Received: 00 Month 0000 | Revised: 00 Month 0000 | Accepted: 00 Month 0000
DOI: xxx/xxxx

Forecasting subnational COVID-19 mortality using a day-of-the-week adjusted Bayesian hierarchical model

Justin J. Slater*[1,2] | Patrick E. Brown[1,2] | Jeffrey S. Rosenthal[2]

[1]Centre for Global Health Research, St. Michael's Hospital
[2]Department of Statistical Sciences, University of Toronto

Correspondence
Email: justin.slater@mail.utoronto.ca

Abstract

As of October 2020, the death toll from the COVID-19 pandemic has risen over 1.1 million deaths worldwide. Reliable estimates of mortality due to COVID-19 are important to guide intervention strategies such as lockdowns and social distancing measures. In this paper, we develop a data-driven model that accurately and consistently estimates COVID-19 mortality at the regional level early in the epidemic, using only daily mortality counts as the input. We use a Bayesian hierarchical skew-normal model with day-of-the-week parameters to provide accurate projections of COVID-19 mortality. We validate our projections by comparing our model to the projections made by the Intitute for Health Metrics and Evaluation, and highlight the importance of hierarchicalization and day-of-the-week effect estimation.

KEYWORDS:
COVID-19, Forecasting, Bayesian inference, Hierarchical Models

1 | INTRODUCTION

As of October 2020, the death toll from the COVID-19 pandemic has to risen over 1.1 million deaths worldwide, with deaths in many regions now rising again following decreases in late spring and early summer. Although this number is likely an under estimate of the true number of deaths, it is more reliable than the reported number of cases, which is largely a function of the number of patients tested. Reliable estimates of mortality due to COVID-19 are useful for guiding intervention strategies such as lockdowns and social distancing measures. Estimates are needed at the regional (e.g provincial or state) level, as the spread of the disease can vary greatly within a particular country.

There have been many attempts at forecasting COVID-19 cases and mortality. Extensions of Susceptible, Exposed, Infectious, or Recovered (SEIR) models have been considered (Anastassopoulou, Russo, Tsakris, and Siettos 2020; Sarkar, Khajanchi, and Nieto 2020). Various time series models have also been considered (Chakraborty and Ghosh 2020; Perc, Gorišek Miksić, Slavinec, and Stožer 2020; Petropoulos and Makridakis 2020). Perhaps most notably, the Institute for Health Metrics and Evaluation (IHME) has made their predictions available since March 25th 2020 (Friedman, Liu, and Gakidou 2020), and have been cited as the gold standard regional level projections. However, in all of these forecasting methods, there has been little attempt at accounting for differences in deaths by day-of-the-week, which if left unaccounted for, can drastically bias long-term forecasts depending on which day of the week the observed data ends. Additionally, making projections for regions can be difficult where there are a relatively small number of deaths. COVID-19 forecasting methodology needs to be able to handle low daily mortality counts, while providing reasonable mortality estimates for each region.

Figure 1 shows the COVID-19 daily death counts in the United States from March 2nd to June 25th. There are several key features of these daily deaths that seem to be prevalent in every country or region's Coronavirus mortality counts. Firstly, note the rapid rise in daily deaths relative to the decline. Capturing this skewness in the daily death counts is essential for accurately forecasting COVID-19 mortality, and is not captured using the Normal density initially used by the IHME (note that the IHME has since switched to an SEIR model). Additionally, notice the weekly periodicity in daily death counts. It appears that certain days of the week tend to have higher death counts than others, which is an important feature to capture, so that our cumulative death forecasts don't depend on what day of the week they are made. This weekly periodicity has been confirmed using spectral analysis (Bukhari, Jameel, Massaro, D'Agostino, and Khan 2020), but the relative risks of mortality between certain days-of-the-week are still indeterminate. Studying day-of-the-week effects has proven useful in other fields such as actuarial science (Crevecoeur, Antonio, and Verbelen 2019) and economics (Berument and Kiymaz 2001).

The goal of this paper is to develop a data-driven model that accurately and reliably estimates COVID-19 mortality at the regional level early in the epidemic, using daily mortality counts as the input. We do so by developing a hierarchical Bayesian model where the daily death counts are assumed to follow a skew-normal density function, and vary by day-of-the-week. In doing so, we estimate the number of daily deaths at the peak of the epidemic, the date of the peak of the epidemic, and other epidemiologically significant features. The hierarchical nature of our model will allow for accurate estimation of cumulative death counts in regions where the epidemic is still in the early stages by borrowing information from regions where the epidemic has matured. We compare the forecasting performance of our model to the projections made by the Institute for Health Metrics and Evaluation (IHME), as well as the observed mortality values and highlight the importance of hierarchicalization and accounting for day-of-the-week effects when projecting mortality counts during an epidemic.

This work is an extension of the model presented in Brown, Jha, et al. 2020, which has shown promise as a national-level forecasting model. The two main contributions of our work are to allow for subnational level data (hierarchicalization), and estimation/modelling of day-of-the-week effects.

2 | METHODS

2.1 | The Skew-Normal Model

We saw in Figure 1 that a key feature of daily coronavirus mortality counts is the rapid rise relative to the fall in daily death counts. A natural choice for the response distribution of the deaths per day is the negative binomial distribution with the trend in deaths per day in a region following a skew normal curve. The negative binomial is often used in infectious disease modelling, where events are positively correlated, causing larger variances than if the events were independent (Lloyd-Smith, Schreiber, Kopp, & Getz 2005). The skew-normal density provides a good base to model the trend in death counts, but is far too simple to capture the full range of shapes of epidemics by itself. Firstly, although a majority of the deaths in a region occur during the main epidemic, a small number of deaths can occur outside of this epidemic. Additionally, in order to be able to compare various regions' mortality counts, we need to "standardize" our estimates based on how many deaths we would expect to see in that region from all causes. Lastly, we need to add a multiplicative term to our skew-normal that accounts for differences in daily mortality counts by day-of-the-week. By including all of these considerations, we arrive at the full model.

The statistical model (referred to as the "DOW model") used to estimate daily mortality, Y_{ij}, in region i of country j is given by:

$$
\begin{aligned}
Y_{ij} &\sim \text{NegBinom}[\lambda_{ij}(t), \tau_j] \\
\lambda_{ij}(t) &= R_{j,m[t]} E_{ij} [C_{ij} f(t; A_{ij}, B_{ij}, K_{ij}) + D_{ij}] \\
C_{ij} &\sim \text{Gamma}(\alpha_j/\theta_C, \theta_C) \\
B_{ij} &\sim \text{Gamma}(\eta_j/\theta_B, \theta_B) \\
K_{ij} &\sim \text{Gamma}(\zeta_j/\theta_K, \theta_K)
\end{aligned}
\tag{1}
$$

Prior distributions for model parameters can be found in Table 1 . The function f is a skew-Normal density function with three parameters: the "location" parameter, A_{ij}, indicate the date at which the daily deaths reaches its peak and is analogous to the mean of a normal distribution; B_{ij} represents the duration of the epidemic, and is analogous to the standard deviation in the Normal density; and K_{ij} is the "skewness" parameter, which describes the ratio of the initial incline relative to the decline. The skewness is a key parameter for capturing the shape of daily mortality trends. The E_{ij}'s are the age standardized death counts in each region of the included countries. This was computed by obtaining age distribution information from census data of each country and comparing it to the deaths-by-age breakdown in Italy on March 29 2020. Note that E_{ij} is a constant, as it is calculated apriori for each region. Inclusion of the E_{ij} allows for comparable heights of peaks between regions, which is captured in the parameter C_{ij}. A high C_{ij} means that there were a large number of coronavirus related deaths, relative to the expected number of deaths in that region. The parameter D_{ij}, known as the "spark" term, captures the few deaths that were outside of the main epidemic. $R_{j,m[t]}$ captures the day-of-the-week effect for country j on day-of-the-week $m[t]$, with $R_{j,Sunday}$ fixed at 1. The overdispersion parameter, τ_j, allows the variance of the daily mean deaths to vary by country by a multiplicative factor.

