and people's perceptions in the EU during the COVID-19 outbreak. *Health Policy*.

Saxonhouse, G. R. (1976). Estimated parameters as dependent variables. *The American Economic Review 66*(1), 178–183.

Tildesley, M. J. and M. J. Keeling (2009). Is R0 a good predictor of final epidemic size: Foot-and-mouth disease in the UK. *Journal of Theoretical Biology 258*(4), 623–629.

Wang, C., L. Liu, X. Hao, H. Guo, Q. Wang, J. Huang, N. He, H. Yu, X. Lin, A. Pan, et al. (2020). Evolving epidemiology and impact of non-pharmaceutical interventions on the outbreak of coronavirus disease 2019 in Wuhan, China. *medRxiv*.

A Appendix

A.1 Degrees of freedom in the SIR model

In the simplified SIR model adopted in this study, the number of infected individuals grows as

$$I_t = I_{t-1} + \frac{\beta I_{t-1} S_{t-1}}{N} - \gamma I_{t-1}. \tag{1}$$

In the initial phases of an outbreak – that is, as long as prevalence is low – we have that $S_t \approx S_0 = N$, obtaining

$$I_t \approx I_{t-1} + \beta I_{t-1} - \gamma I_{t-1} = (1 + \beta - \gamma) I_{t-1}. \tag{2}$$

By iterating Equation (2), we obtain $I_t \approx (1 + \beta - \gamma)^t I_0$. Hence, β and γ cannot be disentangled by looking at the time series of I_t, unless the effect of herd immunity is relevant (that is, $S_t \ll S_0 = N$). In other words, the exponential curve that describes the beginning of the outbreak has only one degree of freedom, which is determined by $\beta - \gamma$ and in turn determines both the size at any given time t, and the curvature.

The number R_t of recovered cases until time t grows instead as a proportion γ of currently infected individuals:

$$R_t = R_{t-1} + \gamma I_{t-1} \approx R_{t-1} + \gamma (1 + \beta - \gamma)^{t-1} I_0. \tag{3}$$

Since $R_t = \sum_{\tau=1}^{t} (R_\tau - R_{\tau-1})$, and assuming, without loss of generality, $R_0 = 0$,

$$\begin{aligned} R_t \approx & \sum_{\tau=1}^{t} \gamma (1 + \beta - \gamma)^{\tau-1} I_0 \\ = & \gamma I_0 \sum_{\tau=0}^{t-1} (1 + \beta - \gamma)^{\tau} \\ \overset{\gamma \neq \beta}{=} & \gamma I_0 \frac{1 - (1 + \beta - \gamma)^{t-1}}{\gamma - \beta}. \end{aligned} \tag{4}$$

From Equation (4) it is evident that β and γ affect the evolution of R_t in different ways: specifically, for a same value of $\beta - \gamma$, a smaller value of γ results in a linear scaling down of the entire series, while for a same value

of γ, changing β results in changing the *rate* of growth $\frac{R_t - R_{t-1}}{R_{t-1}}$. Hence, the parameters β and γ can be obtained from observational data on R_t. The same can be said for the count of total cases $I_t + R_t$: since the first term only depends on $\beta - \gamma$ but the second depends on β and γ separately, their sum does too. Specifically,

$$I_t + R_t \approx I_0 \frac{\gamma - \beta(1 + \beta - \gamma)^{t-1}}{\gamma - \beta}. \tag{5}$$

The argument presented above fails when $\beta = \gamma$, in which case Equation (4) cannot be derived. In this case, however, $I_t = 1^t I_0 = I_0$ and $R_t = \sum_{\tau=1}^{t} \gamma I_{t-1} = \gamma(t-1)I_0$, so that $I_t + R_t = I_0 \cdot ((t-1)\gamma + 1)$ follows a linear trend and the (common) value of the parameters determines its slope.

