Tracking $\mathcal{R}$ of COVID-19

Simas Kučinskas*

Last updated: April 21, 2020

Abstract

We develop a new method for estimating the effective reproduction number of an infectious disease ($\mathcal{R}$) and apply it to track the dynamics of COVID-19. The method is based on the fact that in the SIR model, $\mathcal{R}$ is linearly related to the growth rate of the number of infected individuals. This time-varying growth rate is estimated using Kalman-filtering techniques from data on new cases and recoveries. The method is very easy to apply in practice, and it performs well even when the number of infected individuals is imperfectly measured, or the infection does not follow the SIR model. The estimates of $\mathcal{R}$ for COVID-19 are provided in an online dashboard, and they are used to assess the effectiveness of non-pharmaceutical interventions in a sample of European countries.

1 Introduction

The effective reproduction number ($\mathcal{R}$) plays a central role in the epidemiology of infectious diseases.[1] $\mathcal{R}$ is given by the average number of people infected by a single infectious individual. In benchmark models, the number of infected individuals increases as long as $\mathcal{R} > 1$. In an epidemic, real-time estimates of $\mathcal{R}$ are therefore essential for various policy decisions (Atkeson, 2020; Leung, 2020). Such estimates can be used to study the effectiveness of non-pharmaceutical interventions (NPIs), or assess what fraction of the population needs to be vaccinated to reach herd immunity (Chinazzi et al., 2020; Kucharski et al., 2020; Wang et al., 2020). Some social scientists have argued

*Humboldt University of Berlin (simas.kucinskas@hu-berlin.de, website). For comments and feedback, I am grateful to numerous colleagues. I also thank Chen Lin for excellent research assistance. Replication files are available online.

[1] For overviews of the benchmark models and the role of $\mathcal{R}$, see, among others, Chowell and Brauer (2009), Nishiura and Chowell (2009), Allen (2017), and Stock (2020).

that $\mathcal{R} < 1$ should be viewed as a fundamental constraint on public policy during the current COVID-19 pandemic (Budish, 2020).

In this paper we develop a new way to estimate $\mathcal{R}$ in real time. The method is based on the fact that in the benchmark SIR model (Kermack and McKendrick, 1927), $\mathcal{R}$ is linearly related to the growth rate of the number of infected individuals. In the first step of the procedure, we use data on new cases, recoveries, and deaths to construct a time series of how many individuals are infected at a given point in time. Then, we estimate the growth rate of this time series by Kalman-filtering techniques. In the final step, we use a theoretical relationship implied by the SIR model to invert this growth rate for $\mathcal{R}$. We document that the method is robust in the sense that the estimates of $\mathcal{R}$ remain fairly accurate even when new cases are imperfectly measured, or the true dynamics of the disease do not follow the SIR model.

We apply the new methodology to estimate $\mathcal{R}$ of COVID-19 in real time. As of April 20, The World Health Organization has confirmed more than 2.3 million cases of COVID-19 worldwide (World Health Organization, 2020). We use our estimates to calculate the basic reproduction number ($\mathcal{R}_0$) and take a first pass at evaluating the effectiveness of NPIs in a sample of 14 European countries. We demonstrate that the estimates of $\mathcal{R}$ are predictive of the future growth in COVID-19 deaths. The estimates of $\mathcal{R}$ for COVID-19 are provided in an online dashboard and can be explored interactively (link to dashboard).

While the new method is based on a simple idea—tracking the growth rate of the total number of infected individuals—it appears to be new. There are two broad classes of existing methods that can be used to estimate $\mathcal{R}$ in real time (Chowell and Brauer, 2009; Nishiura and Chowell, 2009). First, one can estimate a fully-specified epidemiological model and then construct a model-implied time series for $\mathcal{R}$. Second, one may use approaches that leverage information on the serial interval of a disease (i.e., time between onset of symptoms in a case and onset of symptoms in his/her secondary cases). Intuitively, imagine a disease with a deterministic serial interval of, say, three days. In that case, we could estimate $\mathcal{R}$ by simply dividing the number of new cases today by the number of new cases three days ago. Cori et al. (2013) exploit this idea to develop a Bayesian estimator that accounts for the randomness in the onset of infections as well as variation in the serial interval; see also Thompson et al. (2019).[2] This method is

[2] Wallinga and Teunis (2004) develop a related maximum-likelihood based method for estimating the so-called case reproduction number (Fraser, 2007). This approach has been implemented in an R software package `R0` by Obadia et al. (2012). However, as discussed by Cori et al. (2013), this approach cannot be used for real-time estimation of $\mathcal{R}$. The method proposed in the current paper estimates what is referred to as the instantaneous reproduction number in the literature.

implemented in a popular R package `EpiEstim`.[3]

The method proposed in this paper attempts to strike a balance between the two approaches mentioned above. Although our estimator is derived from standard epidemiological models, we use the smallest amount of theoretical structure that is necessary to obtain our estimator. In particular, the theoretical relationship that is used to derive our estimator is exactly valid not only in the benchmark SIR model, but also in the SIS model and a generalized SIR model with time-varying parameters and stochastic shocks. In this sense, our approach draws inspiration from the "sufficient statistics" literature in economics (Chetty, 2008). ($\mathcal{R}$ is also a sufficient statistic in itself in the sense used in this literature.) Relative to Cori et al. (2013), our estimator does not need any tuning parameters, and it does not require parametric assumptions on the distribution of new cases (such as assuming that new cases are Poisson distributed).[4] Our approach and its mathematical derivation shares some similarities with the Bayesian estimator proposed by Bettencourt and Ribeiro (2008).[5] Differently from these authors, we use data on the number of infected individuals, rather than the number of new cases, and we do not make parametric assumptions on the distribution of new cases. Finally, we note that in contrast to most existing methods, our method does not require Bayesian methods for estimation.

2 New Real-Time Estimator

We now derive our new estimator for the baseline SIR model (Kermack and McKendrick, 1927). As in Allen and Van Den Driessche (2008) or Stock (2020), we state the model in discrete time. While we develop our estimator for the basic SIR model in this section, in the Appendix we show that the same estimator obtains in the SIS model (Section A.1), and a generalized SIR model (Section A.2).

The standard SIR model describes the evolution of susceptible (S_t), infected (I_t),

[3] A team of researchers led by Timothy Churches and Nicholas Tierney has developed an online dashboard showing estimates of $\mathcal{R}$ for COVID-19 using `EpiEstim` for a number of countries (link).

[4] The method of Cori et al. (2013) assumes that $\mathcal{R}$ is constant over fixed windows of duration τ; the τ parameter effectively becomes a tuning parameter that needs to be chosen by the user.

[5] A team of analysts led by Kevin Systrom has developed an online dashboard tracking $\mathcal{R}$ for the US (both at the state and federal levels) using this estimator (link).

and recovered (R_t) individuals by the following equations:

$$\begin{aligned} S_t &= S_{t-1} - \beta I_{t-1} \frac{S_{t-1}}{N} \\ I_t &= I_{t-1} + \beta I_{t-1} \frac{S_{t-1}}{N} - \gamma I_{t-1} \\ R_t &= R_{t-1} + \gamma I_{t-1} \end{aligned} \tag{1}$$

The model is stated at a daily frequency. Here, N is population size, β is the daily transmission rate, and γ is the daily transition rate from infected to recovered. The recovered group consists of individuals who have either died or fully recovered.

The *basic reproduction number*, $\mathcal{R}_0$, is defined as $\mathcal{R}_0 \equiv \beta/\gamma$, and it gives the expected number of individuals infected by a single infective when everyone else is susceptible. The *effective reproduction number*, $\mathcal{R}_t$, is defined as $\mathcal{R}_t = \mathcal{R}_0 \times (S_{t-1}/N)$, and it equals the expected number of individuals infected by a single infective when a fraction (S_{t-1}/N) of individuals is susceptible.

From Eq. (1) the daily growth rate in the number of infected individuals is

$$\frac{I_t - I_{t-1}}{I_{t-1}} = \gamma \left(\mathcal{R}_0 \frac{S_{t-1}}{N} - 1 \right) = \gamma(\mathcal{R}_t - 1). \tag{2}$$

Denoting the estimated growth rate of infecteds by $\hat{\text{gr}}(I_t)$, and given a value for the transition rate γ, the plug-in estimator for the effective reproduction number is equal to

$$\hat{\mathcal{R}}_t = 1 + \frac{1}{\gamma} \hat{\text{gr}}(I_t). \tag{3}$$

The estimator in Eq. (3) is very intuitive. The estimate of $\mathcal{R}$ is above one if and only if the number of infectives is growing. To understand the role of multiplication by γ^{-1}, a simple numerical example may be helpful. Consider a situation with 100 initially infected individuals, an average duration of infectiousness of two days ($\gamma = 1/2$), and suppose that the true value of the effective reproduction number—which we want to estimate—is $\mathcal{R} = 3$. Under these assumptions, the 100 original infectives will infect on average $3 \times 100 = 300$ new people over the two days of infectiousness (150 infections per day). Starting with 100 infectives today, tomorrow there will therefore be on average

$$\underbrace{100}_{\text{initial infectives}} + \underbrace{150}_{\text{new infectives}} - \underbrace{50}_{\text{recovered initial infectives}} = 200 \text{ infectives},$$

or a growth rate of 100%. Plugging in this number into Eq. (3), we indeed obtain that $\mathcal{R} = 1+2\times1 = 3$. Hence, multiplication by γ^{-1} accounts for both recoveries of infected

individuals as well as the fact that infected individuals infect new people throughout their whole infectious period.

For the estimator in Eq. (3) to be feasible, we need to (i) calibrate the transition rate γ; and (ii) estimate the growth rate of I_t. To calibrate γ, external medical evidence can be used, given that γ^{-1} is the average infectious period.

To estimate the growth rate of I_t, we use data on new cases and new recoveries. For any epidemiological model, including the specific one in Eq. (1), we have by definition that

$$I_t = I_{t-1} + \text{new cases}_t - \text{new recoveries}_t.$$

Here, as in the model above, new recoveries consist of both full recoveries from the disease as well as deaths. We initialize I_t by $I_0 = C_0$ where C_0 is the number of total cases at some initial date (e.g., when 100 cases of the disease have been reached), and then construct subsequent values of I_t recursively.[6]

Given a time series for I_t, we use a standard state-space model to estimate the growth rate of I_t. Specifically, we specify a local-level model for the growth rate:

$$\begin{aligned} \text{gr}(I_t) &= \mu_t + \varepsilon_t, \qquad \varepsilon_t \sim \text{i.i.d. } \mathcal{N}(0, \sigma_\varepsilon^2) \\ \mu_t &= \mu_{t-1} + \eta_t, \quad \eta_t \sim \text{i.i.d. } \mathcal{N}(0, \sigma_\eta^2) \end{aligned} \tag{4}$$

We then estimate $\text{gr}(I_t)$ by the value for $\hat{\mu}_t$ given by the Kalman smoother (see Durbin and Koopman, 2012, Chapter 2), that is, $\hat{\text{gr}}(I_t) = \hat{\mu}_t$. To estimate the unknown parameters (σ_ε^2 and σ_η^2), in principle either classical frequentist methods or Bayesian techniques can be used. However, in practice, sample sizes for estimating $\text{gr}(I_t)$ are usually limited, especially early on in the epidemic. Hence, incorporating prior knowledge can lead to substantially better-behaved estimates. In our empirical application, we therefore use a mixture of frequentist and Bayesian methods, as explained in Section 4.1.