The number of parameters in this model can grow very quickly depending on the number of countries/regions included in analysis, and can be difficult to implement without carefully chosen priors. One of the advantages to Bayesian analysis is the incorporation of prior knowledge to guide parameter estimation. That is, we can use information about duration and severity from epidemics that are already over (e.g Spain) to set priors for regions where the epidemic has yet to peak (e.g Brazil).

Note that C_{ij}, B_{ij} and K_{ij} are all modelled hierarchically. The advantage of this is that for regions low death counts can "borrow" information from other regions in the same country to estimate the severity, duration, and skewness of the epidemic. For example, the estimate of C_{ij} will be a weighted average between the country's mean and the region's mean, but will tend more toward the country's mean when the number of events is small (Gelman & Hill 2006). Modelling A_{ij} hierarchically was considered, however it was deemed inappropriate because in regions with small death counts, the location parameter would tend toward the country average. This is problematic because the reason that the death counts are low in that region is likely because the epidemic has yet to run its course. For this reason, we decided to estimate the location parameter separately for each region. Modelling the day-of-the-week effects, $R_{j.m[t]}$, hierarchically was also considered because it would provide a day-of-the-week effect estimate for each region. However, this was deemed computationally too cumbersome, as this would add over 300 parameters to our model.

2.2 | Daily mortality projections for four countries

We applied our model to daily COVID-19 death counts from 95 regions from four countries: U.S states; Canadian provinces; Spanish Autonomous Communities; and Brazilian states. These four countries were chosen based on demographic data availability at the regional level. Data are from `www.coronavirus.app` (Scriby, Inc. 2020), where any region with 50 or more deaths as of June 25th was included in the analysis. Parameters for our models were estimated using No-U-turn sampling (Hoffman & Gelman 2014) within the Stan software (Carpenter et al. 2017). Four chains were used with 3000 iterations of warm-up and 1000 iterations of sampling, which were then thinned by a factor of 10 (leaving 400 posterior samples for each parameter). Convergence of Markov Chains was assessed using trace plots alongside the Gelman-Rubin Statistic (Rhat < 1.05) (Gelman, Rubin, et al. 1992).

Forecasts were created from the posterior samples of $\lambda_{ij}(t)$ up until October 1st 2020. Given that λ_{ij} has a day-of-the-week effect, the posterior samples of $\lambda_{ij}(t)$ will be oscillatory. Forecasts were also made at the country level by computing

$$\bar{\lambda}_j(t) = \sum_i \lambda_{ij}(t)$$

for each posterior sample.

2.3 | Day-of-the-week effect estimates

The model produces estimates for 24 day of the week effect parameters: one for each day of the week (except for Sunday which is fixed at 1) for each of the four countries. In order to gain some insight as to whether the estimated day-of-the-week effects are simply due to differences in reporting, we pulled proportions of historical deaths by day-of-the-week in Canada from various non-COVID related causes: circulatory, pulmonary, circulatory and pulmonary, and non-circulatory pulmonary. We will compare the estimated mortality rates from these causes to the estimated day-of-the-week effects to see if there is a similar trend. If not, then this suggests that coronavirus may be more likely to cause death on certain days of the week, or simply that coronavirus mortality has its own unique reporting artifacts.

2.4 | Validating Forecasts

Institute for Health Metrics and Evaluation (2020) has made projections for the United States available since March 25th 2020, and has since expanded the number of regions they include in their model. In order to validate our model projections, we ran our model using the same daily mortality data as the IHME, and compared both model's cumulative death counts to the observed death count on June 30th 2020 for 14 different time points ranging from April 1st to June 25th. Additionally, we ran our model without a day-of-the-week effect to see whether the forecasting performance of our model relies on the day-of-the-week, and is not only outperforming the IHME model for other reasons. This model will be referred to as the "non-DOW model".

Models were compared by assessing consistency and accuracy of projections throughout time. Consistency was assessed by examining the amount of overlap between successive intervals, with subsequent intervals hopefully being narrower and mostly contained in previous intervals. Successive interval overlap is important to ensure that a model's results are consistent throughout time. If two successive intervals do not overlap, then at least one of those intervals must not contain the true value. Overlap was measured at 13 time points, since the first time point does not have a previous interval. At each of these 13 time points, the proportion of the interval that is contained in the previous interval is calculated for each region individually, resulting in 13 proportions per region. These proportions are then averaged across all regions to determine which model, on average, had the most consistent predictions.

The first measure used to assess the accuracy of the models was the standardized Root Mean Squared Log Error (sRMSLE) at time t, which was computed as:

$$sRMSLE(t) = \sqrt{\operatorname*{mean}_{ij}\left\{ \frac{\log(P_{ij}(t)/O_{ij}(t))^2}{\log(O_{ij}(t))} \right\}}$$

where $P_{ij}(t)$ and $O_{ij}(t)$ are the predicted/observed number of deaths from region i of country j. Note that each time point has a different number of regions with available data. The mean squared error is often used as a metric to assess accuracy of point estimates. We chose to use the log error because of the skewness in the reported deaths. Additionally, we chose to standardize the result to ensure all included regions have the same weight.

Accuracy of model forecasts was also assessed by computing the proportion of time points that the prediction intervals contain the observed cumulative death counts on June 30th 2020. To avoid the issue of excessively wide intervals appearing the best, we also plot the mean log-length of the intervals for each model, at each time point. The model which contains the true value the most often, relative to the mean log-length of the intervals, was considered the favourable model by this metric.

3 | RESULTS

3.1 | Forecasts of Daily Mortality in Four Countries

Figure 2 visualizes the variation in the population sizes of the regions considered. Note that the regions in Spain tend to be smaller than those of the other three countries and for Canada only the four largest provinces were included. Figures 3 -5 show forecasts for countries and a subset of the regions studied, a full set of results is available in the online supplement to this paper. The forecasts for all of Brazil are shown in Figure 3 a, with the red points representing the data that was used to fit the model, and the purple representing the observed values from June 26th to August 31st. Up until July, our model fits the data quite well, indicated by the red points clustered around the posterior samples, and the day-of-the-week effect is well captured. However, starting in July, our model slightly underestimates the mean daily death counts. Figures 3 b-c show the daily mortality forecasts in Sao Paulo, the region in Brazil with the most deaths, and Acre, a region with few deaths. In Acre, our model captures the trend of daily deaths reasonably well, suggesting that the hierarchical nature of our model is helping provide good forecasts for small regions. Starting in July, we are underestimating daily deaths in Sao Paulo. Although we estimate the mean daily deaths in Sao Paulo before July, it appears that we are slightly underestimating the day-of-the-week effect in Sao Paulo, indicated by the dispersion of the red and purple points having higher variance than the day-of-the-week effect allows for. This is due to the fact that we only allowed for one day-of-the-week effect for each country.

The forecasts for all of the United States are shown in Figure 4 a. Notice that our model under predicts daily mortality counts throughout July and August. This is likely due to the fact that many states in the U.S have loosened their lockdown restrictions, which our model is unable to account for. Figures 4 b-c show plots of the projections for Illinois and California. Our projections for Illinois seem to be quite accurate, as this is one of the states that is yet to experience any repercussions from reopening. However, our model is underestimating deaths in California likely because of loosening COVID-19 related restrictions.

