

Did COVID-19 infections decline before UK lockdown?

Simon N. Wood

University of Bristol, UK. simon.wood@bath.edu

June 1, 2020

The number of new infections per day is a key quantity for effective epidemic management. It can be estimated by testing of random population samples. Without such direct epidemiological measurement, other approaches are required to infer whether the number of new cases is likely to be increasing or decreasing: for example, estimating the pathogen reproductive rate, R , using data gathered from the clinical response to the disease. For COVID-19 such R estimation is heavily dependent on modelling assumptions, because the available clinical case data are opportunistic observational data subject to severe temporal confounding. Given this difficulty it is useful to reconstruct the time course of infections from the least compromised available data, using minimal prior assumptions. A Bayesian inverse problem approach applied to UK data on COVID-19 deaths and the published disease duration distribution suggests that infections were in decline before UK lockdown, and that infections in Sweden started to decline only a short time later.

Clinical data on the number of cases of COVID-19 are subject to severe temporal confounding, as the rate of testing and criteria for testing have been changing rapidly on the same time scale as the infections. Because the ascertainment fraction is changing and unknown, the data can clearly not be used to infer the actual number of infections. Neither, under normal circumstances, would statisticians recommend attempting to estimate the reproductive rate of the pathogen from such data, since given the data problems the estimates must necessarily be driven primarily by the modelling assumptions. Indeed generically it is often very difficult to infer epidemiological parameters from clinical data, without the results being informed more by the prior beliefs encoded in the model than by the data (e.g. Wood et al., 2020).

The exception is when clinical data directly measure the quantity of epidemiological interest. This is the case for deaths with COVID-19 and for fatal disease duration. While not perfect, these data are far less compromised than the data on ‘cases’. Deaths are reliably recorded, clinical grounds for suspecting COVID-19 are clear, and good records are kept for fatal cases. It is of some interest to establish what these high quality data imply about the time course of infections, without strong modelling assumptions.

Two types of daily death data are available. Daily reported deaths (e.g. Worldometer, 2020) typically show marked weekly fluctuations as a result of weekly patterns in reporting delays, and may exclude deaths in some locations (such as nursing homes). Registered death data, such as the ONS data in the UK (Office for National Statistics, 2020), contain deaths in all locations and record exact date of death. The weekly cycle is less pronounced in these data, but their release is necessarily delayed relative to the daily reported deaths. Figure 1 shows data for the UK and Sweden.

Data on the incubation period from infection to onset of symptoms are analysed in Lauer et al. (2020). The period is 2 to 11 days for 95% of people, with a median of 5.2 days. Verity et al. (2020) show that the distribution of time from onset of symptoms to death for fatal cases can be modelled by a gamma density with mean 17.8 and variance 71.2 (s.d. 8.44). Note that Verity et al. (2020) correct for the bias associated with seeing short infections before long ones at the start of an epidemic. Under the assumption that the incubation period does not differ systematically between fatal and non fatal cases the infection to death interval can then be well modelled by a gamma density with mean 23 and variance 76.




Figure 1: Daily reported deaths with COVID-19 (blue) in the UK (left) and Sweden (right) since March 13th. In red is the UK ONS data for England and Wales for all locations of death by registered day of death, illustrating the lag in reported deaths. The grey regions illustrate 68 and 95% confidence regions for the underlying reported death rate from model (1). The scaled bar charts are proportional to the posterior distribution of the day of peak underlying rate according to model (1). The UK lock down started on day 11. Sweden implemented targeted measures short of lock down.

Models

Let y_i denote the deaths or reported deaths on day i . Assume that y_i follows a negative binomial distribution with mean μ_i and variance $\mu_i + \mu_i^2/\theta$. Then let

$$\log(\mu_i) = f(i) + f_w(d_i) \quad (1)$$

where f is a smooth function of time measured in days, and f_w is a zero mean cyclic smooth function of day of the week, $d_i \in \{1, 2, \dots, 7\}$, set up so that $f_w^{[k]}(0) = f_w^{[k]}(7)$, where $k = 0, 1$ or 2 denotes order of derivative. $f(t)$ represents the underlying log death rate, while f_w describes the weekly variation about that rate. The functions f and f_w can be represented using splines with associated smoothing penalties $\lambda \int f''(t)^2 dt$ and $\lambda_w \int f_w''(d)^2 dd$. Hyper-parameters λ and λ_w control the smoothness of the functions, and can be estimated as part of model fitting using a standard empirical Bayes approach (see methods). This model provides a good fit to both the reported deaths and ONS data. As expected f_w is greatly attenuated for the ONS data.