Intuitively, the procedure gives a smoothed version of the growth rate of I_t. The amount of smoothing is informed by the data via the estimation of σ_ε^2 and σ_η^2. Then, the Kalman smoother yields optimal estimates of the underlying growth rate. As is well known, the local-level model provides a model-based justification for exponentially weighted moving average smoothing and forecasting (Muth, 1960).

The local-level model can be viewed as a reduced-form specification for the growth rate of the number of infected individuals. The local-level model is sufficiently flexible

[6] We note that the (lagged) growth rate in the number of infectives can also be inferred from the current growth rate in the number of deaths from the disease. This approach may be less susceptible to problems stemming from ascertainment issues. This approach, however, leads to estimates that are substantially less timely, especially when the mean duration from the onset of symptoms to death is long—as is the case with COVID-19. Hence, we do not use this alternative in our empirical application.

to capture rich dynamic patterns in the data. In addition, in the Appendix (Section A.3), we provide a theoretical rationale for the local-level specification. In particular, we show that Eq. (4) arises naturally in a generalized SIR model (in the early stages of an epidemic) when the transmission rate follows a random walk.

3 Misspecification and Data Issues

We have derived our estimator in the context of the benchmark SIR model. A natural question is whether the method performs well if the dynamics of the disease are misspecified. In addition, our method requires data on new cases and recoveries. In practice, data on new cases and recoveries is often problematic, especially early on in the epidemic. Hence, a natural worry is that our estimates of $\mathcal{R}$ may be unreliable if, say, new cases are underreported. In this section, we investigate the effects of model misspecification and data problems on the performance of our estimator.

3.1 Model Misspecification

Our estimator remains exactly valid if the disease follows the SIS model (Section A.1 in the Appendix) or a generalized SIR model with time-varying transmission rates and stochastic shocks to the number of susceptibles and infectives (Section A.2). In addition, we investigate the performance of the estimator when the disease follows the SEIR model using a Monte Carlo simulation (Section A.4). We find that provided the number of days that individuals are infectious is correctly specified, our estimator yields accurate results even when the true model is SEIR rather than SIR.

3.2 Data Issues

We know discuss the effects of various data issues that are important in practice on the performance of our estimator.

Reporting delays. In practice, data may be subject to significant reporting delays. For example, suppose that due to testing constraints there is a lag of ℓ days between the date that an individual becomes infected and the date on which the case is registered. In this case, the estimates of $\mathcal{R}$ would also be subject to delay of ℓ days. If there are significant reporting delays, one may first obtain, say, one-week-ahead forecasts of new cases, recoveries, and deaths, and then use these forecasts to when constructing the time series for I_t.

Missing data on new recoveries. Sometimes the number of new recoveries is not observed, and only new cases and deaths are recorded. In such cases, one may construct a time series for I_t assuming that recoveries follow a Poisson distribution, using a calibrated recovery rate γ. A related data problem occurs when the number of new recoveries is underreported. In that case, our estimates of $\mathcal{R}$ will tend to be biased towards one in later stages of the epidemic. For instance, consider an extreme case in which recoveries are reported to be zero, and new cases are zero as well (because the epidemic has subsided). In that case, the growth rate of observed I_t is zero, and so our estimator would yield $\mathcal{R} = 1$. Constructing a synthetic time series for new recoveries using the Poisson distribution may serve as a useful robustness check in such situations.

Lumpy reported data. The reported data on new recoveries is often "lumpy." In particular, daily new recoveries are often reported to be zero (or very low), but when positive new recoveries are reported, they appear to count the recoveries that have occurred over a period of several days. In our empirical application, we therefore use a rolling average of new recoveries to construct I_t. One can also use alternative methods to redistribute the reported new cases over an interval of previous days (e.g., distribute equally over the last seven days).

Imperfect ascertainment. A natural worry with any estimator of $\mathcal{R}$ is that it may be substantially biased if ascertainment is imperfect. For COVID-19, the main worry is that of under-ascertainment. Given the simplicity of our estimator, we can analytically assess the effects of imperfect ascertainment.

Suppose that the true numbers of susceptible, infected, and recovered individuals are given by S_t^*, I_t^*, and $\mathcal{R}_t$, respectively. Their evolution is the same as in Eq. (1). However, we only observe $I_t = \alpha_t I_t^*$, where $\alpha_t \equiv I_t / I_t^*$ is the *ascertainment rate*. With this notation, we have that

$$\text{gr}(I_t) = \text{gr}(\alpha_t)[1 + \text{gr}(I_t^*)] + \text{gr}(I_t^*) \approx \text{gr}(\alpha_t) + \text{gr}(I_t^*),$$

since $\text{gr}(\alpha_t) \times \text{gr}(I_t^*) \approx 0$ at a daily frequency; the approximation is exact in continuous time. Using the approximation above and Eq. (2), we therefore obtain that the bias of the estimator under imperfect ascertainment is given by

$$\hat{\mathcal{R}}_t - \mathcal{R}_t \approx \frac{1}{\gamma} \text{gr}(\alpha_t).$$

We now discuss several cases of practical importance:

- Constant ascertainment rate ($\alpha_t = \alpha$). If the ascertainment rate is constant over time, then our estimator is unbiased, and $\hat{\mathcal{R}}_t = \mathcal{R}_t$. Hence, for example, even if we

only detect 10% of the infectives (but the fraction detected remains constant over time), the estimator remains unbiased. Note that if the number of tests increases over time, that is *not* inconsistent with $\alpha_t = \alpha$ given that the number of infectives is likely to be growing at the same time.

- Constant growth in the ascertainment rate ($\text{gr}(\alpha_t) = g_\alpha$). If the growth rate of α_t is constant over time, then our estimate of $\mathcal{R}_t$ is biased upwards if $g_\alpha > 0$ and downwards if $g_\alpha < 0$. Note, however, that we are often mostly interested in the *trend* of $\mathcal{R}$ over time and in particular whether the trend is affected by various policy interventions. The trend in $\mathcal{R}$ is estimated accurately even if $g_\alpha \neq 0$. Intuitively, a constant growth rate leads to a level bias, but the slope is still estimated correctly.
- Ascertainment rate converges over time ($\alpha_t \to \alpha$). The final case of interest occurs when the ascertainment rate converges to a constant over time. For example, if everyone is detected towards the end of the epidemic, we would have $\alpha_t \to 1$. Since our method uses Kalman-filtering techniques to estimate the growth rate of I_t, transient fluctuations in α_t would have a limited effect on the estimates of $\mathcal{R}$ later on in the sample. Given that we are often precisely interested in the behavior of $\mathcal{R}$ in the later stages of the epidemic (when the ascertainment rate is likely fairly constant), our method would still yield reliable estimates.

Overall, we conclude that the method appears to provide reliable estimates in a number of situations with imperfect ascertainment and data issues that are of practical relevance.

4 Results for COVID-19

We now apply the new methodology to track $\mathcal{R}$ of COVID-19 in real time. The estimates are provided in an online dashboard and can be explored interactively (link to dashboard).

We use data from the John Hopkins CSSE repository (Dong et al., 2020).[7] The dataset contains daily observations on COVID-19 cases, recoveries, and deaths for most countries around the world. For each country, we use data after the cumulative number of COVID-19 cases reaches 100. For our baseline results, we use $\gamma^{-1} = 7$ days, similarly to Maier and Brockmann (2020) who assume a mean infectious period of 8 days and Prem et al. (2020) who use a mean infectious period of either 3 or 7 days in their simulations. In addition, we find that $\gamma^{-1} = 7$ leads to estimates of the basic reproduc-

[7] The data is publicly accessible online (link).

Table 1
Priors

Notes: Priors used in the Bayesian estimation of $\mathcal{R}$. See text for description on how the priors for the precision of the irregular component ($1/\sigma_\varepsilon^2$) and the signal-to-noise ratio ($q \equiv \sigma_\eta^2/\sigma_\varepsilon^2$) are calibrated based on cross-country frequentist estimates.

Parameter	Prior
Precision of irregular component ($1/\sigma_\varepsilon^2$)	Gamma$(0.01731, 1.50 \times 10^{-5})$
Signal-to-noise ratio ($q \equiv \sigma_\eta^2/\sigma_\varepsilon^2$)	Gamma$(0.03083, 0.04122)$
Initial value (μ_0)	$\mathcal{N}(0.35, 0.5^2)$

tion number ($\mathcal{R}_0$) that are in line with the current best estimates (Liu et al., 2020). That said, we also investigate the effects of different choices for γ on our results.

The data on new recoveries provided by the John Hopkins CSSE team is "lumpy" in the sense that new recoveries are often reported to be zero (or very low). However, on a few days a high number of new recoveries is recorded. Arguably, the reason for such lumpiness is that new recoveries are only reported on some days, and whenever new recoveries are reported, they sum up recoveries that have occurred over the last few days. To avoid such lumpy observations influencing the estimates unduly, we smooth the time series of new recoveries by calculating 5-day rolling averages. An example of such a lumpy time series (for Denmark) and its 5-day rolling average is provided in the Appendix (Figure A.2). Otherwise, the time series for the number of infected individuals is constructed as described in Section 2. We acknowledge that using rolling averages introduces a certain amount of delay in our estimates $\mathcal{R}$ and makes them somewhat less timely.

The main results are presented in Section 4.2. The later sections apply the estimates to measure the basic reproduction number ($\mathcal{R}_0$), document that the estimated $\mathcal{R}$ numbers are predictive of future deaths, and use the estimates to assess the effectiveness of various policy interventions.