Results for Spain and Canada are presented in Figure 5 , and are less interesting due to the fact that the epidemics are largely over in these countries. When looking at the daily death plot for Spain, we see a small second wave not captured by the model. But the fact that this plot is on the log scale amplifies the apparent size of the second wave. Our model seems to project

Canada's COVID-19 mortality reasonable well, likely due to the relatively firm COVID-19 restrictions in Canada. Ultimately, our model seems to predict daily COVID-19 mortality well in regions with firm COVID-19 restrictions. In 3.3, we will validate our projections under the assumption that COVID-19 related restrictions are held constant.

3.2 | Day-of-the-week effects

The 2.5th, 50th and 97.5th percentiles of the posterior distributions of the day-of-the-week effects are presented in Figure 6 . With the possible exception of Spain, Sunday appears to report the lowest death counts, followed by Monday. In all countries, death counts seem to rise on Tuesday, and remain high until Friday or Saturday, and are still elevated relative to Sunday. The day-of-the-week effect is most pronounced in the United States, where Tuesday - Friday are all very similar, but are vastly different than the other days of the week. Brazil also shows a strong day-of-the-week effect, indicated by Monday's credible interval having almost no overlap to any other day-of-the-week. As expected, Canada's credible intervals are the widest, due to the fewest deaths overall.

The proportion of deaths by day-of-the-week relative to Sunday for historical (non-COVID) data from Canada are shown in Table 2 . Note that there is very little, if any, similarities between our model's day of the week estimates and this data. These numbers only fluctuate by a few percentage points, and do not generally show a large spike on Tuesday as we saw in our COVID-19 day-of-the-week estimates. This could indicate that people are more likely to die due to COVID-19 on particular days, but is more likely explainable by differences in reporting between data sources.

3.3 | Model Validation

Projections made at each of the 14 dates in all regions for the 3 models are shown in the online supplement. Figures 7 -9 show projections for the four regions with the most deaths in each nation. In any plot where the Institute for Health Metrics and Evaluation (2020) results are missing, it is because they were not produced for that region at that time. In the plots where the DOW or non-DOW model were missing, it is because they had not yet achieved the minimum 50 deaths required to be included in our analysis. For regions with a large number of deaths (such as New York), the day-of-the-week model was very consistent, where 95% credible intervals have a large amount of overlap from date-to-date. The IHME projections are somewhat inconsistent for this region, indicated by non-overlapping intervals. However, in regions like Florida, the IHME model seems to outperform the DOW and non-DOW models, requiring a more formal investigation into model projection assessments. The mean proportion of overlap between successive intervals for all regions is shown in Figure 10 b. The DOW model tends to show the highest mean overlap at 8/13 time points, the non-DOW model had the most overlap at 3/13 time points, follow by the IHME which had the most overlap at 2/13 time points. This suggests that the DOW model produces the most consistent projections of the 3 models, and indicates that in the long-run our projections are likely to be better suited than the IHME's.

The sRMSLE for each of the three models at each time point is shown in Figure 10 a. This figure shows that the IHME is drastically underperforming when compared to the other two models between April 7th and May 4th. The sRMSLE's are otherwise comparable. Figure 10 c shows the proportion of regions that each model's interval contained the observed June 30th cumulative death count. Somewhat surprisingly, we see a downward trend as we get closer to June 30th, as shorter term predictions should become easier. This is likely because many regions had started their second uprising in daily deaths just prior to June 30th, causing models to slightly underestimate deaths in the short term. The DOW and non-DOW models seem to do very well early on in the epidemic, capturing over 90% of true values on April 17th. The IHME model tends to do better in the 1-month projection range. All models tend to perform poorly for very short-term projections (i.e <2 weeks).

Although the IHME model appears to be better suited in 1-month projections based on our accuracy metric, Figure 10 d shows that this is likely due to the increased interval width. The IHME's intervals are approximately two times as wide when making 1-month projections. Despite the IHME's wider intervals at almost every time point, the DOW model outperforms the IHME projections in terms of consistency (i.e interval overlap), and is more accurate for longer-term projections. Note that we have not validated these projections for the second uprising in deaths, so although our model outperforms the IHME up until June 30th, an extension of our methodology is likely required to accurately forecast COVID-19 mortality in the second "wave" and beyond.

4 | DISCUSSION

One interpretation of the day of the week effects estimated by our model is that people are more likely to die from COVID-19 on certain days of the week. It may be the case that people are more likely to contract the disease on given days of the week (e.g weekdays), which may cause them to pass away at higher rates in the following days. It is also possible that on certain days of the week, hospitals are more crowded and have fewer available resources, which could increase mortality on those days. However, it is more likely the case that deaths are equally likely to occur on any day-of-the-week, but are simply more likely to be reported on certain days due to hospital administration procedures. Hospital administrative workers likely work less on weekends, so deaths on the weekends may be reported several days after they occur, resulting in lower deaths on Sundays and Mondays. In either case, our results show that regardless of whether or not these are true day-of-the-week effects or are simply an artifact created by inconsistent reporting, accounting for day-of-the-week effects is important when projecting mortality during epidemics. Our non-DOW model also outperformed the IHME projections, which shows the potential benefits of using a skew-Normal and/or hierarchicalizing the model parameters.

One limitation of our model is that although we estimate day-of-the-week effects, the forecasts output from our model are projections of the number of COVID-19 deaths reported on certain days, as opposed to the true death count. Further extensions of our model could echo the methodology of Seaman, Samartsidis, Kall, and De Angelis (2020), allowing for the forecasting of the true number of deaths on any given day.

Since June 30th 2020, many regions such as Florida have seen a second increase in COVID-19 deaths. Further extensions to our model need to be explored to account for a second increase in daily deaths. Prediction of these seconds increases could be performed using region-level covariates such as lock-down severity or mask usage in the region. The second peak itself could be potentially be modelled by adding a second skew-Normal, where the location and height of the peak daily deaths are related to the first skew-Normal's parameters. Additionally, some regions that appear to be undergoing a second increase but may just be having multiple epidemics in different subregions (e.g counties in the U.S). Having smaller-area level data could improve model performance because of this.

Another extension to this model could be to have a different day-of-the-week parameter per region, but this is likely only possible for regions that have at least several weeks of data. Hierarchicalizing the day of the week parameters is also an option. For example, the Monday effect for regions with less mature data could be estimated by "borrowing" information from regions where the epidemic has matured. Another potential extension is to have a day-of-the-week effect changing throughout time, as COVID-19 death reporting may have improved since the beginning of the pandemic. However, in the midst of an epidemic, computational efficiency is of the utmost importance, as these models can easily take over a week to run. Obtaining robust estimates quickly can help aid policy decisions which can ultimately save lives, so having projections within a day or two is largely beneficial.

Further analysis is needed to fully assess the predictive capabilities of our model by expanding the number of countries and regions included. Additionally, including more models for comparison would be ideal, such as the projections made at `https://covid19-projections.com/` (Scriby, Inc. 2020). This work focused on presenting our core methodology, and provide accurate single "wave" projections for four countries that were near or past the peak of their epidemic.

Despite these limitations in the DOW model, it seems to forecast mortality related to COVID-19 in the first "wave" quite well when compared to the alternatives. Our projections made in Section 3.1 can be seen as accurate projections assuming that COVID-19 restrictions were never loosened in each region. The day-of-the-week effect was shown to be important when forecasting COVID-19 mortality, as the IHME and non-DOW model seemed to be inconsistent with their projections over time, which could be because the projections were made on different days of the week. Our skew-Normal Bayesian hierarchical model with day-of-the-week effects could be used as the basis for future COVID-19 mortality predictions where the epidemic is less mature, or perhaps for future outbreaks where accurate and consistent mortality projections are needed, where only daily mortality data is available as the input.