A.2 Optimization procedure

For ease of interpretation, the procedure for fitting the SIR model is implemented over the parameters R_0 and γ rather than β and γ, where $R_0 = \frac{\beta}{\gamma}$. The optimization procedure works as follows (i denotes an iteration):

1. set $i = 0$ and, given the initial value $\underline{\pi}$ for each parameter π, set $\delta_{\pi,i} = \underline{\pi} \cdot 0.5$,
    2. given the current value of π_i, compute $\pi_{i,l} = \pi_i - \delta_{\pi,i}$ and $\pi_{i,r} = \pi_i + \delta_{\pi,i}$ (the *right* and *left candidate values* for parameter π),
    3. measure the fit of the simulation obtained with each of the three candidates $\pi_{i,l}$, π_i and $\pi_{i,r}$ by computing the sum of squared residuals,
    4. pick the best candidate as the new parameter value π_{i+1},
    5. if the value of the parameter did *not* change (that is, $\pi_{i+1} = \pi_i$), set $\delta_{\pi,i+1} = \delta_{\pi,i} \cdot 0.75$; otherwise, leave $\delta_{\pi,i+1} = \delta_{\pi,i}$,
    6. repeat steps from 2 to 5 for each parameter π among R_0, γ, and the initial values for I and R,
7. repeat steps from 2 to 6 until $\frac{\delta_{\pi,i}}{\pi_i} < 0.001$ for all parameters (that is, until each parameter is determined with a precision of 1‰).

We first run the procedure on regional data:

Figure 3: Calibration of R_0 and γ_0

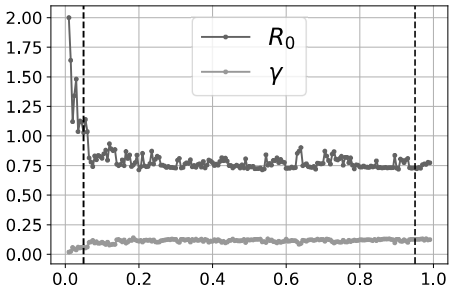

- the initial values for R and I are set to the number of infected cases at the regional level and to the number of recovered or deceased individuals at the regional level, respectively,
- the initial value for R_0 is set to 0.9,
- we repeat the optimization procedure using, as starting value for γ, each value in the set $\{0.050, 0.055, 0.060, \ldots 0.950\}$.

With the exception of very low initial values of γ (for which the optimization procedure misbehaves, given that its speed of exploration of the parameters space is by design related to the absolute size of parameters), the resulting values of γ and R_0 are stable (see Figure 3). In particular, between 0.050 and 0.950 they have a mean of 0.115 and 0.771 and a standard deviation of 0.011 and 0.046, respectively. We take such means as initial parameter values for the SIR fitted at the municipality level.

The initial values for R and I at the municipal level are set instead to the corresponding regional values as of the beginning of the period of interest, multiplied by the share of regional (total) cases reported in the municipality. In other words, given the 20,591 infected and 11,755 recovered patients recorded in Lombardy as of March 25, respectively 64% and 36% of total cases, we assume that a municipality with 100 reported cases would have 64 infected and 36 recovered (or dead) individuals. For each municipality we then follow the optimization procedure described above.

While in principle we could consider a constraint by which the sum of the initial values of I and R adds up to the initial number of cases, this is

not required nor efficient. It is not required because the fitting procedure minimizes the fit error in all periods, including the first; it is not efficient because the first datum might legitimately be affected by random fluctuations that bear no larger importance than subsequent ones.

Figure 4a displays the fit between data aggregated at the regional level and the corresponding simulated SIR model. Figures 4b and 4c are the equivalent for Milan and Castiglione d'Adda: these are the two municipalities which, at the beginning of our period of interest, had been most heavily hit in absolute and per capita terms, respectively. A weekly fluctuation can be observed: this is in line with documented evidence that less tests are processed during the weeekend, and the effect reverberates on the number of positive detected cases with a delay of two to three days. We expect these fluctuations to affect the entire region homogeneously.

