TN2007-1: Data Driven Pandemic Response Policy Evaluation (Sigfor)

George Rebane
31 July 2020

The Covid-19 pandemic has called on hundreds of local jurisdictions to fashion response policies that are suitable for their particular locales and demographics. In this technical note we introduce Sigfor, a computational tool that policy makers can use for monitoring the effectiveness of the current policy and also yield a timely prediction of the effectiveness of a policy change. The tool inputs a continuing stream of daily mortality data, and forecasts the expected maximum number of deaths and when this number would be reached, given that the current response policy remains unchanged.

Pandemic infection and mortality counts are shown in graphics such as in Figure 1 below. Each of these cumulative curves is a member of a family of such curves approximated by what in science and engineering are known as sigmoids. Sigmoid functions are a ubiquitous characteristic of our universe, they turn up in many places to record and illustrate growth in countless processes ranging from pandemic body counts, to the growth of plant and critter populations, to market penetrations of breakfast cereals.

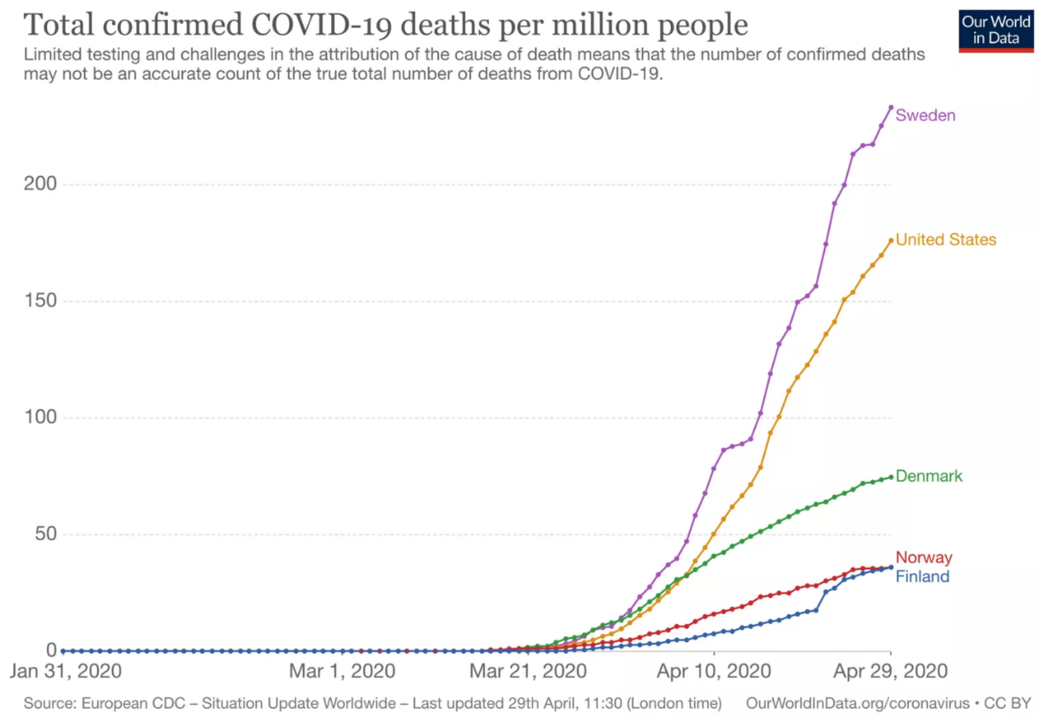

Figure 1

In the current development we will focus on pandemic mortality as experienced within a jurisdiction's target population in which a specific response policy is being implemented. As the daily mortality data comes in, the cumulative number of deaths quickly assume the shape of an increasing noisy sigmoid function of the kind shown above in Figure 1. Taking a more detailed look at the sigmoid, we illustrate its 'clean' basic version in Figure 2 below.