5 | DATA AND CODE AVAILABILITY

Data, code, and results are all available at the Centre for Global Health Research's public github:

https://github.com/cghr-toronto/public/tree/master/covid/DOW

6 | ACKNOWLEDGEMENTS

This research was supported by Canadian Statistical Sciences Institute's rapid response program on COVID-19, and the Natural Sciences and Engineering Research Council.

References

Anastassopoulou, C., Russo, L., Tsakris, A., & Siettos, C. (2020). Data-based analysis, modelling and forecasting of the COVID-19 outbreak. *PloS one*, *15*(3), e0230405.

Berument, H., & Kiymaz, H. (2001). The day of the week effect on stock market volatility. *Journal of economics and finance*, *25*(2), 181–193.

Brown, P., Jha, P., et al. (2020). Mortality from COVID-19 in 12 countries and 6 states of the United States. *medRxiv*. Retrieved from `https://www.medrxiv.org/content/10.1101/2020.04.17.20069161v1`

Bukhari, Q., Jameel, Y., Massaro, J. M., D'Agostino, R. B., & Khan, S. (2020). Periodic oscillations in daily reported infections and deaths for coronavirus disease 2019. *JAMA Network Open*, *3*(8), e2017521–e2017521.

Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., … Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of statistical software*, *76*(1).

Chakraborty, T., & Ghosh, I. (2020). Real-time forecasts and risk assessment of novel coronavirus (COVID-19) cases: A data-driven analysis. *Chaos, Solitons & Fractals*, 109850.

Crevecoeur, J., Antonio, K., & Verbelen, R. (2019). Modeling the number of hidden events subject to observation delay. *European Journal of Operational Research*, *277*(3), 930–944.

Friedman, J., Liu, P., & Gakidou, E. (2020). Predictive performance of international COVID-19 mortality forecasting models. *medRxiv*. Retrieved from `https://www.medrxiv.org/content/10.1101/2020.07.13.20151233v4`

Gelman, A., & Hill, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press.

Gelman, A., Rubin, D. B., et al. (1992). Inference from iterative simulation using multiple sequences. *Statistical science*, *7*(4), 457–472.

Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo.

Journal of Machine Learning Research, *15*(1), 1593–1623.

Institute for Health Metrics and Evaluation. (2020). *Covid-19 projections.* Retrieved from https://covid19.healthdata.org/

Lloyd-Smith, J. O., Schreiber, S. J., Kopp, P. E., & Getz, W. M. (2005). Superspreading and the effect of individual variation on disease emergence. *Nature*, *438*(7066), 355–359.

Perc, M., Gorišek Miksić, N., Slavinec, M., & Stožer, A. (2020). Forecasting COVID-19. *Frontiers in Physics*, *8*, 127.

Petropoulos, F., & Makridakis, S. (2020). Forecasting the novel coronavirus COVID-19. *PloS one*, *15*(3), e0231236.

Sarkar, K., Khajanchi, S., & Nieto, J. J. (2020). Modeling and forecasting the COVID-19 pandemic in India. *Chaos, Solitons & Fractals*, 110049.

Scriby, Inc. (2020). *The coronavirus app.* Retrieved from https://coronavirus.app/map

Seaman, S., Samartsidis, P., Kall, M., & De Angelis, D. (2020). Nowcasting COVID-19 deaths in England by age and region. *medRxiv*. Retrieved from https://www.medrxiv.org/content/10.1101/2020.09.15.20194209v1

	Parameter	Prior	Description
	$A_{i,j}$	$N(\text{Mar } 29, 100^2)$	Location Parameter
	$A_{i,Brazil}$	$N(\text{Jun } 17, 45^2)$	Location Parameter
	$\frac{1}{\sqrt{\tau_j}}$	$\text{Exp}(1/10)$	Overdispersion Parameter
	D_{ij}	$\text{Exp}(1/10000)$	Spark Term
C_{ij}	α_j	$N_+(50, 40^2)$	Mean of C_{ij}, the severity parameters
	θ_C	$N_+(2, 0.66^2)$	Scale of the severity parameters
B_{ij}	η_j	$N_+(60, 30^2)$	Mean of B_{ij}, the duration parameter
	θ_B	$N_+(9, 3^2)$	Scale of the duration parameters
K_{ij}	ζ_j	$N_+(3, 2^2)$	Mean of the K_{ij}, the skewness parameters
	θ_K	$N_+(3, 2^2)$	Scale of the skewness parameters
	$R_{j,m}$	$N_+(1, 2^2)$	Day-of-the-week parameters

TABLE 1 Prior distributions for the DOW model in (1)

Cause	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Circulatory	1.016	1.003	1.005	0.995	1.003	0.997
Pulmonary	0.994	0.996	0.975	0.985	0.996	0.988
Circulatory and Pulmonary	1.011	1.002	0.999	0.993	1.002	0.995
Non-Circulatory Pulmonary	0.988	0.998	1.002	1.007	0.998	1.005

TABLE 2 Relative risks of days-of-the-week (relative to Sunday) for non COVID-19 related causes in Canada

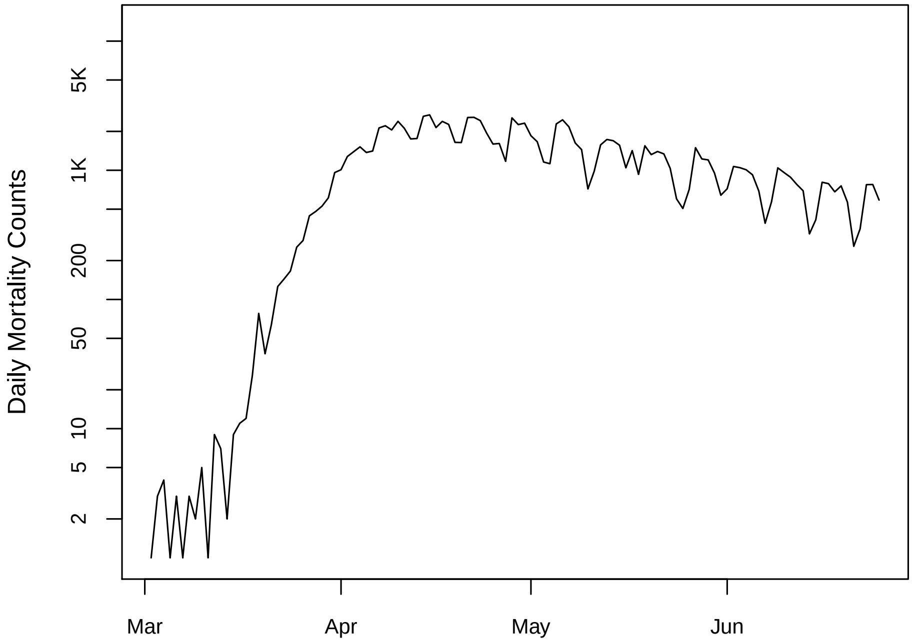

FIGURE 1 United States daily COVID-19 mortality from `www.coronavirus.app` (Scriby, Inc. 2020) from March 1, 2020 to June 25th 2020

