

Estimating and forecasting COVID-19 attack rates and mortality

David I. Ketcheson^{1,4}, Hernando C. Ombao¹, Paula Moraga², Tarig Ballal¹,
and Carlos M. Duarte³

¹*Computer, Electrical, and Mathematical Sciences & Engineering Division, King Abdullah University of Science and Technology, 4700 KAUST, Thuwal 23955, Saudi Arabia*

²*Department of Mathematical Sciences, University of Bath, Bath, BA2 7AY, UK*

³*Biological and Environmental Science and Engineering Division, King Abdullah University of Science and Technology*

⁴*david.ketcheson@kaust.edu.sa*

May 11, 2020

Abstract

We describe a model for estimating past and current infections as well as future deaths due to the ongoing COVID-19 pandemic. The model does not use confirmed case numbers and is based instead on recorded numbers of deaths and on the age-specific population distribution. A regularized deconvolution technique is used to infer past infections from recorded deaths. Forecasting is based on a compartmental SIR-type model, combined with a probability distribution for the time from infection to death. The effect of non-pharmaceutical interventions (NPIs) is modelled empirically, based on recent trends in the death rate. The model can also be used to study counterfactual scenarios based on hypothetical NPI policies.

1 Introduction

The COVID-19 pandemic progressively propagated from Wuhan, China, where it initiated in December 2019 to achieve pandemic status affecting 173 nations of 195 nations worldwide by April 2020¹. In many affected nations, COVID-19 has become a leading cause of mortality, with eight European nations (Belgium, UK, Spain, Italy, The Netherlands, San Marino, Andorra and France) topping the accumulated fatality rate, exceeding 300 individuals per million inhabitants².

¹WHO, <https://www.who.int>

²WHO, <https://www.who.int>, accessed 6 May, 2020

While effective therapies and vaccines are being developed, confinement emerges, as in the Spanish flu pandemic a century ago [26], as the leading policy to mitigate the impacts, by reducing the peak number of infections, hospitalizations and fatalities [21, 20]. However, strict confinement measures to mitigate against COVID-19, as recommended by WHO, are being disputed for their impacts on the economy and jobs. The development of cost-benefit analysis of confinement measures requires reliable projections of both the costs to the economy and the health cost (Mubayi et al. 2010), ultimately dependent on the fatality rate.

Exit strategies that attempt to minimize both public health and economic impacts depend critically on reliable models for both. The fundamentals of analytical models to predict disease spread were laid a century ago (e.g. [14]), allowing models of the health risks associated with COVID-19 to emerge rapidly [11, 13, 17]). The performance of these models depends critically on the quality of the initial data referring to the population affected (total population size, people diagnosed and deaths attributed to COVID-19 over time), the parameters specific to the COVID-19 propagation and impact (e.g. R_0 , Fatality rate), and the impact of mitigation measures on those parameters [22, 24]). As COVID-19 spreads to many countries and other situations, such as vessels with large confined populations, multiple series of data are emerging that offer opportunities to improve the assessment of key parameters. Moreover, as the effort and resolution of testing varies over time, it has become evident that the number of diagnosed cases is unreliable as it depends on effort and particular kits used, which vary among and within nations. Although the number of deaths reported may also be underestimated, as COVID-19 deaths may be undetected or confounded with other causes, it provides a minimum estimate that is reported in a relatively consistent manner across many nations, following WHO guidelines, which require fatalities to be listed only where the disease caused, or is assumed to have caused, or contributed to death.³

1.1 Modeling approach

Most available models make use of confirmed cases as the primary input or calibration; for example [1, 3, 4, 5, 6, 7]. Confirmed case numbers drastically underestimate the real number of infections; furthermore, the ratio of real infections to confirmed cases varies greatly both in time and from region to region, due to differences in testing and in reporting policies (see for example Figure 3 herein). Some models also ignore the fundamentals of epidemiology and are based merely on curve-fitting; for example [2, 3].

Here we describe a model for inferring the past and present attack rate, and predicting the death toll over time due to the ongoing COVID-19 pandemic. The model is based on recorded numbers of deaths and COVID-19 parameters derived from estimates in the literature; it makes predictions of daily numbers of new infected persons and deaths. The effect of non-pharmaceutical interventions (NPIs) such as school closures, prohibition of public events, and restrictions on internal movement is modelled empirically, based on recent trends in the death rate reported for a number of nations and of states in the USA following different levels of normative intervention effected in these case studies.

³https://www.who.int/classifications/icd/Guidelines_Cause_of_Death_COVID-19.pdf

The model can also be used to study counterfactual scenarios based on hypothetical NPI scenarios, and, therefore, be used to explore different proposed policies.

Highly sophisticated epidemiological models involving, for instance, spatial resolution and networks of social connections, can be of great use in making detailed predictions, provided that the relevant parameters are adequately known. This is not yet the case for COVID-19, given that estimates of the basic reproduction number itself differ substantially [31]. As a result, current predictions of COVID-19 infections and deaths differ by orders of magnitude [10]. Most of those predictions are based on confirmed case numbers, which are known to be highly inaccurate. Given these large uncertainties, we propose that a straightforward model based on relatively well-established inputs may be more useful than a complex model requiring poorly-constrained or highly inaccurate data and parameters. The present work is based on one of the simplest epidemiological models available, but with a focus on using the most accurate available inputs. The technique used herein to infer infections from recorded deaths could also be used to provide more accurate input to more sophisticated models, in place of confirmed case numbers.

The model algorithm is described in Section 2. The model data sources and values of model parameters are given in Section 3. In Section 4 we study several model results, including predictions for a counterfactual scenario showing what might have happened if nothing was done to slow the spread of COVID-19. Finally, in Section 5 we discuss the main sources of error and uncertainty in the model.