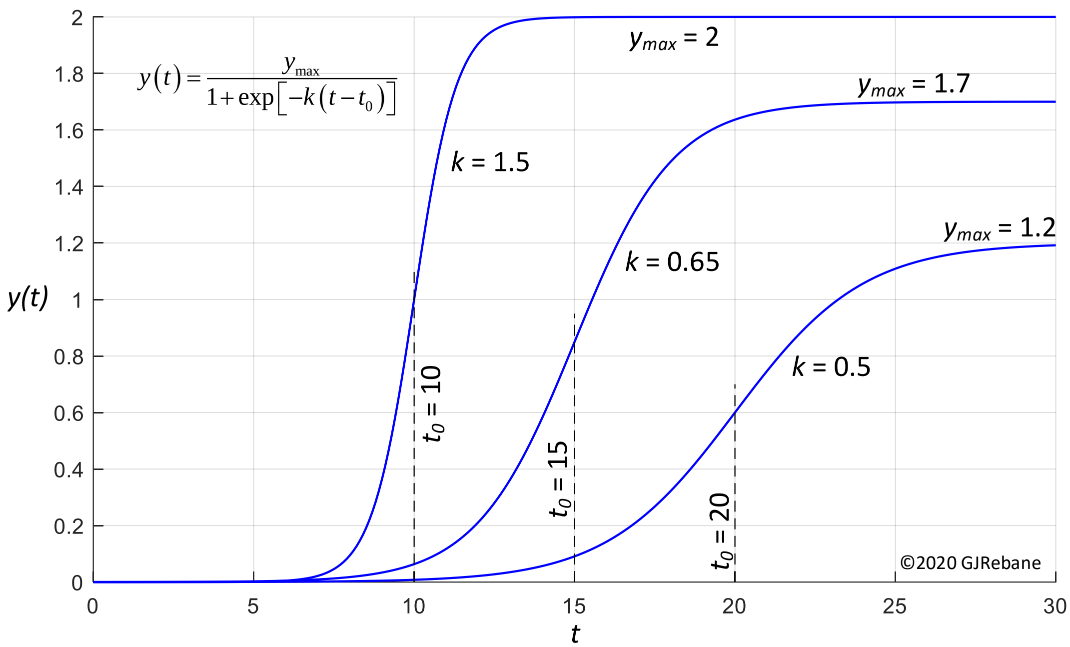

Figure 2

The most general sigmoid function shown above can be defined by specifying the values of three parameters which determine its eventual maximum value, how fast it grows (shape), and the time it goes through its maximum growth rate (at half of its eventual maximum value). From the sigmoid equation shown in the figure above, these are the y_{max}, k, and t_0 parameters respectively. Once these values are calculated, the entire shape of the sigmoid is available, allowing us to answer questions such as 'what is the maximum deaths we can expect as the pandemic runs its course here?', and 'how many weeks from now will the mortality count reach, say, 95% of its maximum?'

Sigfor computes the desired sigmoid's parameters with an estimation filter based on an expanding memory algorithm that has a constrained non-linear programming optimizer at its core. As the days pass, Sigfor accepts an updated time series of the cumulative number of deceased formed by adding each day's deaths count to the previous day's sum. This describes an expanding data window. Each day's Sigfor outputs are the best current estimates of the developing sigmoid's parameters. With these parameters in hand, it is an easy calculation to give updated answers to the above questions.

Figure 3 below illustrates the operation of Sigfor in a typical pandemic scenario. The horizontal axis is in the units of days ranging from zero to 300 (approximately 10 months), and the vertical axis is the number of deceased. The data shown was obtained from a simulation of real data generated from realworld covid mortality data statistics. Sigfor may be exercised and tested with various parameter values that describe the daily 'noisy' mortality count. In the figure, the red step function plots the number of daily deaths (multiplied by 10 for display purposes). The solid green stair-step trace is the actual record of cumulative deaths that forms the early-stage sigmoid input to the Sigfor estimator. The dashed green trace indicates the unknown true cumulative deaths if the current response policy remains unchanged. The dashed black trace indicates the predicted course of deaths calculated from the current input data window.

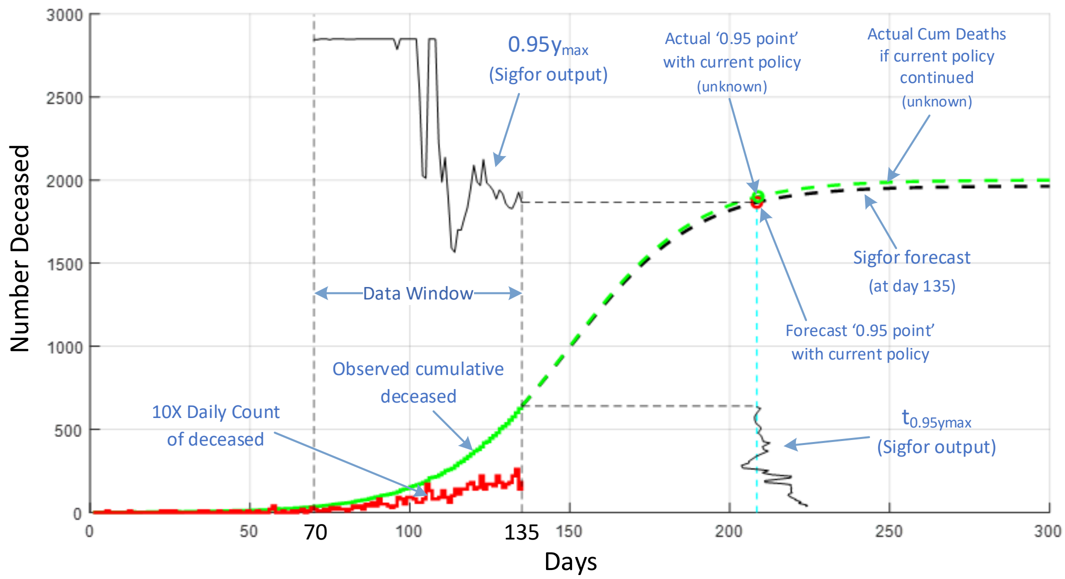

Figure 3

The Sigfor filter's sequential outputs of the 95% maximum deaths (0.95y_{max}) and the week at which this value will be reached are shown as the black (wiggly) traces which converge to their true values as the data window continues to expand. In this simulated scenario, with a data window started at day 70, we see that at day 135 we can already predict fairly accurately that the maximum deaths with the current policy will number almost 2,000 (the black dashed line), and that the 95% level of deceased will occur around day 210. Since the current deaths (at day 135) number less than 700, there is yet time to implement an alternative response policy that would reduce the total deaths expected under the current policy. Sigfor will then be used to predict the efficacy of the new policy as its new expanding data window is established. The impact of such response policy changes are illustrated in Figure 4 below.