Figure 4: Comparison of fitted SIR model and total cases count

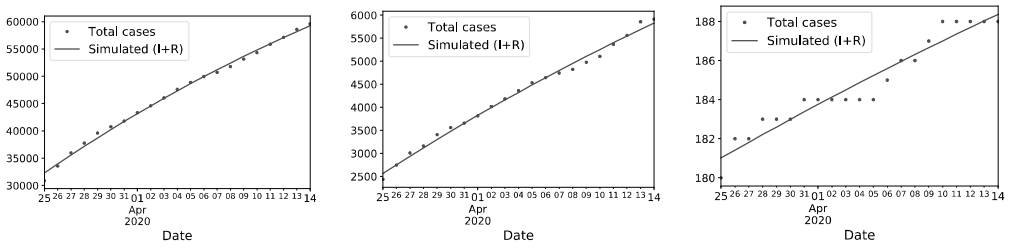

(a) Lombardy (b) Milan (c) Castiglione d'Adda

Note: fit between data and the corresponding SIR model for Lombardy region (left) and the most affected municipalities at the beginning of our period of interest in absolute and per capita terms, respectively (center, right).

A.3 Sensitivity tests

In addition to the quantile analysis described in Section 4, we verify that our main result also holds consistently across the 12 provinces (lower level administrative regions) of which Lombardy is composed. Results are displayed in Figure 5a. We see that, for each province, the effect of *cases‰* on R_0 is negative: although the small sample size results in only few provinces reaching statistical significance, it is clear that no specific area of Lombardy is alone responsible for our findings. Interestingly, the effect is relatively small in the

province of Lodi, which was strongly affected during the initial phase of the outbreak. In fact, ten municipalities in such province were interested by an earlier, very strict lockdown, beginning on February 21 and lasting until the start of the regional lockdown. We verify that four of these municipalities are excluded from our main sample because they do not have new cases in at least two dates; of the remaining six, two are excluded because of trimming, given their very low estimated R_0.

In order to verify that our results do not strictly depend upon the period considered, we replicate our analysis over different 10-days moving windows covering up to April 19 (hence going beyond our main period of analysis). By shrinking the period of analysis, we obtain 17 intervals, allowing us to study the evolution of our main result over time. For each subperiod, we fit the R_0 for each municipality and regress it on the number of cases per thousand individuals at the beginning of that subperiod. In accordance with the selection procedure described in Section 2.1, we reduce this analysis to the municipalities that feature at least one case on March 25 and, *for any given window*, have new cases recorded in at least two dates of that window. For each window, we also trim outlier values of R_0 as for the main analysis. The results are shown in Figure 5b. For comparability, we also display, for each window, the value of the coefficient estimated on the same sample but for the original 21 days time period.

We find that the effect of interest is robust, that is, the coefficient for the *cases‰* variable is consistently negative and significant for each time window covering up to April 7, becoming indistinguishable from zero on April 8. Its absolute value is significantly decreasing over time; that is, the effect of the number of cases on the R_0 in the following days appears to be stronger in the earlier days of the outbreak. While such dynamic might have different explanations, we remark that the speed of diffusion of the epidemic has been consistently decreasing. During the outbreak, citizens of Lombardy had no information on the number of *active* cases in their municipality at a given moment in time: hence, they could only infer this from the total cases counts, which however becomes a weaker proxy as the epidemics progresses. Whether individual behavior responds more to the number of active cases or to updates concerning total cases is a relevant issue for further research.

We also verify that all results reported in Table 1, including statistical significance, are virtually unchanged if we do not trim the data as previously described.