The model code, along with scripts to generate the figures and tables in this paper, are available⁴.

2 Model algorithm

The central element of the forecasting algorithm is based on the classic SIR model [14]:

$$\begin{aligned}\frac{dS}{dt} &= -\beta I \frac{S}{N} \\ \frac{dI}{dt} &= \beta I \frac{S}{N} - \gamma I \\ \frac{dR}{dt} &= \gamma I.\end{aligned}$$

Here $S(t)$, $I(t)$, $R(t)$ represent the susceptible, infected, and removed (recovered or deceased) fractions of the population, respectively. To account for NPIs, we add a time-dependent intervention parameter $q(t)$, which models the change in the contact rate β due to social distancing, school and work closures, and other such measures. This model

⁴https://github.com/ketch/covid_forecasting

takes the following form [15]:

$$\frac{dS}{dt} = -(1 - q(t))\beta I \frac{S}{N} \quad (2a)$$

$$\frac{dI}{dt} = (1 - q(t))\beta I \frac{S}{N} - \gamma I \quad (2b)$$

$$\frac{dR}{dt} = \gamma I \quad (2c)$$

A value $q(t) = 0$ corresponds to no intervention, while $q(t) = 1$ corresponds to a complete elimination of contact with infected individuals. This parameter can model both population-wide measures and those specifically targeting infected or suspected individuals. The time-dependent reproduction number is a measure of transmissibility that denotes the average number of secondary cases caused by an infected individual in a fully-susceptible population and can be expressed as

$$\sigma(t) = \frac{(1 - q(t))\beta}{\gamma}.$$

In the forecast model, β and γ (which are primarily properties of the disease) are fixed for all countries and all time, but $q(t)$ (which is primarily an effect of human actions) is time-dependent and different for each country.

5

Many further extensions of the SIR model exist, in which other groups, such as those under quarantine are directly modeled. Our rationale for using the model (2) is twofold. First, due to the prevalence of asymptomatic cases and broad social-distancing measures, the dominant effect of intervention is through general reduction in the contact rate rather than quarantining of a small group. Second, the existing uncertainties in key parameters for this still-developing pandemic are likely to be much greater than the differences between various refinements of the model.

2.1 Model structure

Let j indicate the date, where for convenience $j = 0$ denotes the present day.

- S_j : Susceptible individuals on day j
- i_j : Newly infected individuals on day j ; $i_j = S_{j-1} - S_j$
- I_j : Currently infectious individuals on day j
- D_j : Cumulative deaths up to day j
- d_j : New deaths on day j ; $d_j = D_j - D_{j-1}$
- R_j : Cumulative removed (including both recovered and deceased) individuals on day j .

As depicted in Figure 1, the forecast model involves three main parts:

⁵Can we also connect the $(1 - q(t))\beta/\gamma$ as the effective reproduction rate at time t which takes into account the impact of an intervention at time t ? This is now done.

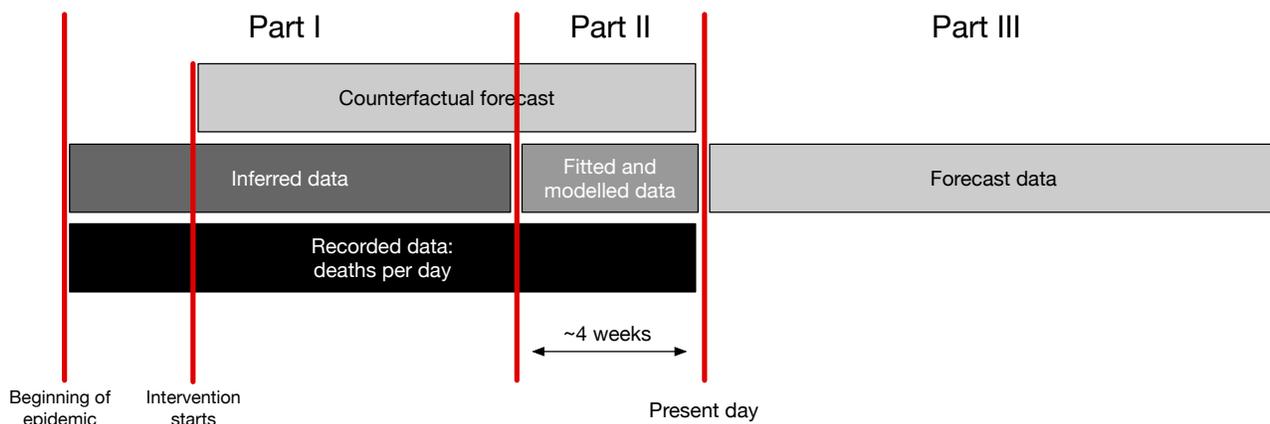


Figure 1: Organization of the elements of the model.

1. Inference of real past infection rates based on recorded numbers of deaths up to a few weeks in the past (to avoid bias due to cases that have not yet died but will do so in the future).
2. Simultaneous modeling and fitting of data over a recent time period approximately equal to the mean time from infection to death.
3. Forecasting future values based on a hypothetical level of intervention.

Each of these parts is described in the following sections.