4.1 Estimation

To estimate $\mathcal{R}$, we use Bayesian filtering methods. We employ the following strategy to calibrate the prior distributions. First, we estimate a local-level model for $\text{gr}(I_t)$ using a frequentist Kalman filter with diffuse initial conditions. The procedure yields maximum likelihood estimates of σ_ε^2 (variance of the irregular component) and the signal-to-noise ratio $q \equiv \sigma_\eta^2/\sigma_\varepsilon^2$ for each country in the sample (with σ_η^2 denoting the variance of the level component). We then use the distribution of $\hat{\sigma}_\varepsilon^2$ and $\hat{q}$ across countries to calibrate

the priors for the precision of the irregular component ($1/\sigma_\varepsilon^2$) and the signal-to-noise ratio (q). To ensure that the priors are not too "dogmatic," we inflate the variance of the estimates by a factor of 3 when calibrating the prior distributions. We use a gamma prior for both the signal-to-noise ratio and the precision of the irregular component, and we calibrate the parameters of the gamma distribution by matching the expected value and variance of the gamma-distributed random variables to their sample counterparts. Finally, we use a fairly uninformative normal prior for the initial value of the smoothed growth rate. The resulting priors are given in Table 1.

Intuitively, these priors shrink the estimates of the precision and signal-to-noise ratio for each country towards their grand mean (average across countries). Such Bayesian shrinkage ensures that the parameter estimates are well behaved even though the sample size for many countries is fairly small, and the data are often noisy. We use the Stan programming language (Gelman, Lee, and Guo, 2015) to specify and estimate the Bayesian model.[8]

4.2 Main Results

Our estimates of $\mathcal{R}$ for COVID-19 for the world as a whole are provided in Figure 1. The graph highlights two distinct phases of the pandemic, first in China and later in Europe and the US. At the beginning of the sample in late January, the estimates of $\mathcal{R}$ are above 3; the estimates decline towards one at the end of February, suggesting a containment of the epidemic in China. Note that there is a moderate upwards jump in the estimated $\mathcal{R}$ around the second week of February. This jump was caused by a temporary change in COVID-19 case definitions in the Hubei province in China; the new definition included clinically-diagnosed COVID-19 cases (Tsang et al., 2020). The estimates of $\mathcal{R}$ start increasing in March, coinciding with the spread of the pandemic to Europe and the US. Our estimates for the world indicate that $\mathcal{R}$ has been trending down at a fairly steady rate since mid-March. However, the rate of decline appears substantially lower than what was observed in late January and early February. The current point estimate of $\mathcal{R}$ for the world as a whole is still above one.

Figure 2 plots the estimated for $\mathcal{R}$ for China, Italy, and the US. For all countries, we only consider dates after the cumulative number of cases has reached 100. In the Appendix, we also provide a graph of the raw data for the growth rate of the number of individuals as well its smoothed version used for estimating $\mathcal{R}$ (Figure A.3). For all three countries, the estimated $\mathcal{R}$ is initially around 3.5. For China, the estimated $\mathcal{R}$ falls below one around the fourth week of February; the 65% credible bounds for $\mathcal{R}$

[8] In particular, we use the `pystan` interface to call Stan from Python (link).

Table 2

Countries With Highest and Lowest Current Values Of $\mathcal{R}$

Notes: Countries with currently highest and lowest values of estimated effective reproduction numbers of COVID-19. Estimates of $\mathcal{R}$ for April 20 are shown.

Highest $\mathcal{R}$		Lowest $\mathcal{R}$	
Singapore	2.24	Mauritius	0.09
Russia	1.94	Iceland	0.36
Mexico	1.89	Australia	0.39
Saudi Arabia	1.85	Austria	0.39
India	1.82	Vietnam	0.41
Peru	1.75	New Zealand	0.51
Oman	1.70	Finland	0.52
Qatar	1.70	Burkina Faso	0.52

Table 3

Estimates of the Basic Reproduction Number ($\mathcal{R}_0$)

Notes: Estimates of the basic reproduction number ($\mathcal{R}_0$) for a sample of 14 European countries. The countries included in the sample are Austria, Belgium, Denmark, France, Germany, Greece, Italy, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and United Kingdom. The basic reproduction number is calculated by averaging our estimates of the effective reproduction number in the first 7 days of the epidemic, where the start of the epidemic is defined as the day when the cumulative number of cases reaches 100.

Number of Days Infectious:	5	6	7	8	9	10
$\hat{\mathcal{R}}_0$	2.52	2.82	3.12	3.42	3.73	4.03
CI Lower Bound (95%)	2.07	2.28	2.49	2.71	2.92	3.13
CI Upper Bound (95%)	2.99	3.39	3.79	4.19	4.59	4.98

remained below unity for the whole month of March. However, $\mathcal{R}$ in China appears to have been drifting upwards since mid-March, and the current point estimate is above one. In Italy, the estimated $\mathcal{R}$ has been steadily falling since March but at a slower rate than what was observed in China. The current point estimate is just slightly above one. In the US, a striking difference is that the estimates of $\mathcal{R}$ remain fairly constant at around 3.3 for more than two weeks. One possibility is that the disease spread without much control during this period. An alternative explanation, however, is that the fraction of under-ascertained cases went down substantially in this period, inflating the estimates of $\mathcal{R}$ upward. The current point estimate of $\mathcal{R}$ in the US is still above one.

Figure 1

$\mathcal{R}$ of COVID-19: Estimates for the World

Notes: Estimates of the effective reproduction rate ($\mathcal{R}$) of COVID-19 for the world as whole. The sample consists of all dates after the total number of reported cases worldwide has reached 100. 65% and 95% credible bounds shown by the shaded areas.

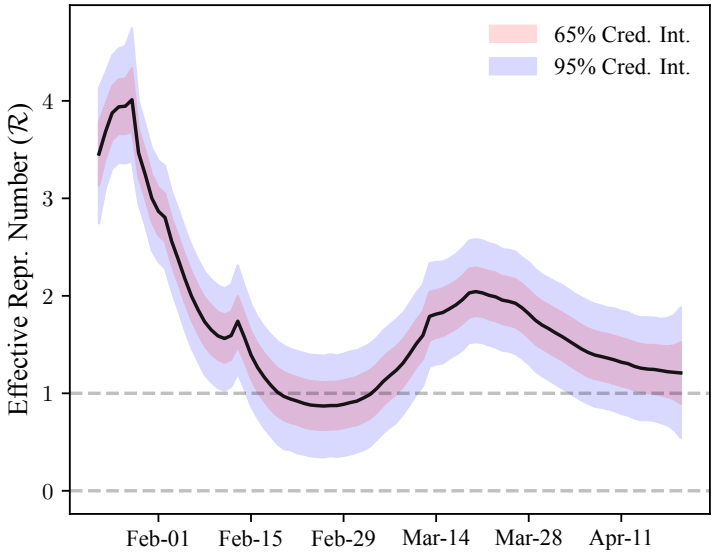

Figure 2

$\mathcal{R}$ of COVID-19: Selected Countries

Notes: Estimates of the effective reproduction rate ($\mathcal{R}$) of COVID-19 for selected countries. The sample consists of all dates after the total number of reported cases in the country has reached 100. 65% credible bounds shown by the shaded areas.

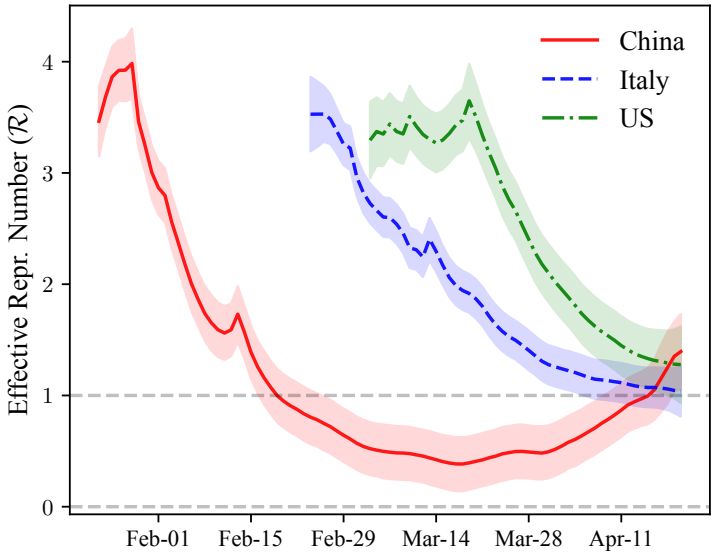

In Table 2, we report countries with the currently highest and lowest estimated effective reproduction numbers. The table illustrates how estimates of $\mathcal{R}$ can inform policymakers in tracking the effectiveness of various NPIs. The results also highlight that there is a substantial amount of heterogeneity in $\mathcal{R}$ across countries. While some of the heterogeneity is due to the epidemic being at different stages in the various countries, some of the differences are likely caused by differences in policy responses.

4.3 Basic Reproduction Number

Our estimates can also be used to measure the basic reproduction number ($\mathcal{R}_0$). The basic reproduction numbers gives the average number of individuals infected by a single infectious individual when the population is fully susceptible.

We estimate $\mathcal{R}_0$ by the average value of $\mathcal{R}_t$ in the first week of the epidemic. As before, we define the start of the epidemic as the date when the cumulative count of confirmed COVID-19 cases reaches 100.

The results for a sample of 14 European counries are shown in Table 3; the set of countries is the same as analyzed by Flaxman et al. (2020a).[9] Under our baseline assumption that the individuals are on average infectious for 7 days ($\gamma = 1/7$), we obtain an estimate of $\mathcal{R}_0 = 3.12$ (95% CI: 2.49–3.79). A recent metastudy has estimated an average $\mathcal{R}_0$ of 3.28 for COVID-19 (Liu et al., 2020) which is very close to our baseline point estimate.

The estimate of $\mathcal{R}_0$ depends on the average number of days an individual is infectious, with an additional day of infectiousness estimated to increase $\mathcal{R}_0$ by approximately 0.30. Such sensitivity, of course, is to be expected from Eq. (3), and it mirrors the sensitivity of the existing estimates of $\mathcal{R}_0$ of COVID-19 to assumptions on the mean serial interval (for example, see Figures 13 and 14 in Flaxman et al. (2020a)).

4.4 $\mathcal{R}$ and Future Deaths

A potential concern with our estimates is that since they are largely based on data on new cases, they may be misleading if new cases are subject to measurement problems. To help assuage this concern, we perform the following exercise. We ask whether *current* values of $\mathcal{R}$ help predict *future* growth in deaths. Since deaths are likely to be measured more accurately, this exercise provide an additional test of whether our estimates contain meaningful information and are not contaminated by data problems.

[9] The updated results of that study are available online (link). See also Flaxman et al. (2020b). The countries included in the sample are Austria, Belgium, Denmark, France, Germany, Greece, Italy, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and United Kingdom.