To estimate the daily infection profile the model must be extended. Consider expressing $f(i)$ in terms of the time course of earlier infections. Let $f_c(i)$ be the function describing the variation in the number of eventually fatal cases over time. Let \mathbf{B} be the square matrix such that $B_{ij} = \gamma(i - j + 1)$ if $i \geq j$ and 0 otherwise, where γ denotes the onset-to-death gamma density (mean 23, variance 76) given above. If $\mathbf{f}_c = [f_c(0), f_c(1), \dots]^T$ and $\boldsymbol{\delta} = [\delta(1), \delta(2), \dots]^T$ then $\boldsymbol{\delta} = \mathbf{B}\mathbf{f}_c$, where $\delta(i)$ is the expected number of deaths on day i . $\log f_c(i)$ can be represented using a spline basis, again with a cubic spline penalty. The final model is then obtained by simply substituting $f(i) = \log \delta(i)$ into (1). \mathbf{B} is rank deficient, so inferring f_c can be viewed as an inverse problem: without regularization multiple solutions that oscillate from day-to-day are possible. This ambiguity is removed by the smoothing penalty on $\log f_c$.

Inference about f_c was conducted in a fully Bayesian manner using Markov Chain Monte Carlo methods, exploiting the fact that smoothing penalties can be induced by the adoption of appropriate Gaussian smoothing priors. It is necessary to infer f_c over a considerable period before the first death occurs. 40 days is clearly sufficient given the form of γ . In fact it makes sense to reduce this interval, after inspecting a pilot run, to avoid a lengthy initial period of zero fatal cases, consequent lack of identifiability of $\log f_c$ and poor MCMC mixing. On this basis a 20 day initial period is more than