2.2 Part I: Inference of past attack rate

In this part we infer the number of newly infected each day, i_j , from the number of deaths, d_j . Let F denote the mean infection fatality ratio F for a population (see Section 3). Based on published estimates we approximate the distribution (among patients who die) of time from infection to death by a discretized gamma distribution with mean 14 days and standard deviation 6.01 days, plotted in Figure 2. Let P_k denote the probability that a newly infected person dies k days after infection. This distribution is normalized based on F so that

$$\sum_k P_k = F.$$

The daily deaths d_j are given by the discrete convolution of the new infections i_j with the distribution P :

$$d_j = \sum_{k=0}^{\infty} i_{j-k} P_k. \quad (3)$$

The new infections can thus be found by a discrete deconvolution, which is ill-conditioned (i.e., the result is highly sensitive to small changes in the data). Due to non-smoothness of real-world recorded death rates, the deconvolved infection rate is typically highly oscillatory and includes negative values.

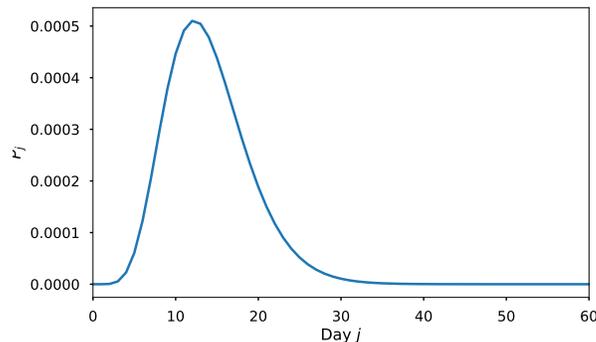


Figure 2: Probability distribution for time from infection to death.

2.2.1 Deconvolution algorithm for inferring infections from deaths

In order to obtain realistic values, we first smooth the daily death values using a 3-day moving average and then use a non-negative least squares approach with regularization [19]. To apply this approach to data recorded over J days, we write (3) in the matrix-vector form

$$\mathbf{d} = \mathbf{M} \mathbf{i}. \quad (4)$$

Here $\mathbf{d} = [d_j], j = 1, \dots, J$; \mathbf{M} is a $J \times J$ Toeplitz matrix with a first column given by $\mathbf{P} = [P_k], k = 1, \dots, J$; and $\mathbf{i} = [i_j], j = 1, \dots, J$. An estimate $\hat{\mathbf{i}}$ of the daily infections can be obtained by pursuing the following optimization [9, 27]:

$$\hat{\mathbf{i}} = \arg \min_{\mathbf{i}} \|\mathbf{M} \mathbf{i} - \mathbf{d}\|_2^2 + \lambda \|\mathbf{i}\|_2^2, \quad (5)$$

where $\|\cdot\|_2$ is the Euclidean norm and λ is a positive real-valued regularization parameter.

A standard non-negative least-squares solver can be used to perform the minimization in (5). We have found that this approach gives satisfactory results except during the very earliest stage of the epidemic, when it tends to significantly underestimate the number of infections. This effect can be attributed to the fact that the probability vector \mathbf{P} represents the average case, which introduces bias by underestimating the contributions from smaller daily infection occurrences. To remedy this, we apply the following perturbation based technique. We perform many random trials of generating the model (4) based on a perturbed version of the probability vector \mathbf{P} . Namely, we generate

$$\tilde{\mathbf{P}} = \mathbf{P} + \boldsymbol{\zeta},$$

where $\boldsymbol{\zeta}$ is a random noise vector. Then we use $\tilde{\mathbf{P}}$ to generate a convolution matrix $\tilde{\mathbf{M}}$ which is used to replace the matrix \mathbf{M} and obtain an estimate of \mathbf{i} using (5). The process is repeated for a sufficient number of trials. The final estimate of \mathbf{i} is formed by taking the average of the estimates obtained from all trials.

This deconvolution gives values of i_j up to the present day, but the most recent values are not accurate since they would be influenced significantly by unknown future values of deaths. We therefore use these values only up to $m = 21$ days in the past; this ensures

that for all dates of inference, at least 90% of the resulting deaths would have occurred by the present date.

Removal is modeled by a Poisson process in which the probability of recovery in a given one-day period is γ . The cumulative number of recovered individuals at day j is then computed as

$$R_j = \sum_{k=-\infty}^j i_k (1 - e^{-\gamma(j-k)}).$$

The number of actively infected persons on day j is then

$$I_j = \sum_{k=-\infty}^j i_k - R_j$$

The number of susceptibles on day j is

$$S_j = N - R_j - I_j$$

2.3 Part II: fitting recent past data

In order to estimate the infection rate up to the present date, we model the most recent m days using the SIR model (2). We fit the recent data by minimizing a weighted sum of squares of the difference between the logarithm of the predicted and observed deaths, where the weighting emphasizes the most recent values:

$$\text{Choose } q(t) \text{ to minimize } \sum_{j=-m}^0 (j+m) |\log(d_j) - \log(\tilde{d}_j)|^2$$

where the modeled number of deaths is

$$\tilde{d}_j = F\gamma I_j$$

with I_j given by the solution of (2) subject to the initial susceptible and infected fractions (S_{-m}, I_{-m}) , obtained from the inference model in part I. Two fits are performed; in the first, intervention is assumed to be constant:

$$q(t) = q_1.$$

In the second the intervention is assumed to be a linear function:

$$q(t) = q_2 + st.$$

The fit parameters are constrained so that $q(t)$ remains in the interval $[0, 1]$. To reduce the sensitivity to noise in the data, this fit includes a penalty proportional to the magnitude of the slope. In both fits, the fit is weighted to be more sensitive to the most recent data. We have found that the constant fit q_1 is more robust but the linear fit is better at adapting to recent changes in intervention.