Formally, we consider the following regression:

$$\text{gr}(d_{i,t+1}) = \alpha_i + \beta(\hat{\mathcal{R}}_{i,t} - 1) + u_{i,t},$$

where i denotes a particular country, and t indexes calendar weeks. Although our original data is daily, we aggregate to a weekly frequency; otherwise, measures of the growth rate of new deaths are too noisy. In addition, we only include weeks after the cumulative number of COVID-19 deaths has reached 50. Given that we have panel data, we can include country fixed effects α_i to account for time-invariant unobserved heterogeneity (such as differences in average age—a key correlate of COVID-19 mortality (Verity et al., 2020)—or family structures). The relationship given above is predicted by the baseline SIR model.[10]

The relationship is shown in Figure 3. In the scatter plot, both variables are residualized to remove country fixed effects. We observe a strong positive relationship between the value of $\mathcal{R}$ this week and the growth in deaths one week later (corr. = 0.67). In the Appendix (Figure A.4), we demonstrate that there is also positive correlation (corr. = 0.41) between $\mathcal{R}$ and deaths two weeks later.

4.5 Assessing Non-Pharmaceutical Interventions

As the final application of our methodology, we take a first pass at assessing the effects of non-pharmaceutical interventions (NPIs) in a sample of 14 European countries (see Footnote 9) studied by Flaxman et al. (2020a). The definitions of NPIs and their introduction dates are provided by Flaxman et al. (2020a).[11]

We first perform an event-study type analysis. Specifically, we consider the behavior of $\mathcal{R}$ one week before and three weeks after a particular NPI is introduced. To ensure that the results are not confounded by selection effects, we only include countries for which we have data on $\mathcal{R}$ throughout the full event window. For our regression analysis later on, we use all the available data (not just three weeks after the introduction of an NPI).

In Figure 4, we plot the estimated values of $\mathcal{R}$ around the introduction of a lockdown. For definitions and dates of lockdowns, we refer to the original study of Flaxman et al. (2020a). We observe that $\mathcal{R}$ declines substantially after a lockdown is put in place, going from around 2.36 on the day of the intervention to around 1.45 two weeks later.

[10] Specifically, consider Eq. (2). Letting $\text{CFR} = d_t/I_{t-\ell}$ denote the case fatality rate (assumed to be constant over time), with ℓ standing for the average time between becoming infected and death, we have that $\text{gr}(d_t) = \text{gr}(I_{t-\ell})$, yielding the regression equation above.

[11] The data are publicly available in an online repository (link).

Figure 3
$\mathcal{R}$ and Future Deaths

Notes: Relationship between current estimates of the effective reproduction number ($\mathcal{R}$) and the growth rate of the number of new deaths in one week. The data is aggregated to a weekly frequency. Both variables are residualized to subtract country fixed effects by performing the within transformation. Only data after the cumulative number of deaths reaches 50 is included in the scatter plot. We remove data for the week of 2020-04-13–2020-04-19 in China that contain a large number of deaths that were previously unrecognized.

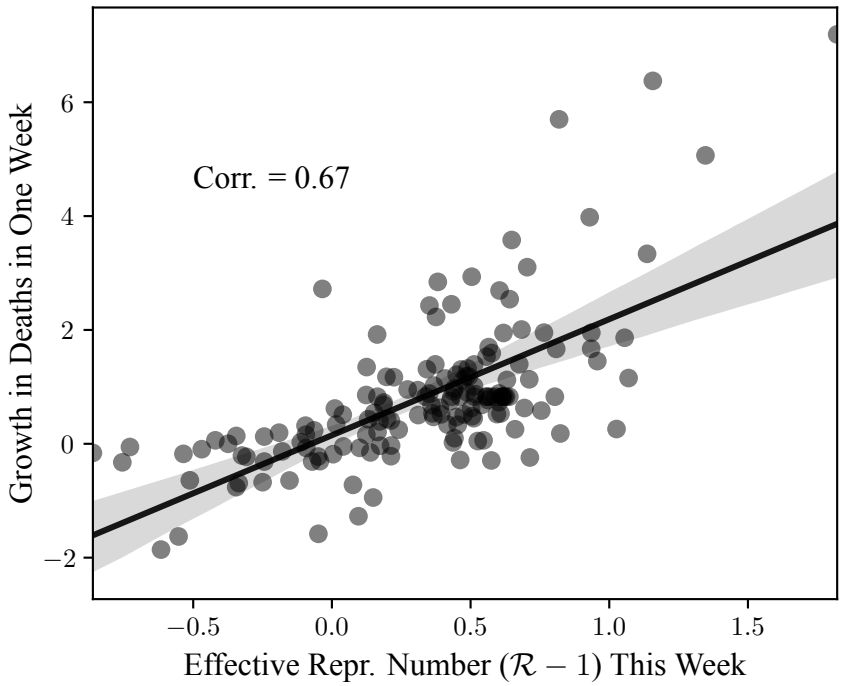

However, the graph shows that $\mathcal{R}$ is also on a downwards trend before the lockdown. In particular, there does not appear to be a visually detectible break in the trend at the date of the lockdown (i.e., there is no "kink" after the introduction of a lockdown). In the Appendix, we show that a similar pattern is observed for the remaining types of NPIs studied by Flaxman et al. (2020a). Specifically, we document the behavior of $\mathcal{R}$ around the introduction of public-event bans (Figure A.5), case-based measures (such as self-isolation whenever feeling ill and experiencing fever; Figure A.6), school closures (Figure A.7), and social-distancing measures (Figure A.8). With the possible exceptions of school closures and public-event bans, there is no visually apparent break in the trend around the date of the policy intervention.

To assess the effects of NPIs more formally and conduct statistical tests, we employ the following panel-data regressions, summarized in Table 4. Specifically, we regress

Figure 4
$\mathcal{R}$ and Policy Interventions: Lockdowns

Notes: The graph plots the estimated effective reproduction number ($\mathcal{R}$) one week before and three weeks after a lockdown is introduced in a country. The original sample consists of 14 European countries studied by Flaxman et al. (2020a). For the event-study graph, we restrict the sample to countries for which data on $\mathcal{R}$ is available for the whole event window. Heteroskedasticity-robust confidence bounds are shown by the shaded areas.

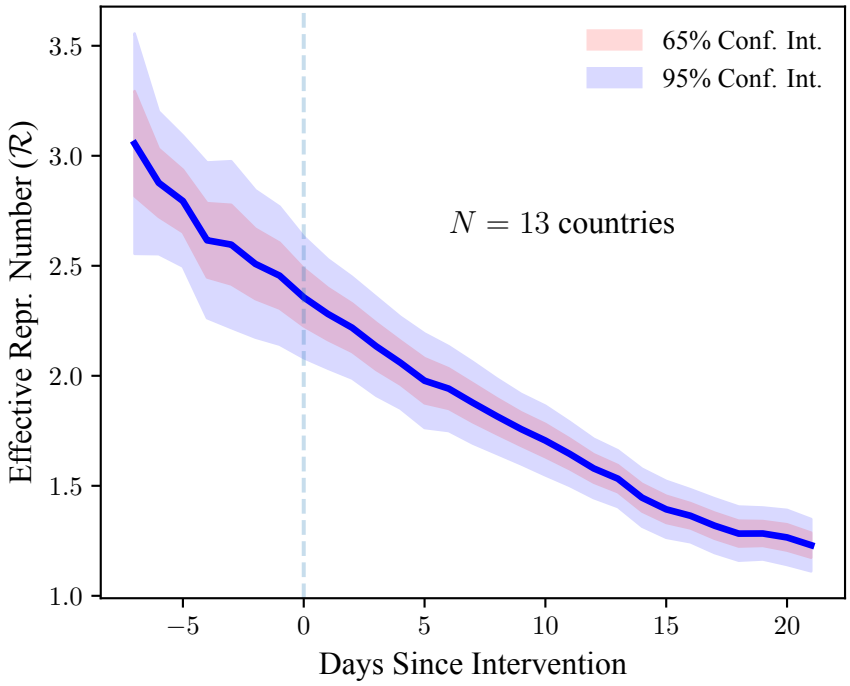

$\mathcal{R}$ (or its log) on a set of indicator variables and different types of fixed effects:

$$R_{i,t} = (\text{fixed effects}) + \sum_{j=1}^{5} \beta_i \text{NPI}_{i,t}^{(j)} + u_{i,t}.$$

Here, $\text{NPI}_{i,t}^{(j)}$ is an indicator variable that equals 1 after the j-th NPI is introduced, and zero before its introduction. We study five types of NPIs for which Flaxman et al. (2020a) have collected data. The index i stands for the country, and t denotes the number of days since the outbreak of the epidemic (i.e., number of days since 100 total confirmed cases have been reached). The regression is estimated using OLS, and heteroskedasticity-robust standard errors are reported.

Columns (1) and (2) of Table 4 provide estimates of the effects of NPIs when only country fixed effects are included. Irrespective of whether a specification in levels or in logs is considered, we find a strong negative effect of lockdowns, social distancing, and measures of self isolation. Bans of public events and school closures are not statis-

Table 4

Effective Reproduction Number After Introduction of NPIs

Notes: Results of panel-data regressions of the effective reproduction number ($\mathcal{R}$) on indicator variables that are equal to 1 after the introduction of a non-pharmaceutical intervention (NPI) and 0 before the introduction. The sample consists of 14 European countries studied by Flaxman et al. (2020a). Regressions always include country fixed effects; regressions in columns (3)–(6) also include days-since-outbreak fixed effects to control for the dynamics of the epidemic. Outbreak is defined as the date on which 100 cases of COVID-19 are reached. The log specification uses $\log(R)$ as the left-hand side variable, while the non-log specification uses the level of $\mathcal{R}$. The regression with testing controls in (6) controls for the change in the number of daily tests per capita conducted in the country. To allow for reasonably precise estimation of days-since-outbreak fixed effects, we only consider days after the outbreak for which we have data for at least 5 countries. Heteroskedasticity-robust standard errors in parentheses.