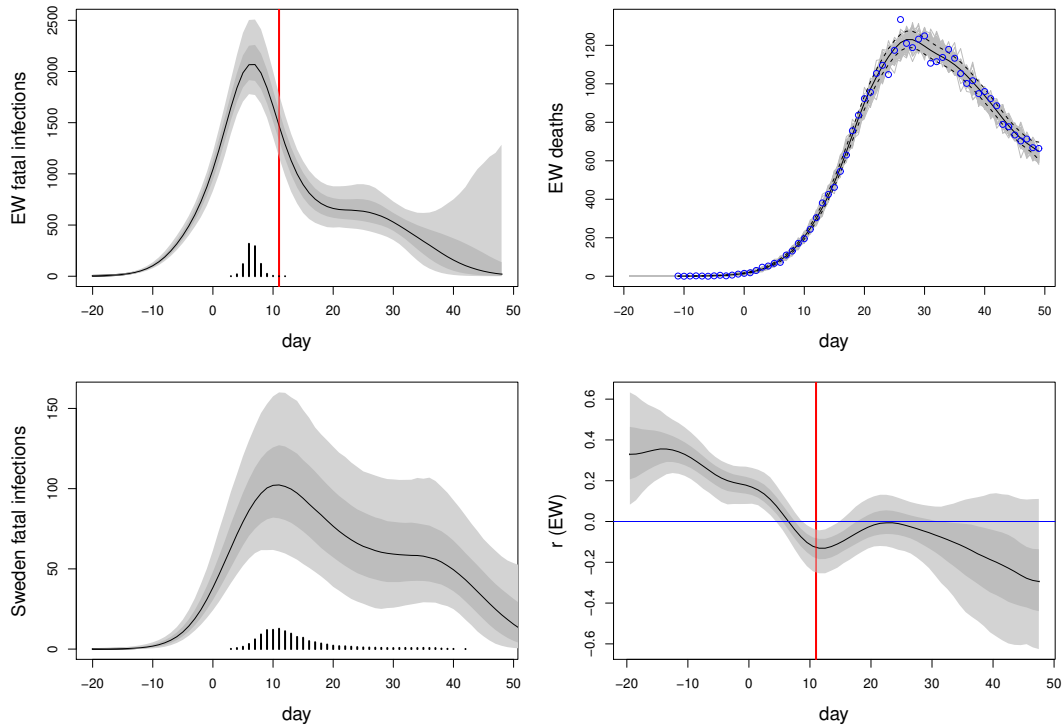


Figure 2: Top left: Inferred daily fatal infection rate, f_c , for England and Wales. Light grey and dark grey regions show 95% and 68% confidence regions, respectively. The black curve is the posterior median profile. Day 0 is 13th March 2020, and the vertical red line marks the first day of UK lockdown. The scaled black barchart shows the posterior distribution for day of peak infection. Top right: Consistency check. In grey are 100 sets of death data simulated forward from the inferred median fatal infection profile. Blue symbols are the ONS daily death data for England and Wales on which inference is based. The dashed curves are 95% confidence intervals for underlying death rate estimated by direct fitting of (1). Bottom left: The inferred fatal infection profile for Sweden, based on reported daily deaths as shown in figure 1, plotted in the same way as for England and Wales. Bottom right: Inferred instantaneous intrinsic growth rate of infections, r , for England and Wales, with confidence regions as in the other plots. The blue line is at $r = 0$, the boundary between increase and decline in the daily infections.

sufficient. For stable inference it also makes sense to explicitly include in the death data the fact that no deaths were observed in this initial period.

Results

Figure 2 shows the results of applying the model to the Office for National Statistics daily COVID-19 death data for England and Wales. The data were truncated at May 1st, since comparison of the data released on 19th of May with that released the previous week suggested that very few more registrations are now likely for the period before May 1st. Including further data risks exaggerating the rate of decline in the death rate at the data's end. The most notable feature of the results is that fatal infections are inferred to be in substantial decline before lockdown. Sweden appears most likely to have peaked only a few days later. The profile for Sweden is considerably less certain, as it is based on the reported death series, which is noisy even allowing for the weekly reporting cycle. It is also likely that the Swedish peak is estimated a day or two too late as a result of reporting delays.

Taken together the results for England and Wales and for Sweden are strongly suggestive that full lockdown may not have been necessary to avoid health service overload, and more limited measures might have been effective. This sharply emphasises the desirability of statistically well founded direct measurement of epidemic size through randomized testing. Had such testing being carried out leading up to lockdown it would have been clear if the measures preceding lockdown were working, or whether stronger restrictions were needed. Instead management was reliant on a complex modelling synthesis of expert judgement and highly problematic clinical case data. Less statistically problematic reconstructions, like the one presented here, are clearly only possible weeks after the fact.

Figure 2 also plots the instantaneous intrinsic growth rate of daily infections, r , (the time derivative of $\log f_c(i)$). Daily infections increase for $r > 0$. Over-interpretation of this quantity should be avoided: conceptually it relates to a single well mixed population: the population was in fact stratified at lockdown.

Model checking

While standard residual checks indicate no problem with the model from the point of view of statistical fit, there are two issues which could potentially undermine the conclusions.

The first is that the infection to death interval distribution could be systematically wrong. The results are based on the Verity et al. (2020) point estimates for the fatal disease duration mean and variance, but shifting the mean infection time down 0.9 to the lower end of their 95% CI simply moves the peak closer to lockdown by much the same amount, as does reducing it by another 0.9. Against this must be set the fact that the (Lauer et al., 2020) incubation period used here is some half a day less than is suggested by a more recent meta-analysis (McAloon et al., 2020), which would tend to push the peak back. Alternatively, it could be that the onset to death distribution from Verity et al. (2020) gives too long durations for the UK context. For example, this could happen if age strongly effects time from onset-to-death and the age structure of the UK cases differs substantially from the age structure of the cases used by Verity et al. (2020). Similarly within hospital transmission to already very ill patients might lead to shorter disease duration, and in turn to the peak being estimated as earlier than it was, but the cross infection rate would have to be quite high for this to be a substantial factor. Another possible problem could arise if fatal cases have substantially shorter incubation periods than non-fatal cases, but no such result appears to have been reported.