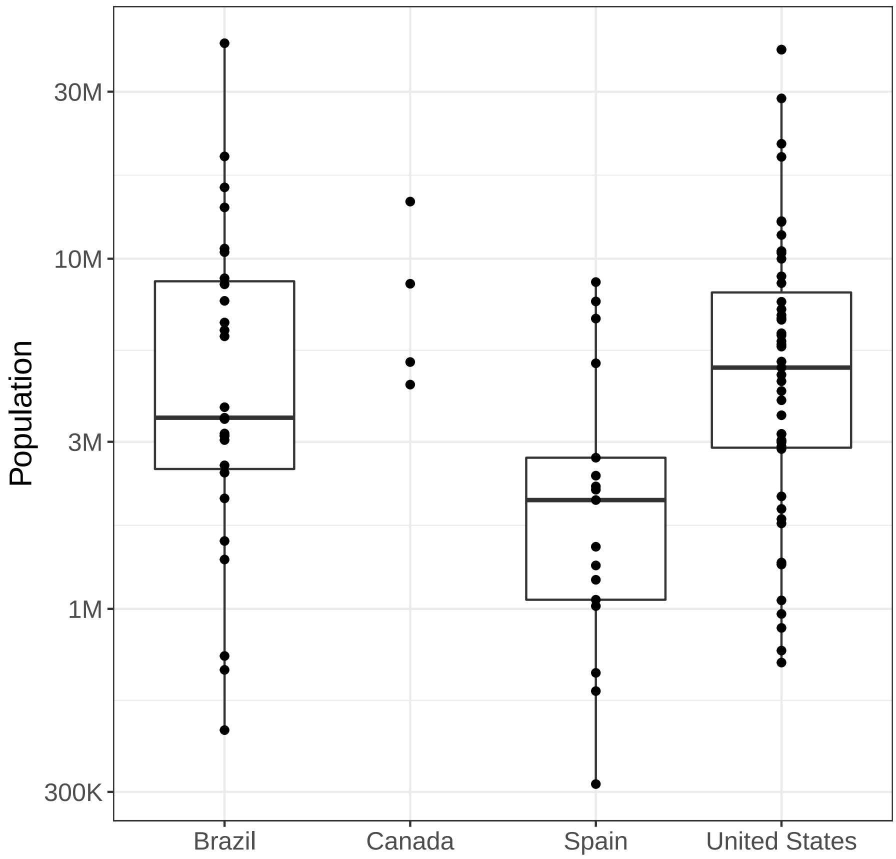

FIGURE 2 Populations of regions (black dots) by country. Note that the boxplot was omitted for Canada, as only 4 Canadian provinces were included

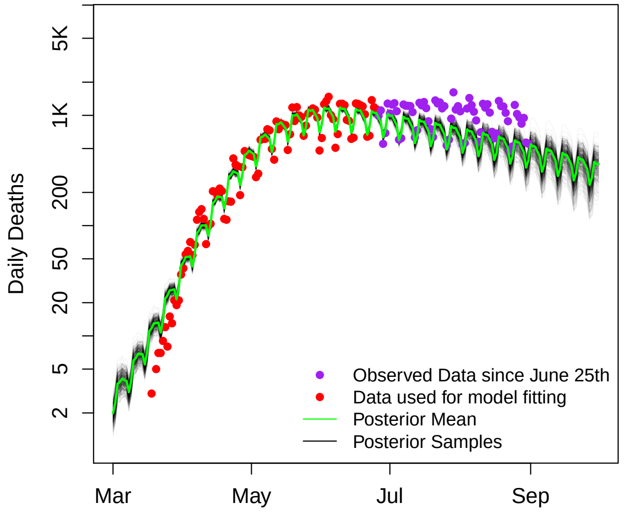

(a) Brazil

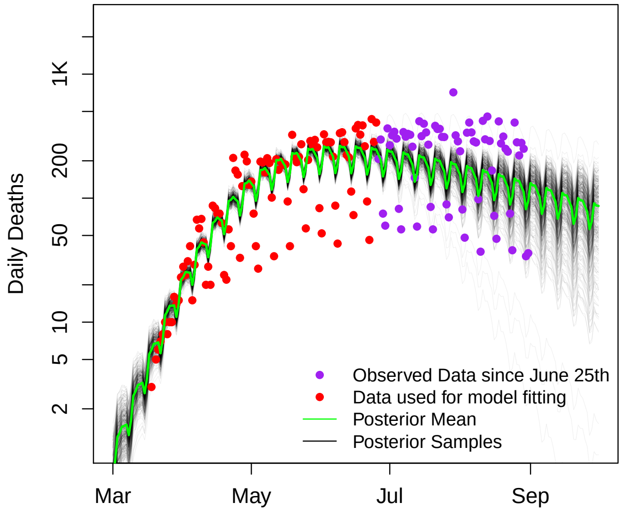

(b) Sao Paulo

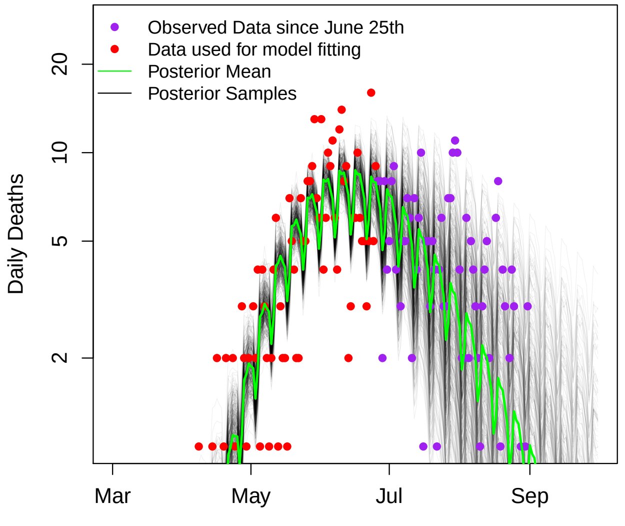

(c) Acre

FIGURE 3 Forecasting daily and cumulative deaths in Brazil as a whole, and the states of Sau Paulo and Acre.