	(1)	(2)	(3)	(4)	(5)	(6)
Lockdown	-0.82***	-0.46***	-0.09	-0.03	0.02	0.08
	(0.06)	(0.03)	(0.07)	(0.03)	(0.08)	(0.05)
Public Events	-0.14*	-0.05	0.2***	0.16***	0.32**	0.26***
	(0.08)	(0.04)	(0.07)	(0.03)	(0.14)	(0.08)
School Closure	-0.09	0.02	-0.18**	-0.05	-0.1	-0.02
	(0.11)	(0.05)	(0.09)	(0.04)	(0.11)	(0.07)
Self Isolation	-0.39***	-0.11**	-0.26***	-0.05	-0.42***	-0.2***
	(0.11)	(0.05)	(0.08)	(0.03)	(0.09)	(0.05)
Social Distancing	-0.37***	-0.21***	-0.07	-0.06	0.2*	0.03
	(0.11)	(0.05)	(0.09)	(0.04)	(0.1)	(0.06)
N	616	616	616	616	353	353
R^2	0.68	0.6	0.87	0.88	0.94	0.91
Country FE	✓	✓	✓	✓	✓	✓
Log Specification		✓		✓		✓
Days-Since-Outbreak FE			✓	✓	✓	✓
Testing Controls					✓	✓

$^{*}p$ <0.1; $^{**}p$ <0.05; $^{***}p$ <0.01

tically significant. These regressions (as well the point estimates) are fairly similar to the statistical analysis performed by Flaxman et al. (2020a).

The regression with country fixed effects only, however, is likely misspecified. Implicitly, such a specification assumes that the only reason why $\mathcal{R}$ can fall is because of an introduction of NPIs (note that the regression has no other time-varying controls). However, $\mathcal{R}$ would likely trend downwards even in the absence of any public policy interventions. First, in standard epidemiological models, $\mathcal{R}$ tends to fall in an epidemic as the number of susceptibles is depleted. Second, people may adjust their behavior

even in the absence of any formal policy measures. Failing to control for the dynamics of $\mathcal{R}$ in the absence of NPIs therefore likely leads to an over-estimation of the effects of NPIs.

We immediately acknowledge that obtaining credible counterfactuals in the present empirical context is extremely challenging. However, we can exploit the panel structure of the dataset to take a first pass at controlling for the dynamics of the epidemic. We do so by introducing days-since-outbreak fixed effects to the baseline regression. Intuitively, with such fixed effects we are comparing $\mathcal{R}$'s in two countries that are both five days from the outbreak (say), with a school closure in one country but not in the other one.

The results from this alternative specification are shown in columns (3) and (4). We observe that the coefficients become substantially smaller in absolute terms and statistically insignificant for lockdowns, school closures, and measures of social distancing. The only coefficient that remains consistently negative is the one for self-isolation. The coefficient for public events is statistically significant but positive rather than negative. A naïve interpretation would suggest that banning public events has a positive effect on $\mathcal{R}$. More likely, however, is that the positive coefficient is due to countries where $\mathcal{R}$ is declining more slowly being faster to ban public events. (Note that the positive coefficient cannot be explained by countries that have a higher average $\mathcal{R}$ banning public events sooner, as time-invariant differences in $\mathcal{R}$ are captured by country fixed effects.)

A natural concern with this regression is that countries may introduce NPIs and simultaneously increase the number of tests for COVID-19 that they perform. The increased number of tests may spuriously show up as an increase in $\mathcal{R}$, masking any potential negative effect of NPIs. While it is very difficult to account for this issue in a fully satisfactory manner, we can at least control for changes in testing intensity. In column (6), we introduce the change in the daily number of tests per capita as an additional explanatory variable.[12] While the sample size is reduced somewhat as we do not have testing data for all countries in the sample, the results are largely unchanged.

We caution the readers against over-interpreting the results of this section. In particular, the results should *not* be taken to mean that the NPIs introduced in Europe have been largely ineffective. While the findings are not inconsistent with such a hypothesis, obtaining good estimates of the true causal impact of NPIs is extremely challenging. As a result, even our best estimates still suffer from important potential endogeneity issues, and more research is called for. In particular, the timing of NPIs is unlikely to be random. Instead, countries that introduced NPIs earlier likely did so because they had observed a stubbornly high $\mathcal{R}$. We also note that our Bayesian estimates of $\mathcal{R}$ use information from the full sample. Hence, estimates of $\mathcal{R}$ *after* the lockdown implicitly

[12] The data on daily tests per capita comes from *Our World in Data* (Our World in Data, 2020).

depend on the estimates of $\mathcal{R}$ *before* the lockdown. This feature of the estimation procedure may attenuate the effects of NPIs. In addition, even if the true causal impacts of NPIs are smaller than currently thought, that does *not* imply that social distancing is ineffective in reducing the spread of COVID-19. Instead, it is possible that people are voluntarily engaging in large amounts of social distancing even in the absence of formal policy interventions.

5 Conclusions and Limitations

In this paper we develop a new way to estimate the effective reproduction number of an infectious disease ($\mathcal{R}$). The new methodology is straightforward to apply in practice, and it appears to yield reliable results. We use the new method to track $\mathcal{R}$ of COVID-19.

The current paper faces several limitations. First, our estimator does not currently impose that the estimated $\mathcal{R}$ should be non-negative (neither for the point estimates nor the confidence intervals). Second, a local-level specification for the growth rate implicitly assumes that the growth rate of the number of infected individuals remains forever in flux. However, in the long-run, this growth rate must converge to zero. Since our model does not capture this feature, it seems likely that our estimated confidence bounds are overly conservative in the late stages of an epidemic. Third, when applying the model to cross-country data, one may achieve important gains in statistical efficiency if the model is estimated jointly for all countries. Finally, for tracking $\mathcal{R}$ of COVID-19, it would be highly desirable to have more timely reporting of new recoveries. For assessing the effects of NPIs more accurately, it is essential to collect data for a wider selection of countries.

Finally, it seems unlikely that any single statistical method would always perform best in estimating $\mathcal{R}$, especially given that real-time epidemiological data are often limited and fairly noisy. Instead of betting on any single model, it is arguably wiser to average estimates obtained from a suite of reasonable models. We hope that the method proposed in this paper may be part of such a suite.

References

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

Allen, L. J., Van Den Driessche, P., 2008. The basic reproduction number in some discrete-time epidemic models. Journal of Difference Equations and Applications 14 (10-11), 1127–1147.

Atkeson, A., 2020. What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios. NBER Working Paper Series, 25.

Bettencourt, L. M., Ribeiro, R. M., 2008. Real time bayesian estimation of the epidemic potential of emerging infectious diseases. PLoS ONE 3 (5).

Budish, E. B., 2020. R < 1 as an economic constraint: Can we "expand the frontier" in the fight against Covid-19? Working paper.

Chetty, R., 2008. Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. NBER Working Paper.

Chinazzi, M., Davis, J. T., Ajelli, M., Gioannini, C., Litvinova, M., Merler, S., Pastore Y Piontti, A., Mu, K., Rossi, L., Sun, K., Viboud, C., Xiong, X., Yu, H., Halloran, M. E., Longini, I. M., Vespignani, A., 2020. The effect of travel restrictions on the spread of the 2019 novel coronavirus (COVID-19) outbreak. Science 9757 (March), 1–12.

Chowell, G., Brauer, F., 2009. The basic reproduction number of infectious diseases: Computation and estimation using compartmental epidemic models. In: Chowell, G., Hyman, J. M., Bettencourt, L. M. A., Castillo-Chavez, C. (Eds.), Mathematical and Statistical Estimation Approaches in Epidemiology. Springer, Ch. 1, pp. 1–30.

Cori, A., Ferguson, N. M., Fraser, C., Cauchemez, S., 2013. A new framework and software to estimate time-varying reproduction numbers during epidemics. American Journal of Epidemiology 178 (9), 1505–1512.

Dong, E., Du, H., Gardner, L., 2020. An interactive web-based dashboard to track COVID-19 in real time. The Lancet 3099 (20), 19–20.

Durbin, J., Koopman, S. J., 2012. Time Series Analysis by State Space Methods. Oxford University Press, Oxford.

Ferguson, N. M., Laydon, D., Nedjati-Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunubá, Z., Cuomo-Dannenburg, G., Dighe, A., Fu, H., Gaythorpe, K., Green, W., Hamlet, A., Hinsley, W., Okell, L. C., Van, S., Thompson, H., Verity, R., Volz, E., Wang, H., Wang, Y., Walker, P. G. T., Walters, C., Winskill, P., Whittaker, C., Donnelly, C. A., Riley, S., Ghani, A. C., 2020. Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Working paper.

Flaxman, S., Mishra, S., Gandy, A., Unwin, J. T., Coupland, H., Mellan, T. A., Zhu, H., Berah, T., Eaton, J. W., Guzman, P. N. P., Schmit, N., Cilloni, L., Ainslie, K. E. C.,

Baguelin, M., Blake, I., Boonyasiri, A., Boyd, O., Cattarino, L., Ciavarella, C., Cooper, L., Cucunubá, Z., Cuomo-Dannenburg, G., Dighe, A., Djaafara, B., Dorigatti, I., Van Elsland, S., Fitzjohn, R., Fu, H., Gaythorpe, K., Geidelberg, L., Grassly, N., Green, W., Hallett, T., Hamlet, A., Hinsley, W., Jeffrey, B., Jorgensen, D., Knock, E., Laydon, D., Nedjati-Gilani, G., Nouvellet, P., Parag, K., Siveroni, I., Thompson, H., Verity, R., Volz, E., Gt Walker, P., Walters, C., Wang, H., Wang, Y., Watson, O., Xi, X., Winskill, P., Whittaker, C., Ghani, A., Donnelly, C. A., Riley, S., Okell, L. C., Vollmer, M. A. C., Ferguson, N. M., Bhatt, S., 2020a. Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries. Imperial College London (March).

Flaxman, S., Mishra, S., Gandy, A., Unwin, J. T., Coupland, H., Mellan, T. A., Zhu, H., Berah, T., Eaton, J. W., Guzman, P. N. P., Schmit, N., Cilloni, L., Ainslie, K. E. C., Baguelin, M., Blake, I., Boonyasiri, A., Boyd, O., Cattarino, L., Ciavarella, C., Cooper, L., Cucunubá, Z., Cuomo-Dannenburg, G., Dighe, A., Djaafara, B., Dorigatti, I., Van Elsland, S., Fitzjohn, R., Fu, H., Gaythorpe, K., Geidelberg, L., Grassly, N., Green, W., Hallett, T., Hamlet, A., Hinsley, W., Jeffrey, B., Jorgensen, D., Knock, E., Laydon, D., Nedjati-Gilani, G., Nouvellet, P., Parag, K., Siveroni, I., Thompson, H., Verity, R., Volz, E., Gt Walker, P., Walters, C., Wang, H., Wang, Y., Watson, O., Xi, X., Winskill, P., Whittaker, C., Ghani, A., Donnelly, C. A., Riley, S., Okell, L. C., Vollmer, M. A. C., Ferguson, N. M., Bhatt, S., 2020b. Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 14 European countries: Technical description update. Imperial College London (March).

Fraser, C., 2007. Estimating individual and household reproduction numbers in an emerging epidemic. PLoS ONE 2 (8).