The second issue is whether the smoothing penalty on $\log f_c$ would lead to systematic mis-timing of the estimated peak under the scenario of a very asymmetric peak in the true infection profile around lockdown. To investigate this, data were simulated from a model in which the underlying infection rate increased geometrically, doubling every 3 days until lockdown, when the rate dropped immediately to 0.2 of its peak value, shrinking thereafter by 5% per day. Fatal infections were simulated as Poisson deviates with the given underlying rate. This model is an extreme scenario, in which pre-lockdown measures had no effect, and the effect of lockdown was instant, as if the locked down population (i.e. those not in essential work) had isolated alone, rather than increasing their contact with members of their household while drastically reducing it with everyone else. Under this scenario, the method does indeed tend to incorrectly estimate the infection peak as 2 to 3 days before lockdown, rather than the day before, as it struggles to accommodate the drop.

The naive approach to this issue is to introduce a parameter at lockdown representing an instantaneous drop in infections. However doing so introduces a very strong structural assumption into the model, undermining the aim of avoiding strong assumptions. This approach also has the serious side effect of introducing non-parametric smoothing boundary effects on both sides of the break. These boundary effects severely compromise inference in the most interesting region of the infection profile, while simultaneously increasing the importance of the structural assumption at the expense of the data. Indeed when such a model is built it estimates a large drop even from data simulated from a smooth

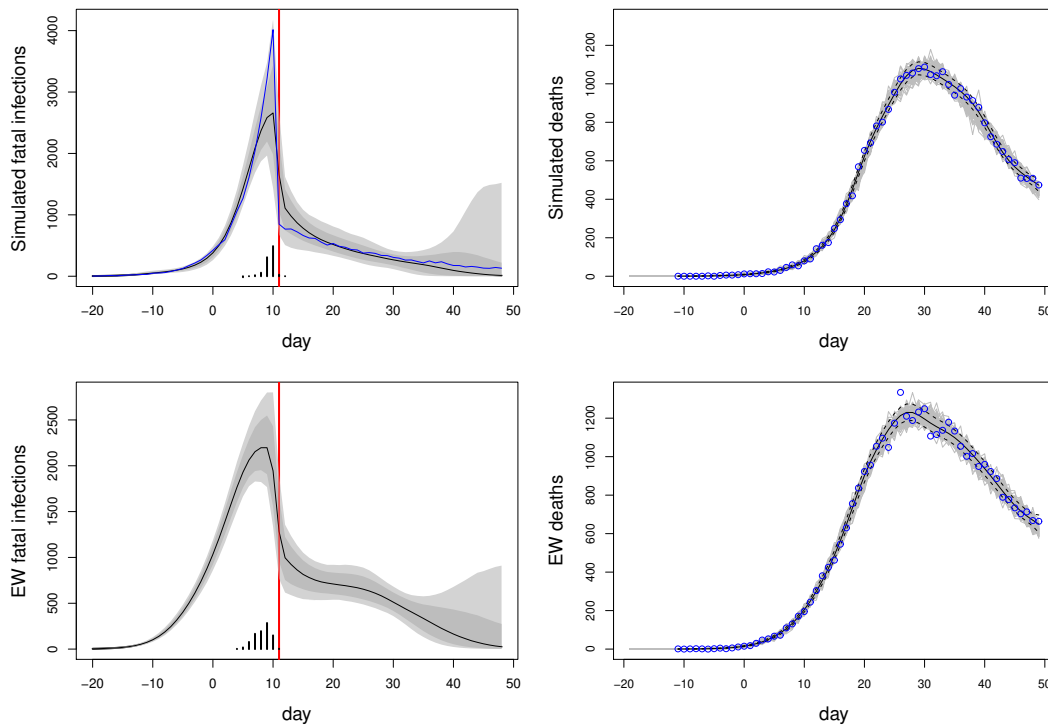


Figure 3: Model checking plots in which the smoothness assumptions are relaxed around lockdown by a time dilation, in order to allow accurate capture of any extremely discontinuous infection profile in this region. The top row shows the method reconstructing an extreme simulation scenario in which there was no reduction in transmission rate up until lockdown, and then an instantaneous drop. Left: the reconstruction (plot meaning as figure 2) with the true simulated daily infections shown in blue. Right: forward simulation from the median profile as in figure 2. The blue symbols are the simulated death data used for inference. The bottom row is for the ONS data under the time dilated model. Even this model deliberately modified to promote a very abrupt change at lockdown suggests that the infection rate was probably declining before lockdown.

infection profile. It also estimates such a drop if we move the drop's location.