2.4 Part 3: forecasting

From Part 2 we obtain estimate values (S_0, I_0, R_0) at the present day. These are combined with a hypothesized future intervention effectiveness $q(t)$ and fed into the augmented SIR model (2) to predict future dynamics of the pandemic. The model is not intended to predict the long-term intervention effectiveness, which is largely a matter of government policy. In the short term, as a compromise between robustness and flexibility, we estimate the intervention effectiveness by the average of the two fits from part 2, evaluated at the present day:

$$q(t) = \frac{q_1 + q_2}{2}.$$

Forecasts can also be generated using some assumed future intervention policy that changes in time.

3 Parameter values and data sources

Here we describe the main inputs and parameters of the model, along with the values used. For both inputs and parameters, the model can easily be used with other values than the specific sets chosen here. For instance, it might be useful to replace the officially reported COVID-19 deaths with excess mortality data.

3.1 Data

The primary model input is the number of confirmed COVID-19 deaths per day in a given region (a country or US state) since the start of the pandemic. The recorded deaths are obtained from:

- For countries: <https://github.com/CSSEGISandData>
- For states: <https://github.com/nytimes>

Because death counts are noisy in practice, we smooth these time series by using a 3-day moving average. In order to compute the mean IFR, we use population data from the UN [8].

3.2 Parameters

Next we discuss the key parameters and give the values used in the model.

- Mean infectious period: 14 days [18, 28]. This enters the model through $\gamma = 1/14$.
- Basic reproduction number σ_0 : 3.0 (2.0-4.0). This enters the model through $\beta = \sigma_0\gamma$. Early estimates of the basic reproduction number were low, with some even less than two [31], but some more recent estimates are closer to four [25].
- The infection fatality ratio F . This is the fraction of infected persons that die from COVID-19. Since the IFR for COVID-19 is known to vary strongly with age, we use age-specific IFR estimates from [29]. These are combined with UN demographic

Region	Reported	Estimated	Difference	Ref. [16]
China	80,763	409,353	328,590	12-89 million
Italy	8,889	381,011	372,122	30,223
Iran	7,248	398,766	391,518	266,213
South Korea	7,416	30,111	22,695	18,809
France	1,369	46,214	44,845	7,931
Spain	1,140	121,806	120,666	87,405
Germany	1,215	5,180	3,965	2,277
US	675	40,341	39,666	1.21 million

Table 1: Comparison of reported and estimated cumulative infections for several countries, as of March 9th. The last column shows, for comparison, estimated under-reporting from a different model in [16].

data for each country [8] to determine an effective IFR, where we assume a uniform attack rate across age groups:

$$F = \frac{1}{N} \sum_k F_k N_k.$$

Here F_k is the IFR for age group k and N_k is the population in age group k . For regions in which we lack detailed demographic data, we use a value of 0.6% based on estimates from [23, 31, 29].

4 Exploration of model outputs

4.1 Real vs. confirmed infections

We can use the inferred infection rates from Part I above to estimate the degree to which cases in a given region are under-reported. Table 1 shows this comparison for the date of March 9th. We compare the values from our model with those of [16], which used a very different approach.

In Figure 3 we compare confirmed cases with inferred numbers of infections. The values plotted are the daily increase in each number. For Spain, Italy, and the USA we observe a similar trend, with a huge gap (on the order of a 100x ratio) between the real and confirmed values, which later narrows to a factor of four or five. In contrast, South Korea (where the pandemic was rapidly controlled) and Saudi Arabia (where the pandemic has grown at a relatively slower rate) show a much higher detection rate early on, but similar rates later.

Note that the confirmed case counts include some negative values, due to after-the-fact revisions made by some reporting organizations. The avoidance of this kind of artifact is another benefit of using recorded deaths as the basis (which so far have been much less affected by such artifacts) for modeling. Still, confirmed deaths are not completely

accurate and may be supplemented by statistical estimations of excess mortality. In Figure 3 we show for Spain only the rate of infection inferred based on excess mortality⁶. As one might expect, it is more difficult to capture the early rise in infection as the deviation from expected mortality is at first not statistically significant. On the other hand, when the infection rate is high the excess mortality gives a more accurate accounting and implies higher rates of infection.

4.2 Estimating current immunity

Next we investigate the fraction of the population estimated to have antibodies (i.e., all of those who have been infected). Based on the behavior of other viruses, it is reasonable to expect that this portion of the population has some degree of immunity at present. Results for the 10 most-affected countries and the 10 most-affected US states are shown in Figure 4. Currently-infected individuals are included in this count.

We remark that the values given are no more accurate than the total confirmed deaths for each region. Since excess mortality records suggest that deaths are substantially under-reported, we expect that actual antibody levels are significantly higher than what is computed here. This also means that the most hard-hit regions might not even appear in these figures if they have greatly under-reported fatalities.

Nevertheless, the values obtained here are broadly in agreement with rough serological surveys that have been conducted recently, indicating for instance that around 3% of the Netherlands and 15% of New York residents have antibodies. While these surveys did not necessarily adhere to scientific standards of sampling, the general agreement gives some confidence in these estimates.

4.3 Estimating intervention effectiveness

Here we try to infer values of the intervention effectiveness parameter $q(t)$ based on the recorded data. The inference model above gives an estimate i of the number of new infections on day j , as well as estimates of S_j, I_j, R_j . According to the SIR model with intervention (2), the number of new infections each day should be approximately

$$i_j = (1 - q_j)\beta I_j \frac{S_j}{N}.$$

Thus

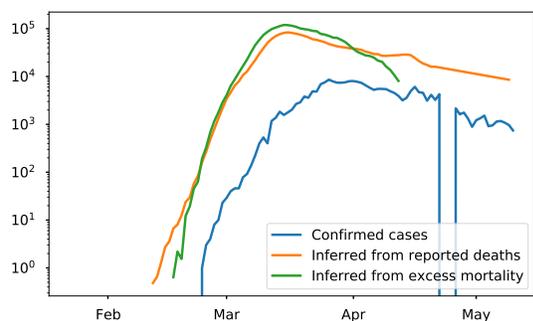
$$q_j = 1 - \frac{iN}{\beta S_j I_j}.$$

Some estimates of $q(t)$ are plotted in Figure 5.