Gelman, A., Lee, D., Guo, J., 2015. Stan: A probabilistic programming language for bayesian inference and optimization. Journal of Educational and Behavioral Statistics 40 (5), 530–543.

Kermack, W. O., McKendrick, A. G., 1927. A contribution to the mathematical theory of epidemics. Proceedings of the Royal Society A 115 (772), 700—-721.

Kucharski, A. J., Russell, T. W., Diamond, C., Liu, Y., Edmunds, J., Funk, S., Eggo, R. M., Sun, F., Jit, M., Munday, J. D., Davies, N., Gimma, A., van Zandvoort, K., Gibbs, H., Hellewell, J., Jarvis, C. I., Clifford, S., Quilty, B. J., Bosse, N. I., Abbott, S., Klepac, P., Flasche, S., 2020. Early dynamics of transmission and control of COVID-19: a mathematical modelling study. The Lancet Infectious Diseases 3099 (20), 1–7.

Leung, G., 2020. Lockdown can't last forever. Here's how to lift it. The New York Times (https://nyti.ms/3dWXHZR).

Liu, Y., Gayle, A. A., Wilder-Smith, A., Rocklöv, J., 2020. The reproductive number of COVID-19 is higher compared to SARS coronavirus. Journal of Travel Medicine 27 (2), 1–4.

Maier, B. F., Brockmann, D., 2020. Effective containment explains sub-exponential growth in confirmed cases of recent COVID-19 cases in China. Science 4557 (April), 1–8.

Muth, J. F., 1960. Optimal properties of exponentially weighted forecasts. Journal of the American Statistical Association 55 (290), 299–306.

Nishiura, H., Chowell, G., 2009. The effective reproduction number as a prelude to statistical estimation of time-dependent epidemic trends. In: Chowell, G., Hyman, J. M., Bettencourt, L. M. A., Castillo-Chavez, C. (Eds.), Mathematical and Statistical Estimation Approaches in Epidemiology. Springer, Ch. 5, pp. 103–122.

Obadia, T., Haneef, R., Boëlle, P. Y., 2012. The R0 package: A toolbox to estimate reproduction numbers for epidemic outbreaks. BMC Medical Informatics and Decision Making 12 (1).

Our World in Data, 2020. Coronavirus disease (COVID-19) – statistics and research. https://ourworldindata.org/coronavirus.

Prem, K., Liu, Y., Russell, T. W., Kucharski, A. J., Eggo, R. M., Davies, N., Jit, M., Klepac, P., Flasche, S., Clifford, S., Pearson, C. A. B., Munday, J. D., Abbott, S., Gibbs, H., Rosello, A., Quilty, B. J., Jombart, T., Sun, F., Diamond, C., Gimma, A., van Zandvoort, K., Funk, S., Jarvis, C. I., Edmunds, W. J., Bosse, N. I., Hellewell, J., 2020. The effect of control strategies to reduce social mixing on outcomes of the COVID-19 epidemic in Wuhan, China: a modelling study. The Lancet Public Health 2667 (20), 1–10.

Stock, J. H., 2020. Data gaps and the policy response to the novel coronavirus. https://www.nber.org/papers/w26902.

Thompson, R. N., Stockwin, J. E., van Gaalen, R. D., Polonsky, J. A., Kamvar, Z. N., Demarsh, P. A., Dahlqwist, E., Li, S., Miguel, E., Jombart, T., Lessler, J., Cauchemez, S., Cori, A., 2019. Improved inference of time-varying reproduction numbers during infectious disease outbreaks. Epidemics 29 (March).

Tsang, T. K., Wu, P., Lin, Y., Lau, E. H. Y., Leung, G. M., Benjamin, J., 2020. Impact of changing case definitions for COVID-19 on the epidemic curve and transmission parameters in mainland China. https://doi.org/10.1101/2020.03.23.20041319.

Verity, R., Okell, L. C., Dorigatti, I., Winskill, P., Whittaker, C., Imai, N., Cuomo-dannenburg, G., Thompson, H., Walker, P. G. T., Fu, H., Dighe, A., Griffin, J. T., Baguelin, M., Bhatia, S., Boonyasiri, A., Cori, A., Cucunubá, Z., Fitzjohn, R., Gaythorpe, K., Green, W., Hamlet, A., Hinsley, W., Laydon, D., Nedjati-gilani, G., Riley, S., Elsland, S. V., Volz, E., Wang, H., Wang, Y., Xi, X., Donnelly, C. A., Ghani, A. C., Ferguson, N. M., 2020. Estimates of the severity of coronavirus disease 2019: a model-based analysis. Lancet Infectious Diseases 3099 (20), 1–9.

Wallinga, J., Teunis, P., 2004. Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. American Journal of Epidemiology 160 (6), 509–516.

Wang, H., Wang, Z., Dong, Y., Chang, R., Xu, C., Yu, X., Zhang, S., Tsamlag, L., Shang, M., Huang, J., Wang, Y., Xu, G., Shen, T., Zhang, X., Cai, Y., 2020. Phase-adjusted estimation of the number of coronavirus disease 2019 cases in Wuhan, China. Cell Discovery 6 (1), 4–11.

World Health Organization, 2020. Coronavirus disease 2019 (COVID-19) situation report – 91. https://www.who.int/docs/default-source/coronaviruse/situation-reports/20200420-sitrep-91-covid-19.pdf.

Appendix A Additional Theoretical Results

A.1 SIS Model

We now show that the estimator in Eq. (3) also obtains when the dynamics of the disease follow the SIS model. The SIS model, again in discrete time, is given by

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

The only difference from the SIR model in Eq. (1) is that formerly infected individuals do not obtain immunity after recovery and instead again join the pool of susceptibles. As is well known, the basic reproduction number $\mathcal{R}_0$ in the SIS model is the same as in the SIR model (e.g., Chowell and Brauer, 2009) and given by $\mathcal{R}_0 = \beta/\gamma$. Since the law of motion for I_t in the SIS model is the same as in the SIR model, we can repeat the same steps as in the benchmark analysis to arrive at Eq. (3).

A.2 Generalized SIR Model

In this section, we show that the estimator in Eq. (3) also obtains in a much more general version of the SIR model. Specifically, we consider the following generalized SIR model:

$$S_t = S_{t-1} - \beta_t I_{t-1}\frac{S_{t-1}}{N} - v_{1,t}$$
$$I_t = I_{t-1} + \beta_t I_{t-1}\frac{S_{t-1}}{N} - \gamma I_{t-1} + v_{1,t} - v_{2,t}$$
$$R_t = R_{t-1} + \gamma I_{t-1} + v_{2,t}$$

Differently from the baseline model in Eq. (1), we now allow β_t to vary arbitrarily over time. In addition, we introduce random shocks $v_{1,t}$ and $v_{2,t}$. The shocks are i.i.d., and the time-varying support of $v_{1,t}$ is $[0, S_{t-1} - \beta_t I_{t-1}/N]$, while the support of $v_{2,t}$ is $[0, I_{t-1} + \beta_t I_{t-1} S_{t-1}/N]$. We also assume that $\mathbb{E}_{t-1}[v_{1,t} - v_{2,t}] = 0$, so that the conditional expectation $\mathbb{E}_{t-1}[I_t]$ coincides with the value for I_t given by the noiseless SIR model. With these modifications, the model can capture rich patterns of infectious disease dynamics. For example, "super spreader events" can be modeled either as $v_{1,t}$ shocks or as a spike in β_t. The model can also capture richer forms of population structures than the baseline SIR model. For example, if individuals who are more infectious (e.g., those with more connections in a network model) are more likely to become in-

fected first, that can be captured by assuming that β_t becomes lower over time.

Defining the (time-varying) basic reproduction number as $\mathcal{R}_0^{(t)} = \beta_t/\gamma$, and $\mathcal{R}_t \equiv \mathcal{R}_0^{(t)} S_{t-1}/N$, we obtain that

$$\text{gr}(I_t) = \gamma(\mathcal{R}_t - 1) + v_t,$$

where $v_t \equiv (v_{1,t} - v_{2,t})/I_{t-1}$. Taking expectations on both sides of the equation, we arrive at

$$\mathbb{E}[\mathcal{R}_t] = 1 + \frac{1}{\gamma}\,\mathbb{E}[\text{gr}(I_t)].$$

Hence, the generalized SIR model of the present section leads to the same estimator as the baseline SIR model in Eq. (1).

Finally, we note that if γ varies deterministically over time, the equation above remains essentially unchanged, the only difference being that γ is replaced by γ_t. If γ_t follows a non-degenerate stochastic process, then the estimator for $\mathbb{E}[\mathcal{R}_t]$ would need to correct for the covariance between γ_t and $\mathcal{R}_t$.

A.3 Foundation for the Local-Level Model

When estimating $\mathcal{R}$, we use a local-level specification for the growth rate of the number of infectives. In this section, we show that the local-level model arises naturally in an SIR model in the early stages of an epidemic when the transmission rate follows a random walk.

Specifically, consider the generalized SIR model of Section A.2. We now specialize the process for the transmission rate β_t to be a random walk:

$$\beta_t = \beta_{t-1} + \eta_t, \eta_t \sim \text{i.i.d. } \mathcal{N}(0, \sigma_\eta^2),$$

with a given initial value $\beta_0 > 0$. Using the law of motion for β_t, we calculate that

$$\text{gr}(I_t) = \left(\frac{S_{t-1}}{N}\right)\beta_t - \gamma + v_t \approx \beta_t - \gamma + v_t$$

in the early stages of the epidemic when $S_t \approx N$. Now define $\tilde{\beta}_t$ recursively by $\tilde{\beta}_t = \tilde{\beta}_{t-1} + \eta_t$ with the initial condition $\tilde{\beta}_0 = \beta_0 - \gamma$. Then, the growth rate of I_t follows a local-level model with

$$\begin{aligned} \text{gr}(I_t) &= \tilde{\beta}_t + v_t \\ \tilde{\beta}_t &= \tilde{\beta}_{t-1} + \eta_t \end{aligned}$$

Provided that the distribution of v_t can be approximated with a normal distribution, we directly obtain the specification in Eq. (4). Alternatively, to obtain an exact normal local-level model, we could assume that $v_{1,t} = v_{2,t} = 0$ (no shocks in the original model, just as in Eq. (1)) but that instead of observing the true growth rate $\text{gr}(I_t)$, we only observe $\text{gr}(I_t) + \varepsilon_t$ where ε_t is i.i.d. normally distributed mean-zero measurement error.