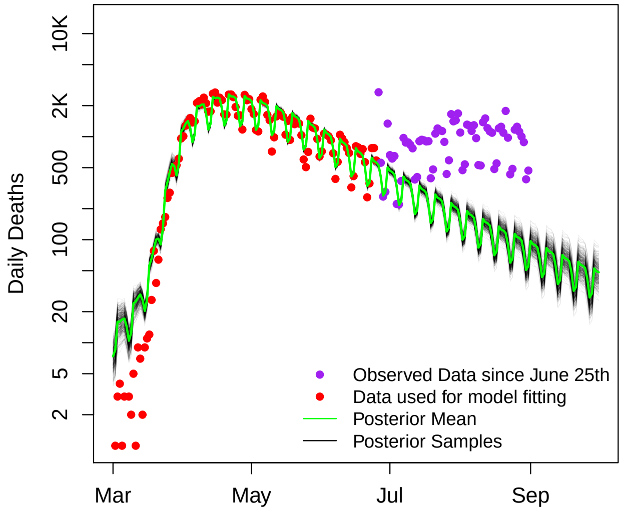

(a) United States

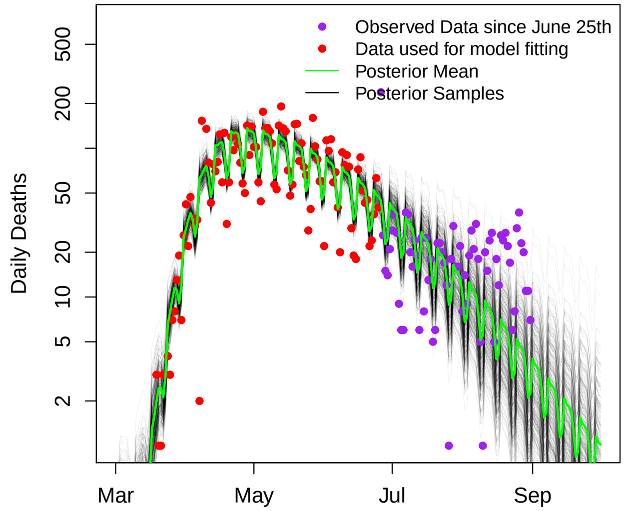

(b) Illinois

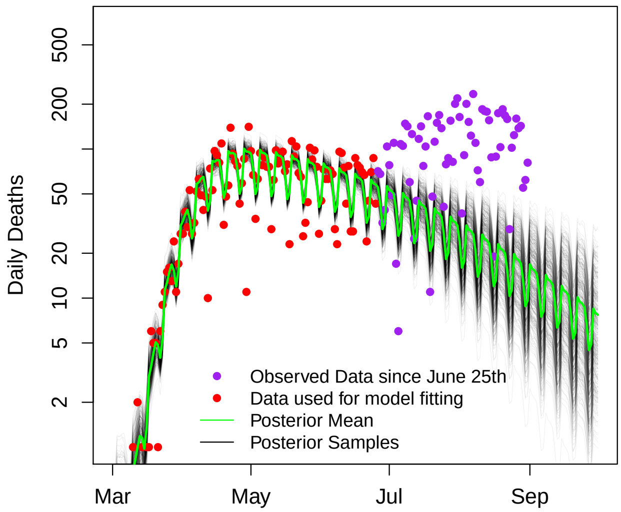

(c) California

FIGURE 4 Forecasting daily and cumulative deaths in the United States

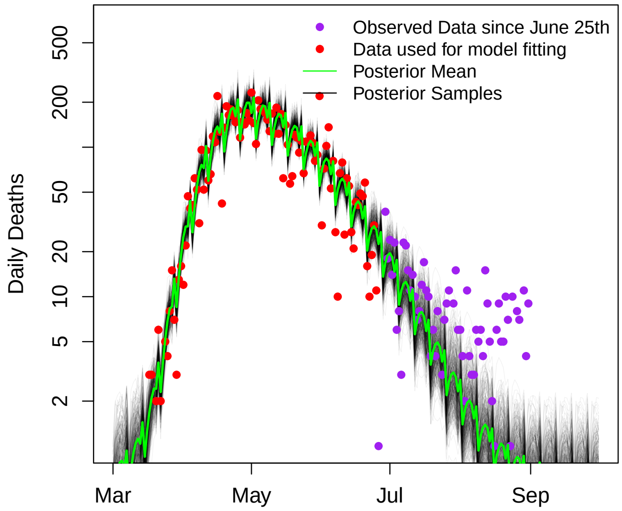

(a) Canada

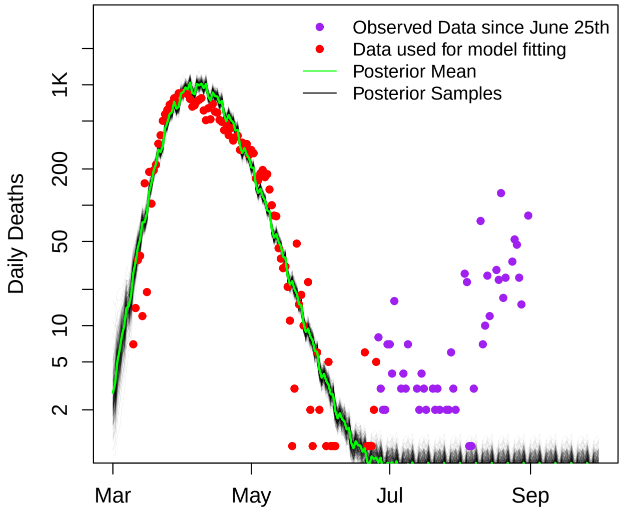

(b) Spain

FIGURE 5 Forecasting daily deaths in the Canada and Spain

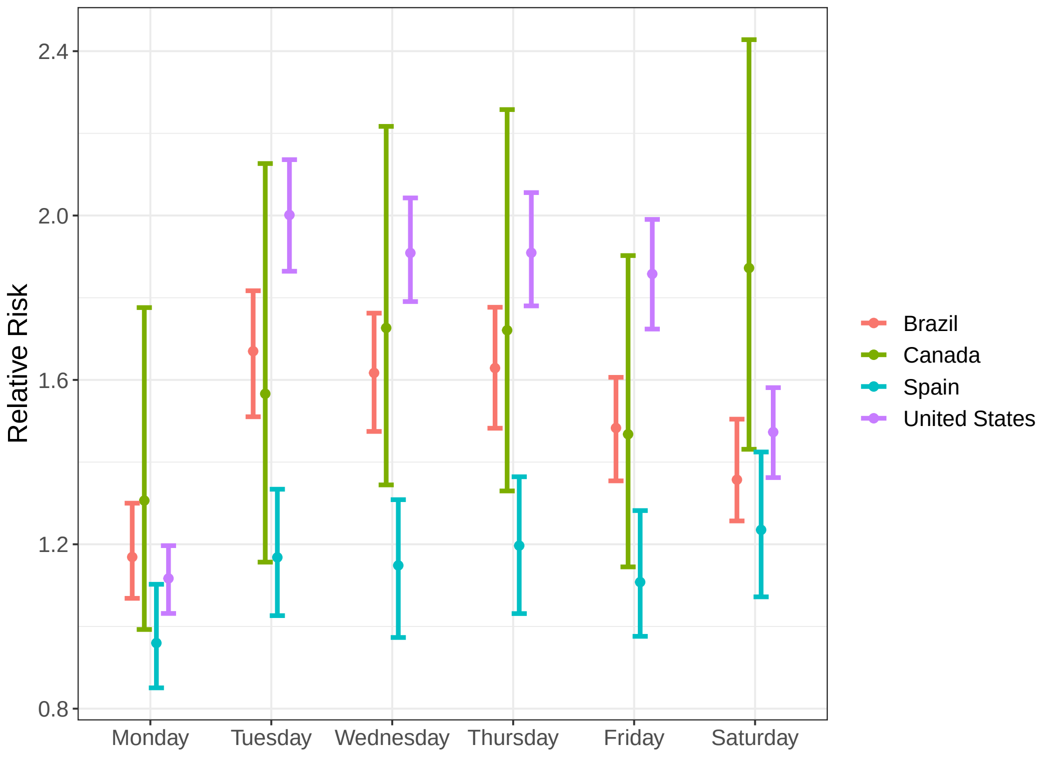

FIGURE 6 2.5th, 50th, and 97.5th percentiles of posterior distributions for day of the week parameters relative to Sunday (which is fixed at 1.)