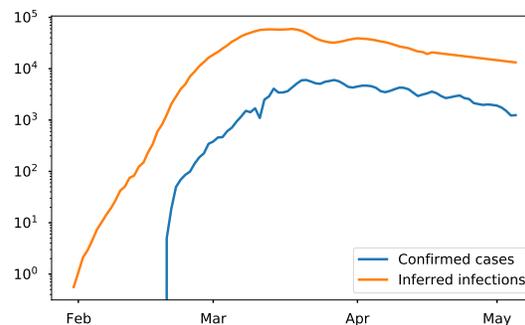
4.4 Comparison of forecasts with data

The model should not be viewed as a tool for giving absolute predictions of future infections and deaths. Rather, it provides predictions *conditioned* on an assumed future

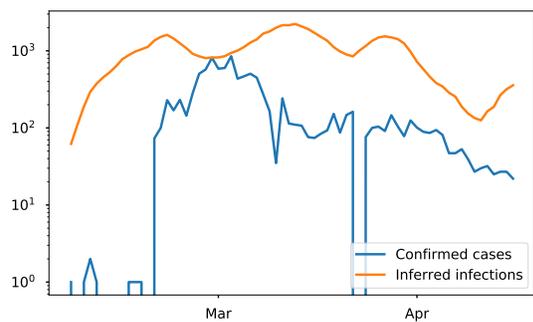
⁶Obtained from <https://momo.isciii.es/public/momo/data>. Accessed May 11th.



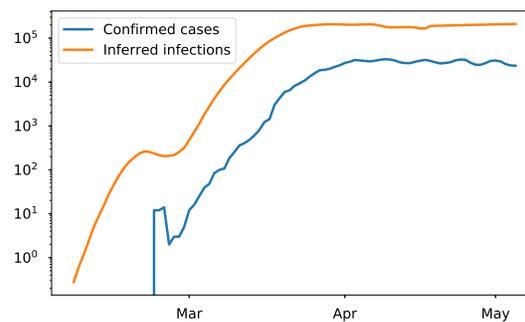
(a) Spain



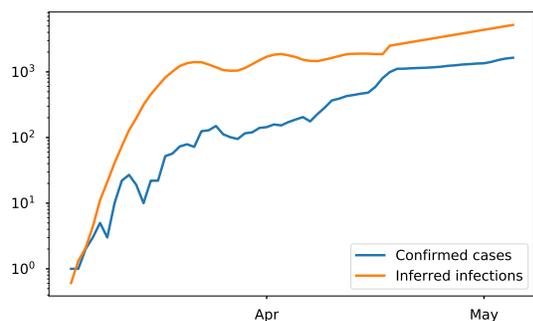
(b) Italy



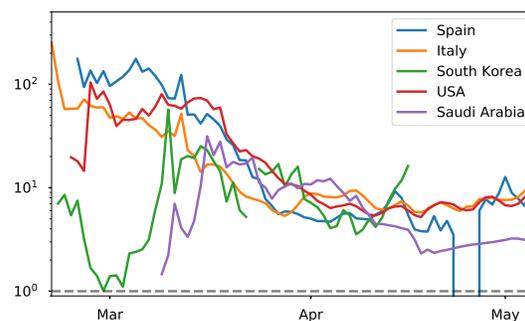
(c) South Korea. Recorded death numbers over the last 3 weeks are too small to provide a statistically significant basis to infer infections.



(d) USA



(e) Saudi Arabia



(f) Ratio of real to confirmed cases

Figure 3: Inferred infections vs. confirmed infections.

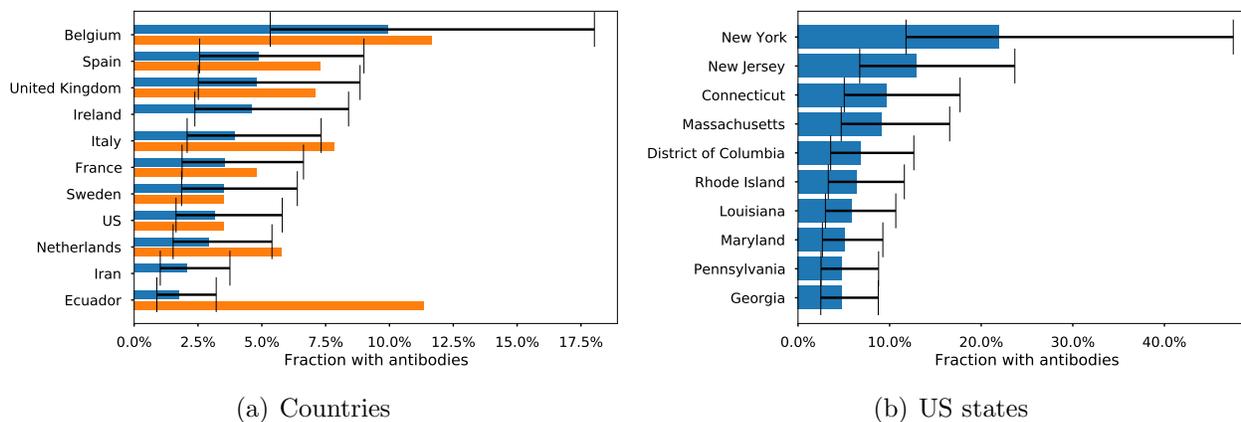


Figure 4: Estimated fraction of population with antibodies, showing the 10 countries and US states with the highest proportion of antibodies as of May 7th. Expected values and upper/lower bounds are based on infection fatality rates from [29]. Note the different horizontal scales. For countries, the orange bars show expected value corrected based on the ratio of excess mortality to recorded COVID-19 deaths. These numbers are not complete and are not available from all sources, but give some idea of how underreporting affects the estimates.