A.4 SEIR Model: Monte Carlo Simulation

Our estimation method uses a structural mapping between $\mathcal{R}_t$ and $\text{gr}(I_t)$ derived from the basic SIR model. While we can generalize the basic SIR model in several directions (Section A.2), and the estimator remains valid in an SIS model (Section A.1), the model is nevertheless restrictive. In particular, it ignores incubation periods as well as transmission during the incubation period. These features are likely especially important when modeling COVID-19.

We now perform a simulation exercise to see how our estimator of $\mathcal{R}_t$ performs in a richer model that accounts for these additional features. Specifically, we consider an SEIR model in which the exposed are infectious:

$$\begin{aligned} S_t &= S_{t-1} - \beta I_{t-1}\frac{S_{t-1}}{N} - \beta\epsilon E_{t-1}\frac{S_{t-1}}{N} \\ E_t &= E_{t-1} + \beta I_{t-1}\frac{S_{t-1}}{N} + \beta\epsilon E_{t-1}\frac{S_{t-1}}{N} - \kappa E_{t-1} \\ I_t &= I_{t-1} + \kappa E_{t-1} - \gamma I_{t-1} \\ R_t &= R_{t-1} + \gamma I_{t-1} \end{aligned} \tag{5}$$

Here, E_t denotes the number of individuals that are exposed at day t, κ is the daily transition rate from exposed to infected, and $\epsilon \in [0, 1]$ measures the degree to which the exposed are less infectious than the infected. If $\epsilon = 0$, the exposed are not infectious at all, and we obtain the benchmark SEIR model. If $\epsilon = 1$, the exposed are as infectious as the infected, and the model is isomorphic to the standard SIR model.

We calibrate the parameters following Wang et al. (2020) who apply the benchmark SEIR model (with $\epsilon = 0$) to study the dynamics of COVID-19 in Wuhan. In particular, we use $\kappa = 1/5.2$ and $\gamma = 1/18$ as in Wang et al. (2020). Then, we set $\epsilon = 2/3$, following Ferguson et al. (2020) who assume that symptomatic individuals are 50% more infectious than the asymptomatic (that is, $\epsilon^{-1} = 1.5$). Finally, we choose β by targeting a basic reproduction number of $\mathcal{R}_0 = 2.6$, again as in Wang et al. (2020). In the model above, $\mathcal{R}_0$ is given by $\mathcal{R}_0 = \beta/\gamma + \beta\epsilon/\kappa$, implying $\beta = \mathcal{R}_0\gamma\kappa/(\gamma\epsilon + \kappa)$. The formula yields $\beta \approx 0.12$. Finally, we set $S_0 = 11 \times 10^6$ (approximating the population

size of Wuhan), $E_0 = R_0 = 0$, and $I_0 = 1$.

The Monte Carlo design is as follows. First, we simulate the deterministic system in Eq. (5) using the parameters above. Then, we calculate the growth rate in the true number of infected individuals, i.e., $\text{gr}(I_t) = I_t/I_{t-1} - 1$. However, instead of knowing the true growth rate, the statistician is assumed to observe a noisy version of it given by $\widetilde{\text{gr}}(I_t) = \text{gr}(I_t) + \varepsilon_t$. Here, ε is an i.i.d. normal disturbance with mean zero and standard deviation of 0.10. The standard deviation of the disturbances is roughly equal to the range of the true growth rates. Hence, the amount of noise used in the simulation is fairly large. For each realization of the disturbances, we estimate $\mathcal{R}_t$ using our method. As in our empirical application, only data after 100 total cases have been reached is used.

We investigate two values for $\gamma_{\text{est.}}$ that are used when estimating $\mathcal{R}_t$ via Eq. (3). First, we consider a situation in which the statistician uses the correct time that individuals are infected, given by $\gamma_{\text{est.}} = (\gamma^{-1} + \kappa^{-1})^{-1}$ where γ and κ are the true parameter values of the SEIR model. Second, we investigate a case in which the statistician incorrectly things that individuals are infectious only for ten days ($\gamma_{\text{est.}} = 1/10$). We repeat the process for 10,000 Monte Carlo replications.

The results of the Monte Carlo simulation are shown in Figure A.1. When the statistician uses the correct number of days that an individual is infectious (that is, taking into account the incubation time), the estimates of $\mathcal{R}_t$ from our method are very close to their true theoretical values. That is in spite of the fact that our estimator for $\mathcal{R}_t$ is derived assuming that the dynamics of the disease are described by an SIR model. However, we also show that if the statistician misspecifies the number of days than an individual is infectious (assuming 10 days instead of the true number of 23.2 days), the estimates of $\mathcal{R}_t$ are substantially biased, especially in the early stages of the epidemic. As is to be expected from Eq. (3), underestimating the number of days that an individual is infectious leads to a downwards bias in the estimates of $\mathcal{R}_t$ early on in the epidemic (when $\mathcal{R}_t > 1$), and upwards bias when the true $\mathcal{R}_t$ falls below one. Overall, the results imply that the new method performs well when estimating $\mathcal{R}_t$ even when the true dynamics of the disease do not follow the SIR model, provided that the duration of infectiousness used in the estimation is sufficiently accurate.

Figure A.1

Monte Carlo Simulation: Effects of Misspecification

Notes: Estimates of the effective reproduction rate ($\mathcal{R}_t$) when the true dynamics of the disease follow an SEIR model. We investigate two values for $\gamma_{\text{est.}}$, the transition rate from infected to recovered, that are used when estimating $\mathcal{R}_t$. First, we use the correct value of $\gamma_{\text{est.}} = (\gamma^{-1} + \kappa^{-1})^{-1} \approx 0.043$. Second, we use a misspecified values of $\gamma_{\text{est.}} = 1/10$. Average values from 10,000 Monte Carlo replications are shown. See text for more details.

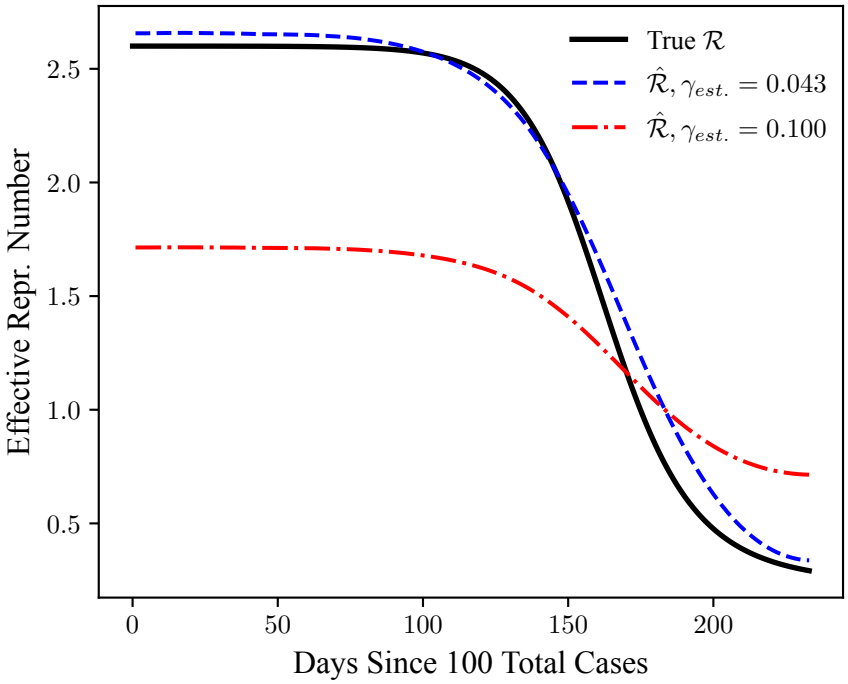

Appendix B Additional Figures

Figure A.2
New Recoveries in Denmark

Notes: An example of "lumpiness" in the original data for new recoveries in Denmark. The original time series and a 5-day rolling average are shown.

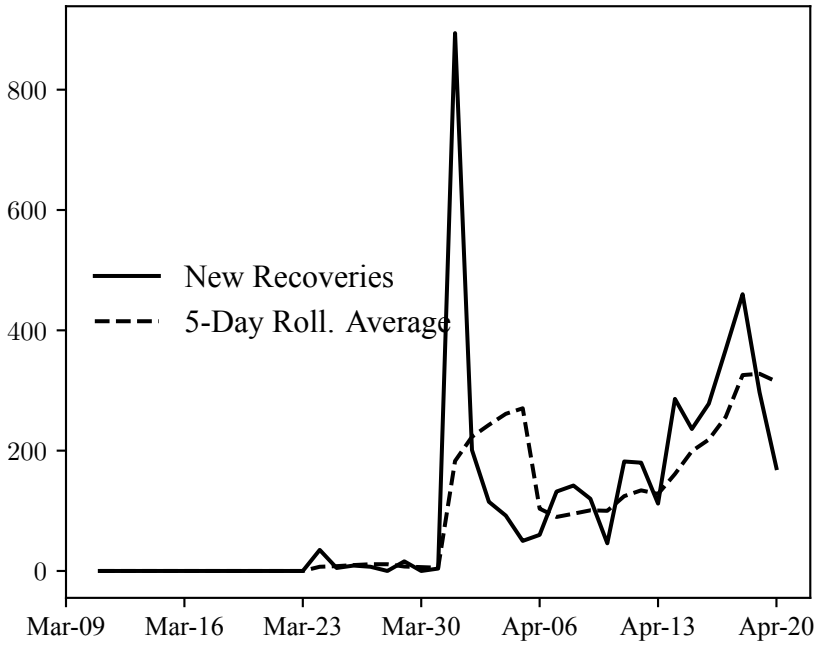

Figure A.3

Growth Rate of the Number of Infected Individuals

Notes: Raw data for the growth rate of the number of infected individuals (solid lines) and our estimate of its time-varying average (dashed lines) for China, Italy, and the US.