A better approach is to use a smooth time-dilation to relax, but not eliminate, the model smoothness assumptions in the vicinity of lockdown. The dilation is made sufficient that the model can accurately capture the extreme scenario in the simulation, but without imposing a break and boundary effects. In particular f_c and its smoothing penalty are computed with respect to a version of time which makes the day before, of and after lockdown count as 3, 5 and 3 days, respectively. Obviously regular un-dilated time is used for the mapping infections to deaths. For the extreme simulation, the model then gives over 50% posterior probability to the day before lockdown as the peak. In contrast the same model for the real data has 13% probability of the peak being the day before lockdown, with the remainder earlier. The posterior expectation is then for a peak 4 days before lockdown.

Figure 3 shows the results from fitting the time dilated model to the extreme simulation scenario and to the England and Wales ONS data. Even this model, deliberately modified to favour a very abrupt change at lockdown, suggests that infections started to decline before lockdown.

Discussion

This paper does not prove that the peak in fatal infections in England and Wales preceded lockdown by several days. Indeed the failure to undertake the sampling that could have gathered data to directly measure infections early in the epidemic means that it will never be possible to be certain about timings, given the severe biases in clinical data other than deaths and fatal disease duration. What the results show is that, in the absence of strong assumptions, the currently most reliable data strongly suggest that the decline in infections in England and Wales began before lockdown. Furthermore, such a scenario would be consistent with the infection profile in Sweden, which began its decline in fatal infections shortly after the UK, but did so on the basis of measures well short of lockdown.

These facts have implications for the policies to be adopted in the coming autumn, particularly given the peculiar ethical issues associated with lockdown. For example, plausible estimates of the life loss burden from an unmitigated COVID-19 epidemic in the UK are about 2 weeks per person¹. A plausible lower bound on the UK life loss from the 2008 financial crisis and its aftermath is 7 weeks per person². The economic shock from lockdown is substantially larger than 2008. Similarly the implied willingness to pay to save a life year from COVID-19 appears to be an order of magnitude higher than the usual UK National Institute for Health and Care Excellence threshold used for any other disease.

Methods

Direct inference about (1) uses the empirical Bayes approach of Wood et al. (2016) in which the smooth functions are estimated by penalized likelihood maximisation (e.g. Green and Silverman, 1994), with the smoothing parameters and θ estimated by Laplace approximate marginal likelihood maximization. Writing β for the combined vector of basis coefficients for f and f_w , the penalized version of the log likelihood, $l(\beta)$, can be written

$$l(\beta) - \frac{\lambda_f}{2} \int f^{[2]}(t)^2 dt - \frac{\lambda_w}{2} \int f_w^{[2]}(d)^2 dd = l(\beta) - \frac{1}{2} \beta^\top \mathbf{S}_\lambda \beta$$

where $\mathbf{S}_\lambda = \lambda_f \mathbf{S}_f + \lambda_w \mathbf{S}_w$: \mathbf{S}_f and \mathbf{S}_w are known constant positive semi-definite matrices. Smoothing parameters, λ_f and λ_w , control the smoothness of f and f_w . Let $\hat{\beta}$ be the maximizer of the penalized log likelihood, and \mathbf{H} its negative Hessian at $\hat{\beta}$. Viewing the penalty as being induced by an improper Gaussian prior, $\beta \sim N(\mathbf{0}, \mathbf{S}_\lambda^-)$, $\hat{\beta}$ is also the MAP estimate of β . Furthermore in the large sample limit

$$\beta | \mathbf{y} \sim N(\hat{\beta}, (\mathbf{H} + \mathbf{S}_\lambda)^{-1}). \quad (2)$$

Writing the density in (2) as π_g , and the joint density of \mathbf{y} and β as $\pi(\mathbf{y}, \beta)$, the Laplace approximation to the marginal likelihood for the smoothing parameters λ and θ is $\pi(\lambda, \theta) = \pi(\mathbf{y}, \beta) / \pi_g(\beta | \mathbf{y})$. Nested Newton iterations are used to find the values of $\log(\lambda), \theta$ maximizing $\pi(\lambda, \theta)$ and the corresponding $\hat{\beta}$ (for details see Wood et al., 2016).