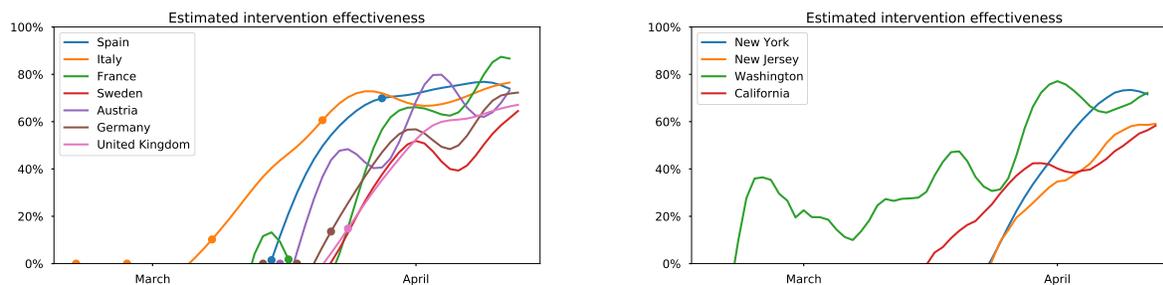


Figure 5: Estimated intervention effectiveness over time.

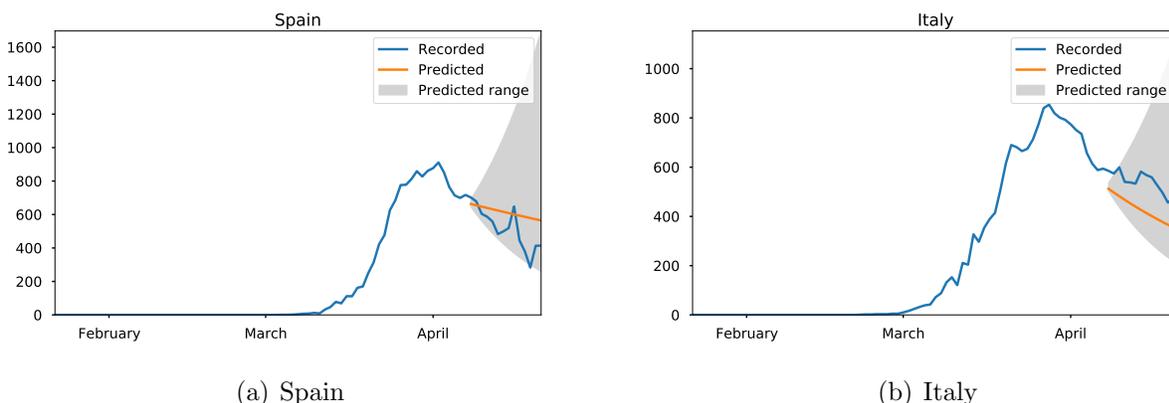


Figure 6: Predicted values obtained by omitting the most recent 14 days of data, compared with recorded values for daily deaths.

intervention effectiveness. We have equipped the model with a way to empirically estimate the intervention effectiveness q over the last 3 weeks, and a natural approach to prediction is to assume that the same value of q will apply in the future while interventions are maintained. This can also be used as a way to validate the model, by applying the model to a restricted data set in which the most recent values are omitted, and then *predicting* values for that recent time period. We expect this to give accurate results if there has been no policy change in the last few weeks. If intervention has changed significantly in the last few weeks, the results will be less accurate.

In Figure 6 we show results of this approach. We have omitted the last 14 days of data. We have intentionally chosen a difficult period for prediction, just after the death rate peaked in Spain and Italy.

4.5 Modeling a no-intervention scenario

The model can also be used to generate a hypothetical scenario starting at some time in the past. The main purpose of this feature is to examine how the epidemic might have progressed in a given region in the absence of intervention. This involves the following steps:

1. Choose a starting date for the model. This should be a date before when intervention started, but after a statistically significant number of deaths had occurred. For some regions (those that are very small or where the pandemic started very late) these two requirements may be contradictory.
2. Determine initial numbers of susceptible, infected, and recovered for the starting date. These are determined using the inference model described in Section 2.2.
3. Solve the initial value problem consisting of the SIR model (1) (without intervention), from the chosen starting date.

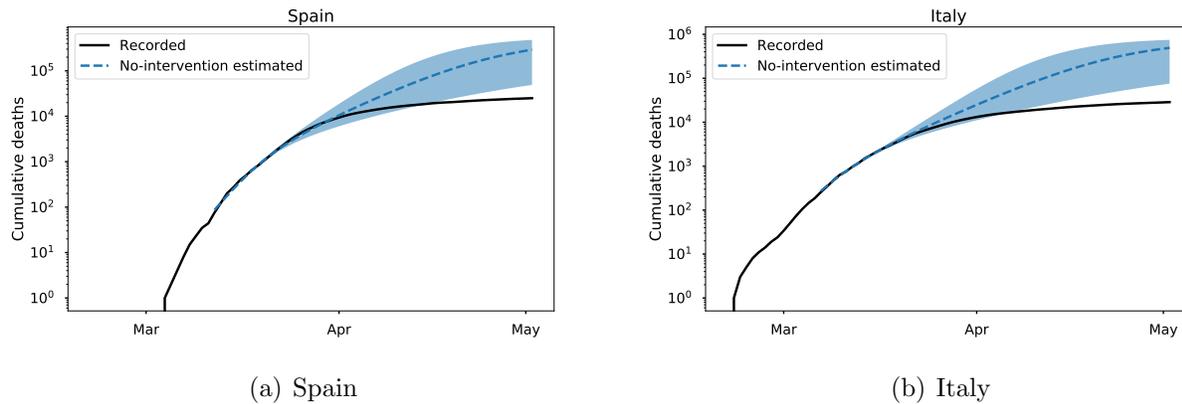


Figure 7: Estimated deaths in the absence of intervention, compared with recorded deaths.