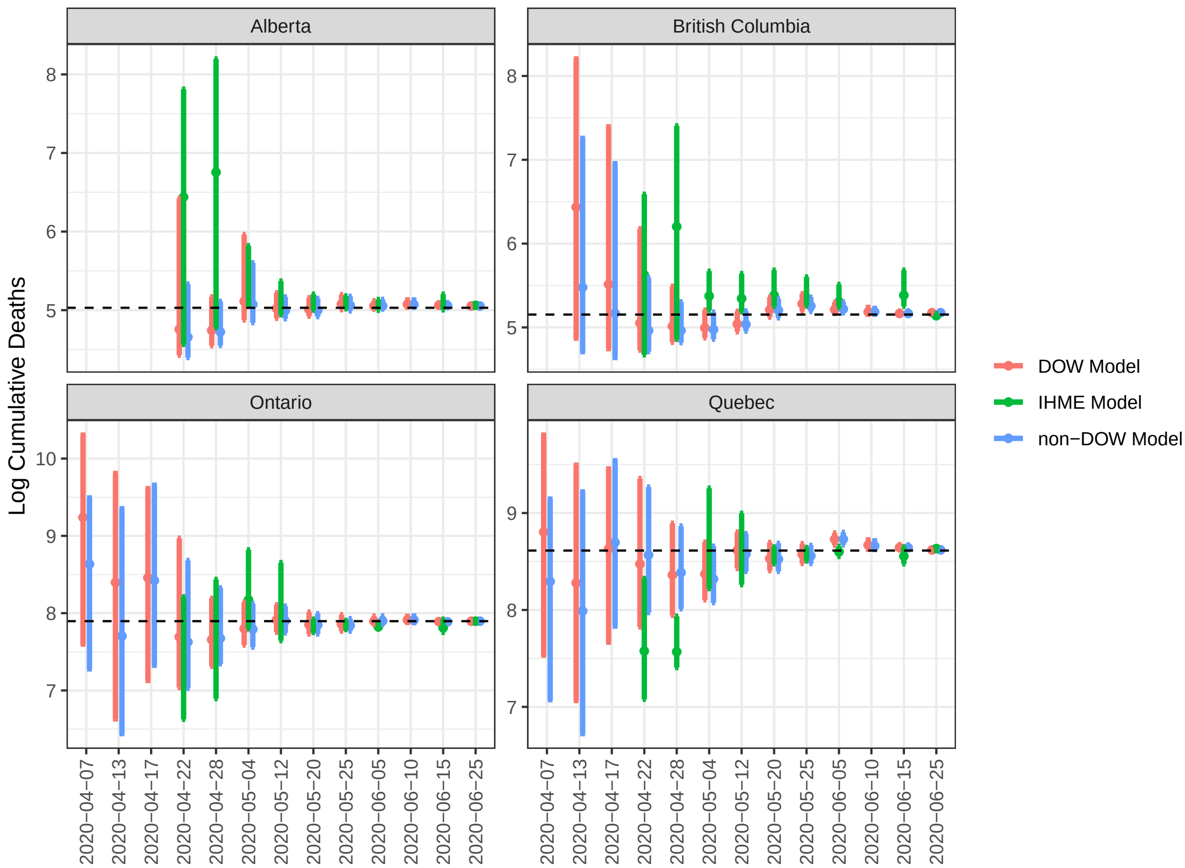

FIGURE 7 Cumulative mortality projections for June 30th made at 14 different dates starting April 1st 2020. The log of the cumulative mortality counts for June 30th are represented by the dashed line. Results are shown for the four Canadian provinces with the most COVID-19 deaths.

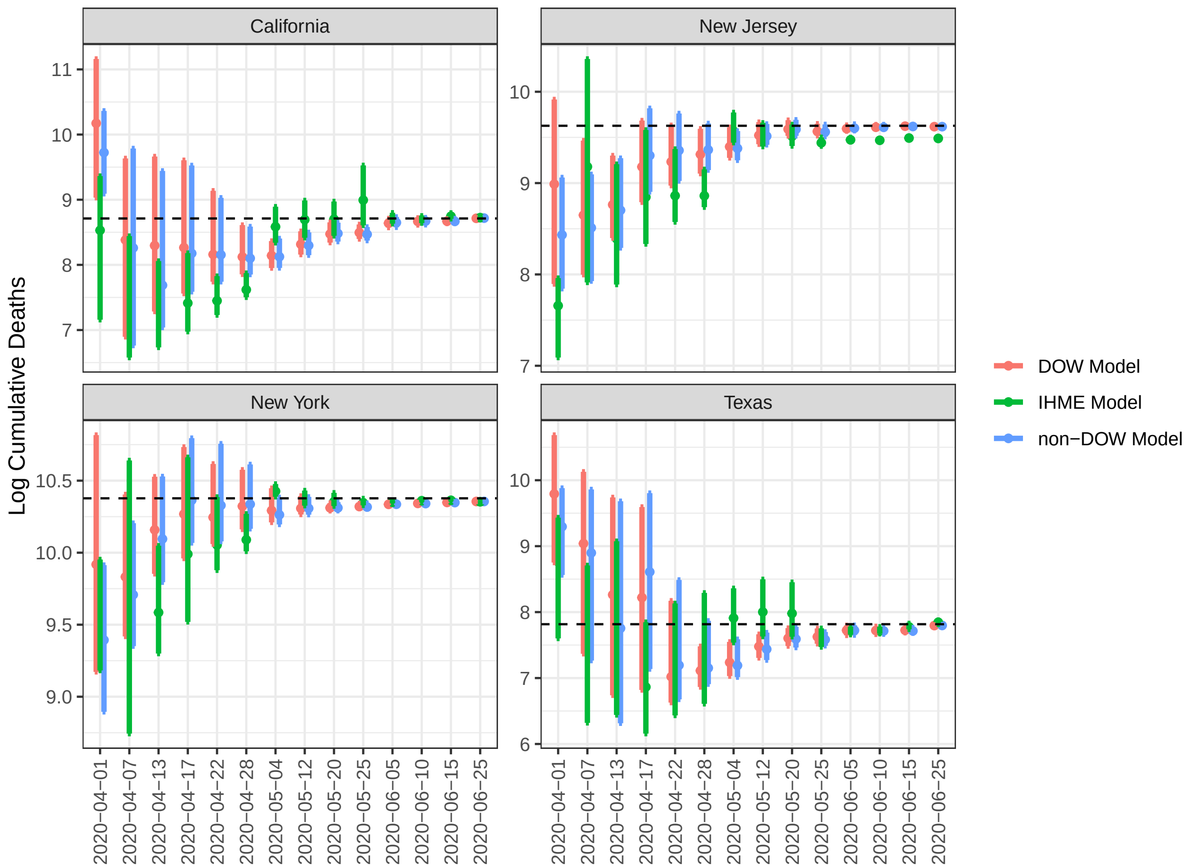

FIGURE 8 Cumulative mortality projections for June 30th made at 14 different dates starting April 1st 2020. The log of the cumulative mortality counts for June 30th are represented by the dashed line. Results are shown for the four U.S states with the most COVID-19 deaths.