Given (2) credible intervals for f are readily computed, but it is also straightforward to make inferences about when the peak in f occurs. Simply simulate replicate coefficient vectors from (2) and find the day of occurrence of the peak for each corresponding underlying death rate function, f .

In principle the model formulated in terms of f_c could be estimated using the framework of Wood et al. (2016), but some non-trivial computational work would be required to set it up. Instead a fully

¹Based on Office for National Statistics (2019) lifetables, the age specific infection fatality ratios of Wood et al. (2020), a herd immunity fraction of 80% and a lower bound adjustment for co-morbidities based on Hanlon et al. (2020).

²The life expectancy gap between those in the upper and lower half of the UK income scale grew by 14 weeks in the aftermath of 2008, a loss of life that is difficult to attribute to confounders. See Marmot et al. (2020) especially figure 2.5.

Bayesian approach was taken and the model implemented using the JAGS software for Gibbs sampling (Plummer, 2003; Plummer et al., 2006), making use of the automatic code template generation described in Wood (2016) for reliable implementation of spline smoothers in JAGS. 5×10^6 samples were generated, retaining every 500th sample. This was sufficient to ensure effective sample sizes in the hundreds for even the slowest mixing parameters, while most parameters had effective sample sizes close to 10000. Trace plots suggested rapid convergence.

References

- Green, P. J. and B. W. Silverman (1994). *Nonparametric Regression and Generalized Linear Models*. Chapman & Hall.
- Hanlon, P., F. Chadwick, A. Shah, R. Wood, J. Minton, G. McCartney, C. Fischbacher, F. S. Mair, D. Husmeier, and J. Matthiopoulos (2020). COVID-19 exploring the implications of long-term condition type and extent of multimorbidity on years of life lost: a modelling study. <https://wellcomeopenresearch.org/articles/5-75>.
- Lauer, S. A., K. H. Grantz, Q. Bi, F. K. Jones, Q. Zheng, H. R. Meredith, A. S. Azman, N. G. Reich, and J. Lessler (2020). The incubation period of coronavirus disease 2019 (covid-19) from publicly reported confirmed cases: estimation and application. *Annals of internal medicine*.
- Marmot, M., J. Allen, T. Boyce, P. Goldblatt, and J. Morrison (2020). *Health Equity in England: The Marmot Review 10 Years On*. The Health Foundation.
- McAloon, C. G., A. Collins, K. Hunt, A. Barber, A. Byrne, F. Butler, M. Casey, J. M. Griffin, E. Lane, D. McEvoy, P. Wall, M. J. Green, L. O’Grady, and S. J. More (2020). The incubation period of covid-19: A rapid systematic review and meta-analysis of observational research. *medRxiv*.
- Office for National Statistics (2019). National life tables, UK: 2016 to 2018. <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/lifeexpectancies/bulletins/nationallifetablesunitedkingdom/2016to2018>.
- Office for National Statistics (2020, May). Deaths registered weekly in England and Wales. <https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/bulletins/deathsregisteredweeklyinenglandandwalesprovisional/latest>.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003)*, pp. 20–22.
- Plummer, M., N. Best, K. Cowles, and K. Vines (2006). coda: Convergence diagnosis and output analysis for MCMC. *R News* 6(1), 7–11.
- Verity, R., L. C. Okell, I. Dorigatti, P. Winskill, C. Whittaker, N. Imai, G. Cuomo-Dannenburg, H. Thompson, P. G. Walker, H. Fu, et al. (2020). Estimates of the severity of coronavirus disease 2019: a model-based analysis. *The Lancet Infectious Diseases*.
- Wood, S. N. (2016). Just another Gibbs additive modeller: Interfacing JAGS and mgcv. *Journal of Statistical Software* 75(7).
- Wood, S. N., N. Pya, and B. Säfken (2016). Smoothing parameter and model selection for general smooth models (with discussion). *Journal of the American Statistical Association* 111, 1548–1575.
- Wood, S. N., E. C. Wit, M. Fasiolo, and P. J. Green (2020). COVID-19 and the difficulty of inferring epidemiological parameters from clinical data. *The Lancet Infectious Diseases in press*.
- Worldometer (2020). <https://www.worldometers.info/coronavirus/>.