Estimated deaths in the no-intervention scenario are compared with actual deaths in Figure 7 and Table 2.

5 Sources of error and uncertainty

The quote from statistician G.E.P. Box that "all models are wrong, but some models are useful" is certainly applicable here. The difficulty of modeling a pandemic of a new disease while it is ongoing is high, and, given the uncertainties involved, all results should be taken as rough estimates. The primary sources of error and uncertainty include:

- Uncertainty in the key parameters β, γ, F . Published scientific estimates of each of these values vary by at least a factor of 2, and two of these parameters enter into the exponent of the exponential growth of the pandemic. Therefore, even small uncertainties in these parameters can lead to very large uncertainties in the model outputs.
- Inaccuracies in the input data. Our model intentionally disregards confirmed case numbers because they are a very poor indicator of real infections due to inadequate testing and prevalence of mild or asymptomatic cases, as well as non-canonical testing standards [30] and false negatives depending on the diagnostic method used [16]. Numbers of deaths are typically much more accurate, but still contain substantial errors.
- Inaccurate model assumptions. The SIR model is one of the simplest epidemiological models. One of its key assumptions is that of homogeneous mixing among the population – i.e., that any two individuals are equally likely to have contact. This is of course far from true in the real world. More detailed models involving spatial structure and human mobility networks can partially remedy this deficiency. We have stuck to the SIR model so far because it seems that the parameter uncertainties already listed above – which will impact any epidemiological model – may be more significant than this modeling error.

Region	Starting date	Difference in estimated deaths
Austria	03-16-2020	10956 (796-48759)
Belgium	03-20-2020	34889 (554-82118)
Denmark	03-18-2020	7551 (651-30487)
France	03-17-2020	191034 (9088-502624)
Germany	03-14-2020	47879 (0-361409)
Italy	03-07-2020	463653 (46781-718723)
Norway	03-24-2020	3106 (441-13621)
Spain	03-15-2020	282825 (36757-453079)
Switzerland	03-20-2020	25562 (3361-61360)
United Kingdom	03-24-2020	164636 (15864-400730)

Table 2: Estimation of deaths averted as of May 4th, 2020, due to intervention. The start date for each simulation is based on when the epidemic began to grow and when nationwide intervention was put in place. The values in parentheses represent the Bayesian credible interval assuming $\sigma_0 \in [2.0, 4.0]$.

- Modeling the impact of intervention. Other models, such as that of [12], have attempted to explicitly model the effect of officially-announced intervention policies, lowering the contact rate by some amount starting when the policy is put in place. This requires an inherently uncertain estimation of the quantitative impact of a given policy, which will certainly vary in time and between different societies. To avoid this, we use an empirical assessment of intervention, choosing a contact rate that reproduces the data. This contact rate is somewhat sensitive to the noisy daily death rate data. Additionally, the model is slow to adapt to changes in intervention policy, since their effects do not show up in the death rate until at least 2-3 weeks after they are implemented.
- Assumptions about future policy. Our primary forecast assumes that the current level of intervention will be maintained over the entire forecast period. For long-range forecasts, this is very unlikely. The model is not intended to show what will likely happen in reality, since future intervention policy is a matter of political choices and not susceptible to detailed mathematical modeling. Rather, the forecast values should be interpreted as showing what will likely happen under a specified intervention policy.

References

- [1] COVID-19 Confirmed and Forecasted Case Data. <https://covid-19.bsvgateway.org/>. Accessed: 2020-09-10.
- [2] COVID-19 Epidemic Forecasting. <https://renkulab.shinyapps.io/COVID-19-Epidemic-Forecasting/>. Accessed: 2020-09-10.
- [3] COVID-19 Projections. <https://covid19.healthdata.org/>. Accessed: 2020-09-10.