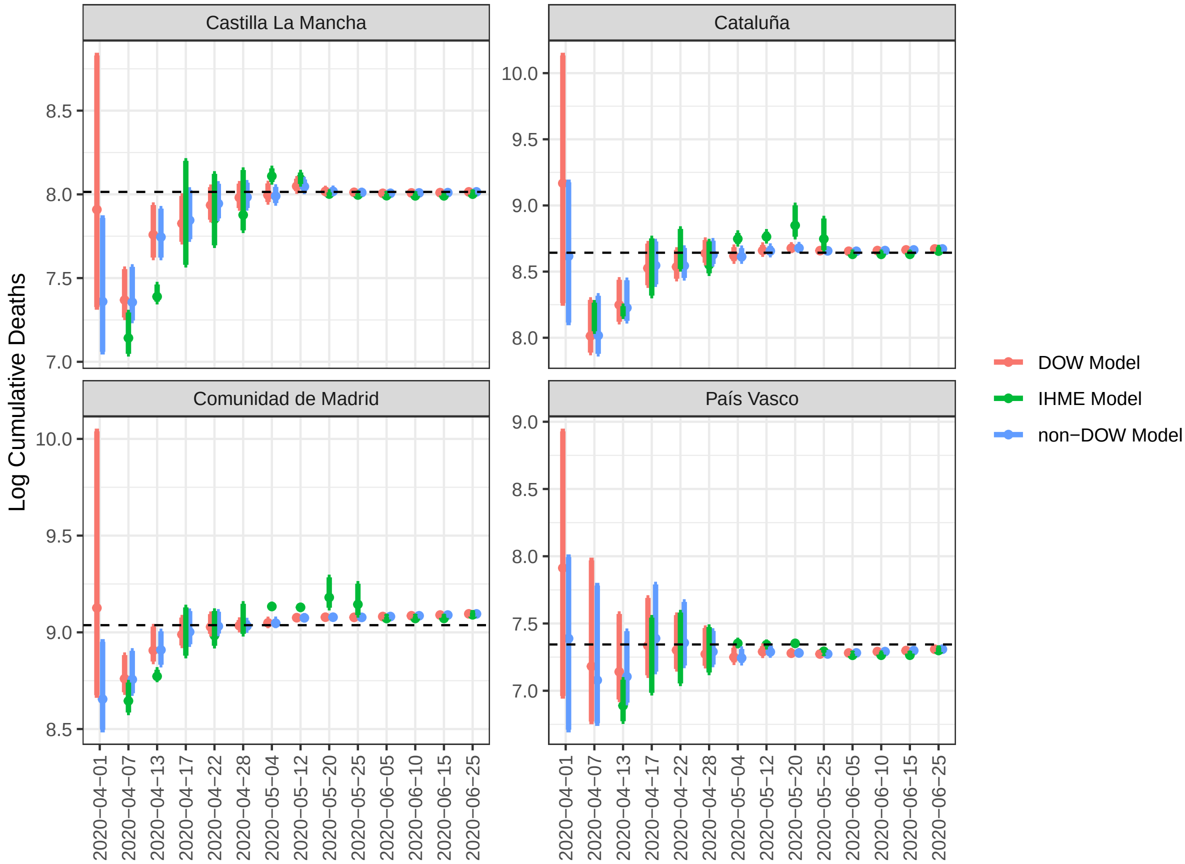

FIGURE 9 Cumulative mortality projections for June 30th made at 14 different dates starting April 1st 2020. The log of the cumulative mortality counts for June 30th are represented by the dashed line. Results are shown for the four Spanish Autonomous Communities with the most COVID-19 deaths.

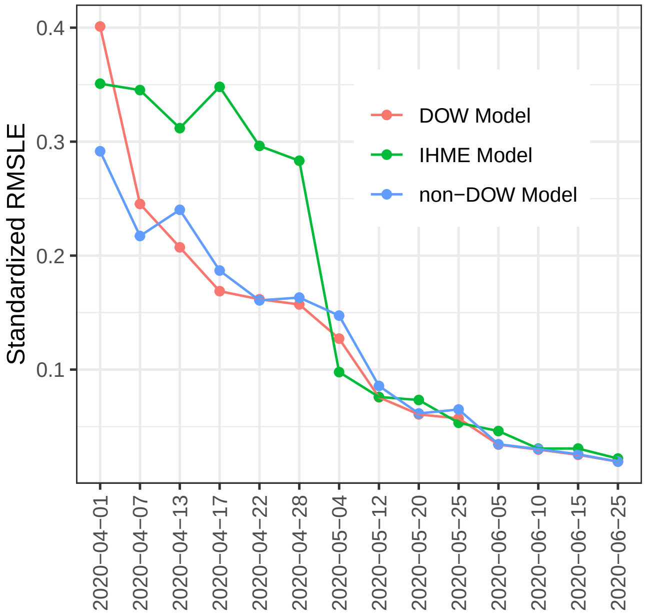

(a) Standardized Root Mean Squared Log Error

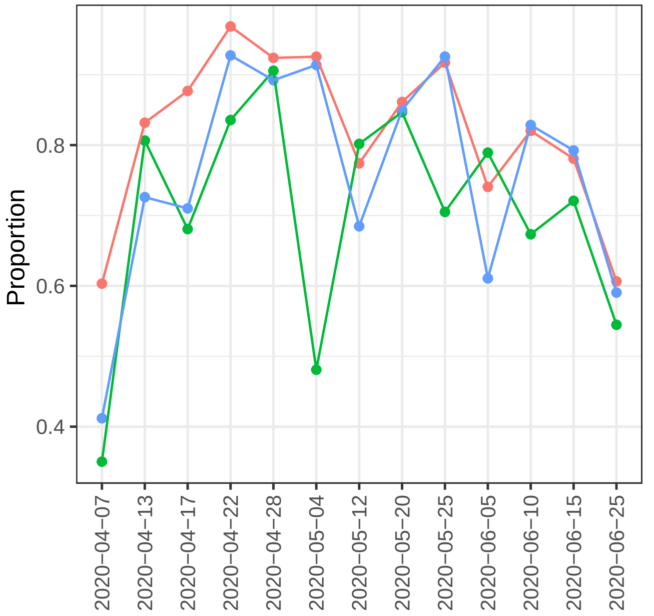

(b) Mean proportion of overlap between successive intervals

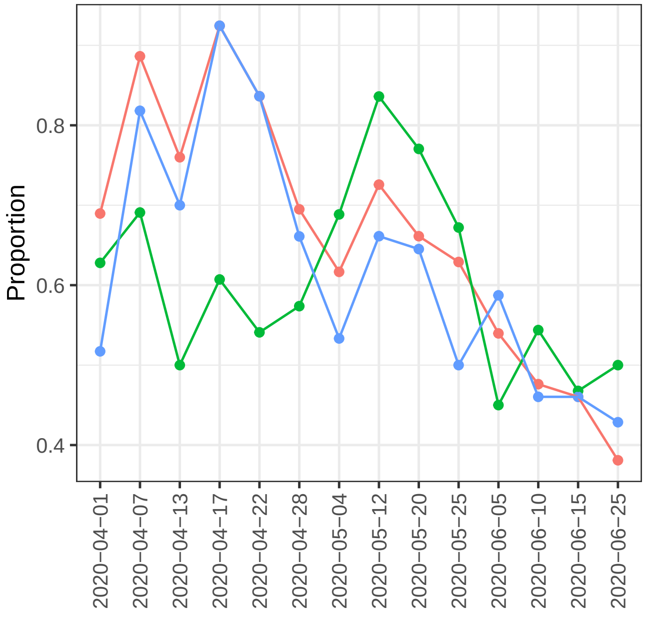

(c) Proportion of intervals that contained the true value for June 30th.

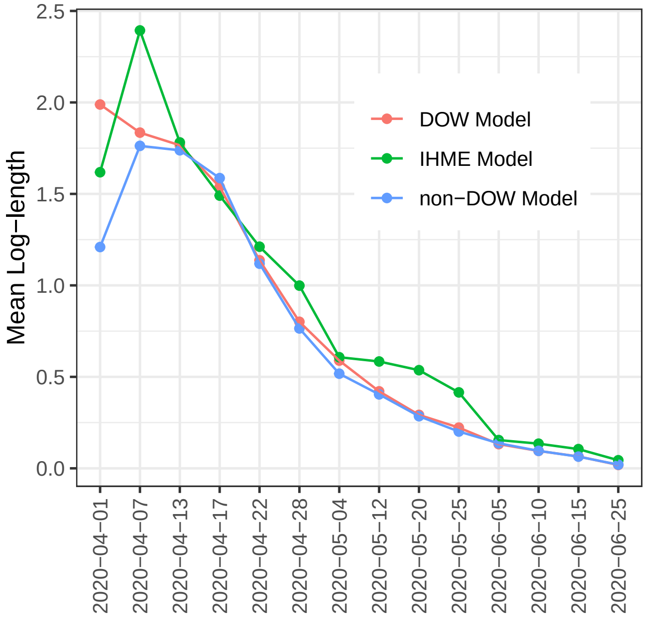

(d) Mean log-length of intervals

FIGURE 10 Comparing the DOW, non-DOW, and IHME models based on standardized Root Mean Squared Log Error, how much overlap there was between successive intervals, how often their intervals contained the true value, and their interval lengths