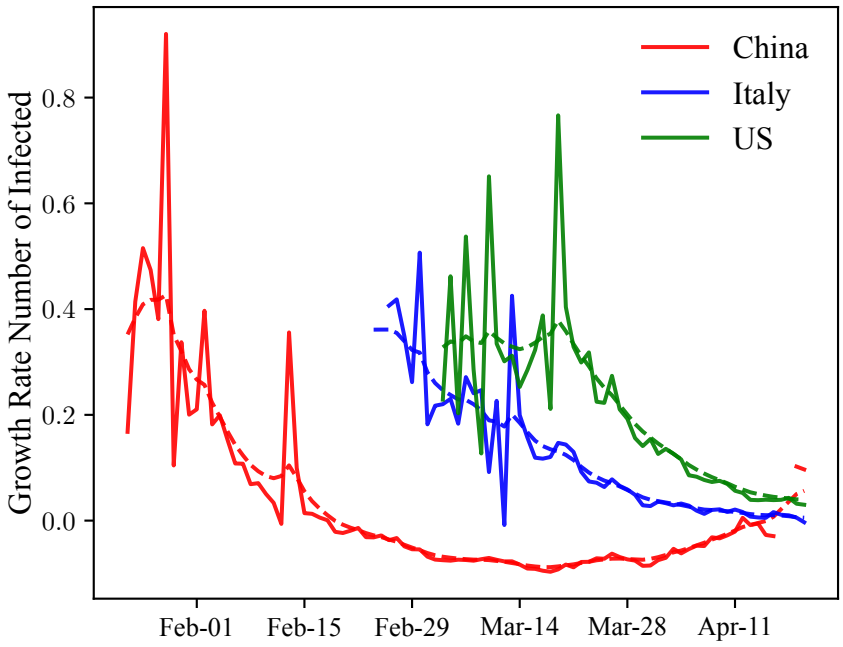

Figure A.4

$\mathcal{R}$ and Deaths in Two Weeks

Notes: Relationship between current estimates of the effective reproduction number ($\mathcal{R}$) and the growth rate of the number of new deaths in two weeks. The data is aggregated to a weekly frequency. Both variables are residualized to subtract country fixed effects by performing the within transformation.

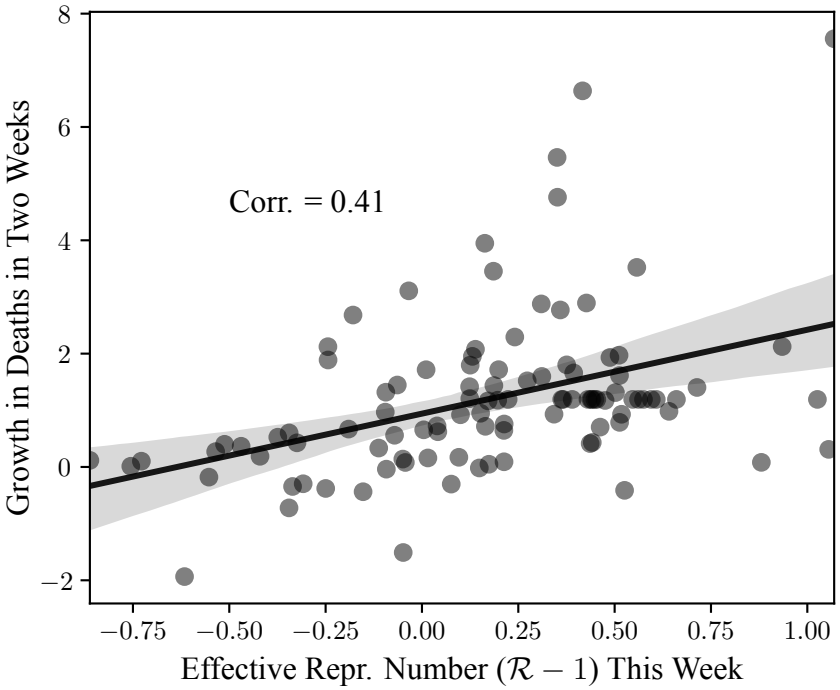

Figure A.5

$\mathcal{R}$ and Policy Interventions: Bans of Public Events

Notes: The graph plots the estimated effective reproduction number ($\mathcal{R}$) one week before and three weeks after public events are banned in a country. The original sample consists of 14 European countries studied by Flaxman et al. (2020a). For the event-study graph, we restrict the sample to countries for which data on $\mathcal{R}$ is available for the whole event window. Heteroskedasticity-robust confidence bounds are shown by the shaded areas.

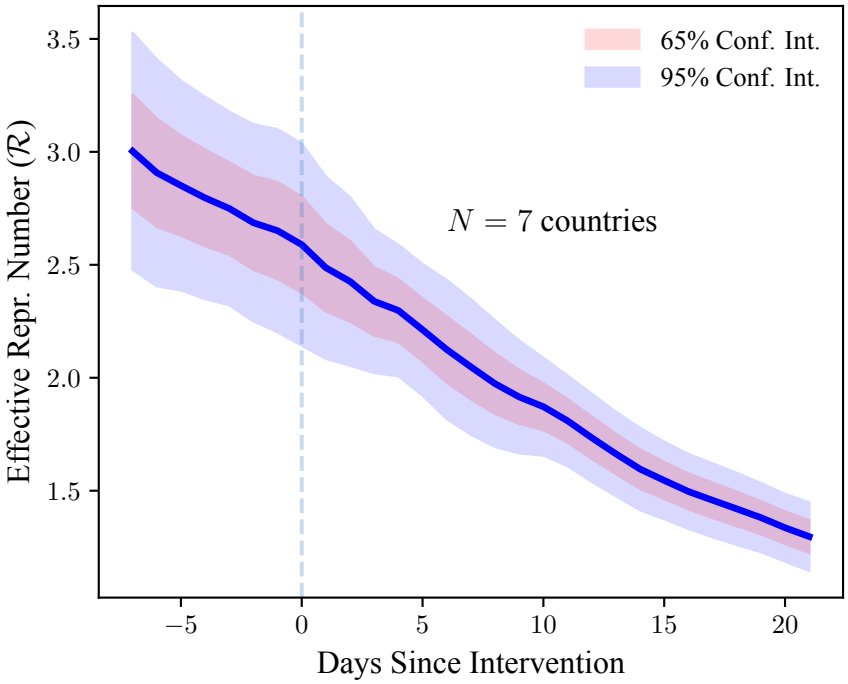

Figure A.6

$\mathcal{R}$ and Policy Interventions: Case-Based Measures

Notes: The graph plots the estimated effective reproduction number ($\mathcal{R}$) one week before and three weeks after case-based measures are introduced in a country. The original sample consists of 14 European countries studied by Flaxman et al. (2020a). For the event-study graph, we restrict the sample to countries for which data on $\mathcal{R}$ is available for the whole event window. Heteroskedasticity-robust confidence bounds are shown by the shaded areas.

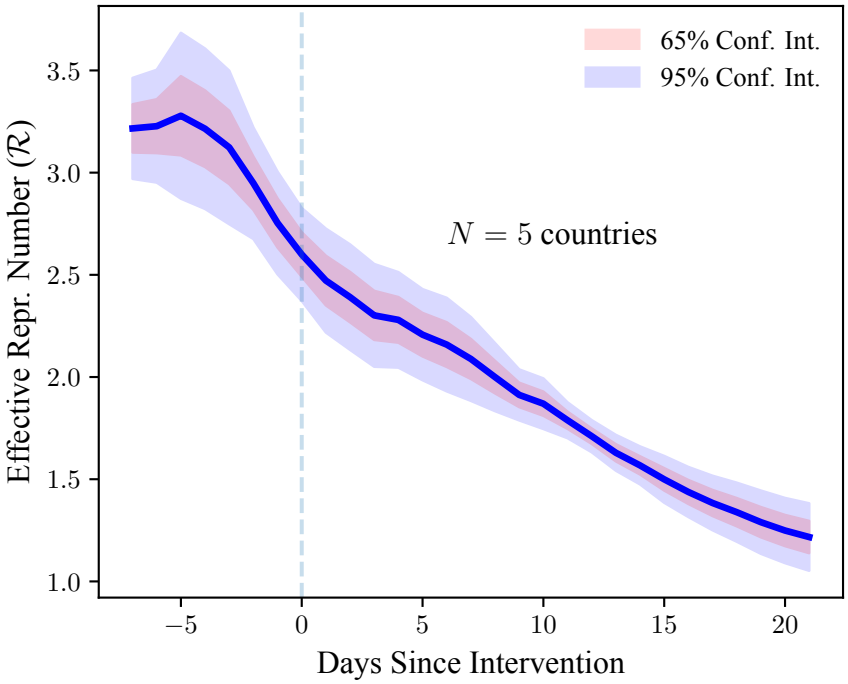

Figure A.7

$\mathcal{R}$ and Policy Interventions: School Closures

Notes: The graph plots the estimated effective reproduction number ($\mathcal{R}$) one week before and three weeks after school closures are ordered in a country. The original sample consists of 14 European countries studied by Flaxman et al. (2020a). For the event-study graph, we restrict the sample to countries for which data on $\mathcal{R}$ is available for the whole event window. Heteroskedasticity-robust confidence bounds are shown by the shaded areas.

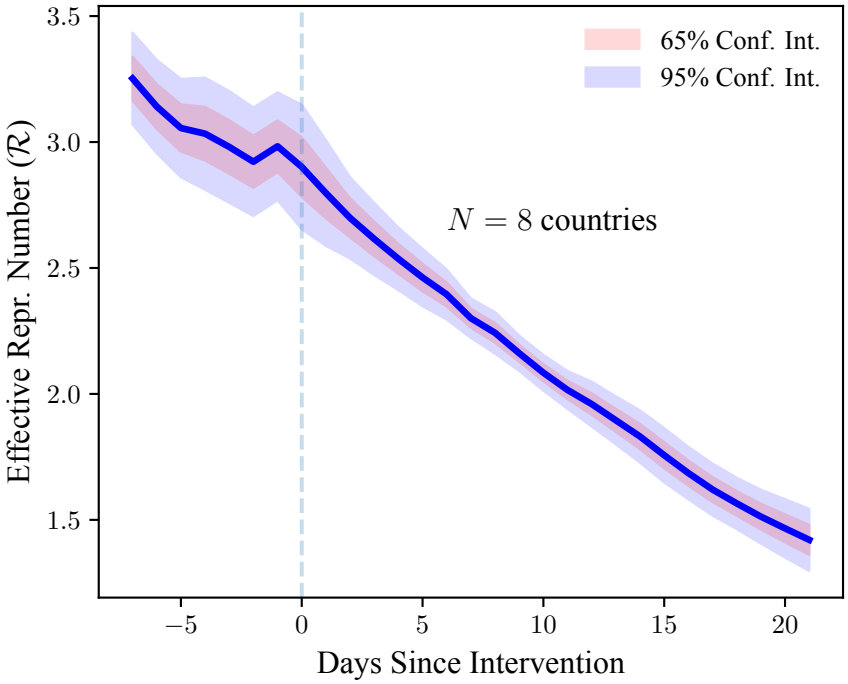

Figure A.8

$\mathcal{R}$ and Policy Interventions: Social Distancing

Notes: The graph plots the estimated effective reproduction number ($\mathcal{R}$) one week before and three weeks after social distancing is encouraged in a country. The original sample consists of 14 European countries studied by Flaxman et al. (2020a). For the event-study graph, we restrict the sample to countries for which data on $\mathcal{R}$ is available for the whole event window. Heteroskedasticity-robust confidence bounds are shown by the shaded areas.