We also fit a SIR model on each of the 106 Italian provinces (data at the

Figure 5

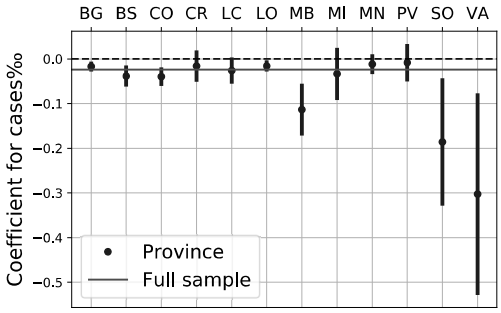

(a) Disaggregated estimation across provinces

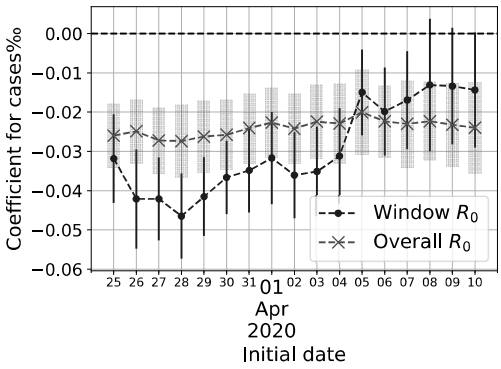

(b) Estimation over 10-days moving windows

Note: estimates are run controlling for population, as in column (2) of Table 1. Error bars represent robust 95% confidence intervals. In (a), the horizontal line denotes the corresponding coefficient estimated on the entire sample under analysis. In (b), crosses denote the corresponding coefficient estimated on the window-specific sample but over the original time period (from March 25 to April 14).

municipality level are not available for other regions). These are observed over the same period of time (since the day after its introduction, the regional lockdown was extended to the entire country). We find again a strongly significant negative relationship between the estimated R_0 (mean value: 0.85) and the number of initial cases per one thousand inhabitants. The coefficient (-0.041 when controlling for population size) is actually larger in absolute terms than that appearing in column (2) of Table 1. However, it should be noticed that the interpretation of this result is less straightforward, because different Italian regions have different health systems: those facing severe outbreaks might implement specific strategies to fight the diffusion. Hence, it is not possible to disentangle the effect of an institutional response from the effect of the population's change in behavior.

Our main results are obtained by estimating both β and γ (simultaneously) for each municipality. While this is a feasible approach, given that we focus on the total number of cases, we also replicate our main analysis by keeping γ fixed (hence only estimating β and the initial values for I and R). We hence derive beforehand γ from a SIR model fitted on data from the entire Lombardy region in the same period of analysis (see Appendix A.2 for details), and use the resulting value for all municipalities. While our main approach is more conservative – as a faster growth in the number of cases might in principle be due to heterogeneity in both β and γ – this robustness test is based on the plausible assumption that the average time to recovery (which is $\frac{1}{\gamma}$) is roughly constant across municipalities. Results confirm the negative relationship between *cases‰* and R_0: the coefficient is actually larger in absolute value (-0.040) than with our main specification, and it is still strongly significant ($p = 0.000$).

Finally, given that our dependent variable (R_0) is itself a parameter estimated at the municipality level, we re-estimate our main model via a WLS regression where the weights are given by the inverse of the variance of the normalized municipality-specific errors (Saxonhouse, 1976; Hornstein and Greene, 2012). That is, defining as $C_{m,t}$ the total cases reported until day t in municipality m (with $\bar{C}_m$ its mean), and as $\hat{I}_{m,t}$ and $\hat{R}_{m,t}$ the values estimated according to the fitted SIR model, observation m is attributed weight

$$w_m = \frac{1}{\sum_t (r_{m,t})^2}, \text{ where } \quad r_{m,t} = \frac{C_{m,t} - (\hat{I}_{m,t} + \hat{R}_{m,t})}{\bar{C}_{m,t}}.$$

With this WLS estimation, we still obtain a strongly significant effect ($p =$

0.000), albeit with a smaller coefficient in absolute terms (-0.012), presumably reflecting the disproportionate weight given to more affected municipalities, where data are more robust and the effect of interest smaller.