Here we see the plots of cumulative deaths that would be the result of implementing three increasingly stringent response policies each starting at day 0. Of course, the shape of these three cumulative mortality sigmoids is not known to policy makers who decide to implement Policy 1 at day 0. The number of the deceased starts accumulating as shown by the red daily figures (X10), until at around day 80 they have enough data from Sigfor runs to get an acceptably reliable total death count of about 2,000. The policy makers consider this to be unacceptably high, and implement the more stringent Policy 2 at day 80. The number of daily deaths is seen to respond to this policy change, and cumulative deaths now continue to increase from the previous cumulative level at the pace indicated by the shifted Policy 2 sigmoid (shown by the cyan dotted line). This policy is continued until a new data window is built up while running Sigfor, and seeing its estimates converge to the Policy 2 values which indicate that total deaths now will be reduced to around 1,700. While this is lower than the 2,000 deaths predicted from Policy 1, the 1,700 number is still considered too high, and on day 120 an even more stringent Policy 3 is implemented.

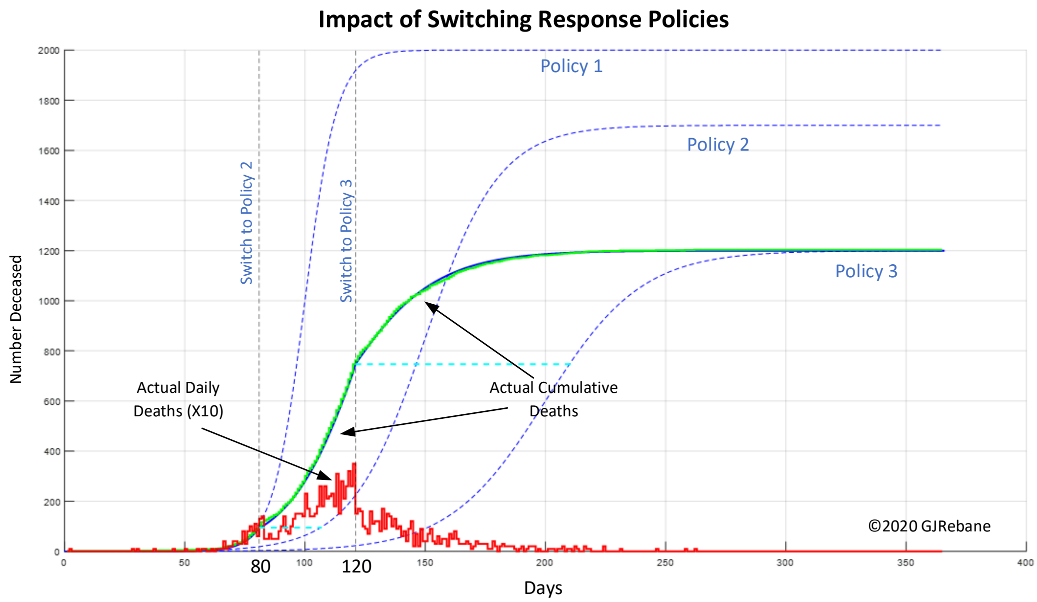

Figure 4

At this duration into the pandemic's outbreak (or resurgence) enough time has passed and deaths have occurred to reach near the top of the daily death rate hump for Policy 2, and fortuitously beyond the hump for Policy 3. Therefore, under Policy 3 the death rate begins responding markedly at day 120 as the most stringent policy is put into effect. Running Sigfor with the expanding data window from day 120 quickly shows that total deaths will now level out at around 1,200. This markedly reduced total is deemed acceptable, and shows that with the timely use of the Sigfor with incoming mortality field data, policy makers will be able to save 800 lives, or around 40% of the mortalities they would have experienced had they continued with Policy 1.

From a technical viewpoint it should be clear that Sigfor can also be run in parallel with several sigmoid functions which differ from the basic one illustrated here. Such contending functions can be readily derived from mortality curves already experienced in other jurisdictions. In such a multi-sigmoid estimation process, the 'best one' reveals itself by virtue of its minimum error when compared with the errors using its competitor sigmoid functions. In Figure 4, this error is the difference between the staircase green curve compared to the underlying projected smooth black curve.

Sigfor was developed by George Rebane, PhD. He is a systems scientist who may be contacted at gjrebane@gmail.com. Sigfor was developed in the Matlab™ environment, and may also be implemented in the MS Excel™ spreadsheet that includes the Solver add-in.