- [4] delphi epidemiological case predictions.
- [5] MODELING COVID-19 IN THE UNITED STATES. <https://covid19. gleamproject.org/>. Accessed: 2020-09-10.
- [6] Predictive Monitoring of COVID-19. <https://ddi.sutd.edu.sg/>. Accessed: 2020-09-10.
- [7] UCLAML Combating COVID-19. <https://covid19.uclaml.org/>. Accessed: 2020-09-10.
- [8] UN world population prospects 2019. <https://population.un.org/wpp/Download/Standard/Population/>. Accessed: 2020-04-29.
- [9] T. Ballal, M. A. Suliman, and T. Y. Al-Naffouri. Bounded perturbation regularization for linear least squares estimation. *IEEE Access*, 5:27551–27562, 2017.
- [10] Quoc Trung Bui, Josh Katz, Alicia Parlapiano, and Margot Sanger-Katz. What 5 coronavirus models say the next month will look like. *The New York Times*, Apr 2020.
- [11] Matteo Chinazzi, Jessica T Davis, Marco Ajelli, Corrado Gioannini, Maria Litvinova, Stefano Merler, Ana Pastore y Piontti, Kumpeng Mu, Luca Rossi, Kaiyuan Sun, et al. The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. *Science*, 2020.
- [12] Seth Flaxman, Swapnil Mishra, Axel Gandy, H Unwin, H Coupland, T Mellan, H Zhu, T Berah, J Eaton, P Perez Guzman, et al. Report 13: Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries. 2020.
- [13] Joel Hellewell, Sam Abbott, Amy Gimma, Nikos I Bosse, Christopher I Jarvis, Timothy W Russell, James D Munday, Adam J Kucharski, W John Edmunds, Fiona Sun, et al. Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts. *The Lancet Global Health*, 2020.
- [14] William Ogilvy Kermack and Anderson G McKendrick. 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, 1927.
- [15] David I. Ketcheson. Optimal control of an sir epidemic through finite-time non-pharmaceutical intervention, 2020.
- [16] Steven G Krantz and Arni SR Srinivasa Rao. Level of under-reporting including under-diagnosis before the first peak of covid-19 in various countries: Preliminary retrospective results based on wavelets and deterministic modeling. *Infection Control & Hospital Epidemiology*, pages 1–8, 2020.
- [17] Adam J Kucharski, Timothy W Russell, Charlie Diamond, Yang Liu, John Edmunds, Sebastian Funk, Rosalind M Eggo, Fiona Sun, Mark Jit, James D Munday, et al. Early dynamics of transmission and control of COVID-19: a mathematical modelling study. *The Lancet infectious diseases*, 2020.

- [18] Stephen A Lauer, Kyra H Grantz, Qifang Bi, Forrest K Jones, Qulu Zheng, Hannah R Meredith, Andrew S Azman, Nicholas G Reich, and Justin Lessler. The incubation period of coronavirus disease 2019 (covid-19) from publicly reported confirmed cases: estimation and application. *Annals of internal medicine*, 2020.
- [19] Charles L Lawson and Richard J Hanson. *Solving least squares problems*, volume 15. SIAM, 1995.
- [20] Mariano Martini, Valentina Gazzaniga, Nicola Luigi Bragazzi, and Ilaria Barberis. The Spanish influenza pandemic: a lesson from history 100 years after 1918. *Journal of Preventive Medicine and Hygiene*, 60(1):E64, 2019.
- [21] Christina E Mills, James M Robins, and Marc Lipsitch. Transmissibility of 1918 pandemic influenza. *Nature*, 432(7019):904–906, 2004.
- [22] K Roosa, Y Lee, R Luo, A Kirpich, R Rothenberg, JM Hyman, P Yan, and G Chowell. Real-time forecasts of the COVID-19 epidemic in China from February 5th to February 24th, 2020. *Infectious Disease Modelling*, 5:256–263, 2020.
- [23] Timothy W Russell, Joel Hellewell, Christopher I Jarvis, Kevin Van-Zandvoort, Sam Abbott, Ruwan Ratnayake, Stefan Flasche, Rosalind M Eggo, Adam J Kucharski, CMMID nCov working group, et al. Estimating the infection and case fatality ratio for COVID-19 using age-adjusted data from the outbreak on the diamond princess cruise ship. *medRxiv*, 2020.
- [24] Eunha Shim, Amna Tariq, Wongyeong Choi, Yiseul Lee, and Gerardo Chowell. Transmission potential and severity of COVID-19 in South Korea. *International Journal of Infectious Diseases*, 2020.
- [25] Marlena M Siwiak, Pawel Szczesny, and Marian P Siwiak. From a single host to global spread. the global mobility based modelling of the COVID-19 pandemic implies higher infection and lower detection rates than current estimates. *The Global Mobility Based Modelling of the COVID-19 Pandemic Implies Higher Infection and Lower Detection Rates than Current Estimates (3/23/2020)*, 2020.
- [26] George Soper. The lessons of the pandemic. In *Global Health*, pages 185–190. Routledge, 2017.
- [27] Mohamed A. Suliman, Tarig Ballal, and Tareq Y. Al-Naffouri. Perturbation-based regularization for signal estimation in linear discrete ill-posed problems. *Signal Processing*, 152:35 – 46, 2018.
- [28] Kelvin Kai-Wang To, Owen Tak-Yin Tsang, Wai-Shing Leung, Anthony Raymond Tam, Tak-Chiu Wu, David Christopher Lung, Cyril Chik-Yan Yip, Jian-Piao Cai, Jacky Man-Chun Chan, Thomas Shiu-Hong Chik, et al. Temporal profiles of viral load in posterior oropharyngeal saliva samples and serum antibody responses during infection by sars-cov-2: an observational cohort study. *The Lancet Infectious Diseases*, 2020.
- [29] Robert Verity, Lucy C Okell, Ilaria Dorigatti, Peter Winskill, Charles Whittaker, Natsuko Imai, Gina Cuomo-Dannenburg, Hayley Thompson, Patrick GT Walker, Han Fu, et al. Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet Infectious Diseases*, 2020.

- [30] Yishan Wang, Hanyujie Kang, Xuefeng Liu, and Zhaohui Tong. Combination of rt-qpcr testing and clinical features for diagnosis of covid-19 facilitates management of sars-cov-2 outbreak. *Journal of medical virology*, 2020.
- [31] Joseph T Wu, Kathy Leung, Mary Bushman, Nishant Kishore, Rene Niehus, Pablo M de Salazar, Benjamin J Cowling, Marc Lipsitch, and Gabriel M Leung. Estimating clinical severity of COVID-19 from the transmission dynamics in wuhan, china. *Nature Medicine*, pages 1–5, 2020.