Is the Imperial College pandemic modelling work 'fit for purpose'?

By Matthew Edwards FIA for

COVID-19 Actuaries Support Group – Learn. Share. Educate. Influence.

Summary

Government responses to the Coronavirus are heavily informed by models. In the UK, one of the main models used is that developed by Professor Neil Ferguson's team at the MRC Centre for Global Infectious Disease Analysis, Imperial College London. This type of model incorporates behavioural elements to help users consider the impacts of different policies (for instance, school closures). It is a truism that all models are wrong in some respects, but models of the spread of COVID-19 are particularly susceptible to 'parameter error' because of the lack of reliable data so far, and the fact that changes (such as new 'lockdowns') will materially change some of the parameters in hard-to-predict ways.

Having studied the paper 'Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand' with knowledge of many of the associated modelling issues, we believe that the Government's response is based on an appropriate methodology, where the modellers have taken reasonable care in calibrating the model parameters.

Introduction

We assume that readers of this note are already familiar with the paper 'Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand' (Ferguson et al., 2020), and this note does not summarise any of the content.

Given how important it is that the Government's policies are effective in controlling the pandemic, there has naturally been considerable debate about modelling on which the Government's decisions have been made. This short note outlines the COVID-19 Actuaries Response Group's thoughts about the paper, and some of the criticisms that have been made of it.

Objective of the modelling

It is vital to be clear about what the modelling is trying to. It is this:

To approximately quantify different strategies to delay spread

It is also important to be open about the fact that all models are wrong to some extent, given that they are simplified representations of the underlying, very complex reality. A model of this type is particularly prone to 'parameter error', given the lack of good data available to decide on values for the parameters – or even the formulations that should be used.

How wrong could it be?

Given the aim noted above, many of the apparent weaknesses in the modelling become irrelevant in the sense that they would have no effect on the decisions made.

For instance, if the Infection Fatality Ratios assumed in the work (Table 1 of the paper) were to have been over- or under-estimated by, say, 25% or even more, the relative impact of the various strategies considered would not change. In fact, as we would expect, the authors have looked at the sensitivity of their results to changes in some of the (seemingly) key assumptions (choice of local triggers, R0 value, IFR values) and confirm that their results regarding relative effectiveness of different policies are insensitive to such parameter error.

We could instead ask where the model is likely to be wrong in other ways, and consider possible impacts – in particular, might they make the final decisions (based on the relative impact of different policies) wrong?

We note the following areas where the modelling looks open to challenge:

- *It is assumed that symptomatic people are more infectious than asymptomatic people)*: One of the bigger unknowns of the coronavirus is the extent of asymptomatic transmission, and it would be sensible to consider a reverse scenario where asymptomatic people are just as (if not more) infectious.
- *There will be material non-compliance with the policies when implemented:* the modellers have assumed material non-compliance, for instance (table 2 of the paper), they assume 25% non-compliance on social distancing of the over 70s.
- *Many of the deaths are 'accelerations' of deaths that would have occurred anyway.* It is true that a material (but small) proportion of the projected COVID-19 deaths would have died in any case in 2020, but allowing for this would not materially change the results. (There will also be a conversely acting effect whereby the swamping of hospitals by COVID-19 patients leads to increased mortality among people who might otherwise have survived.)
- *Lack of external validation or evaluation:* we comment on this below.

Overall, working through the paper with a view to 'what could go wrong' aspects of the modelling, we do not feel that there are any material features in the modelling that invalidate the UK Government's recent policy decisions.

We also considered criticisms raised by Chen Sen, Taleb et al., and do not consider any of these to be material. For instance, one of their major points was the contention that the virus would not return post-suppression; but, no country is completely unconnected to the rest of the world. If the virus survives in any country, unless we have developed herd immunity by that stage, the virus would return and start spreading again.

Modelling improvements

One area of interest to us is the challenge of heterogeneity in responses to policies across subsets of society (with particular regard to a socio-economic stratification). If non-compliance is higher in more deprived areas, and/or spreading is exacerbated in more deprived areas owing to inner city crowding, then this could materially alter the projected numbers of cases (but almost certainly not the relative attractiveness of the policies).

In the next round of model rework done by the Imperial College team, we think that some investigation of this variation would be valuable. This concern is triggered also by a concern that the IFR itself may exhibit similar socio-economic variation.

We would also expect to see some form of external validation of the model's results. For a model of this complexity, results can be surprisingly sensitive to particular combinations of parameters and data. One very useful form of validation would be evidence that a simpler model with similar structure and identical parameters leads to similar conclusions.

What the modelling scope excludes

A key part of the Government's decision making will have been to balance the mortality gains from severe containment policies with the economic cost of such policies. This balancing entails consideration of the likely medium-term impact on the health sector of reduced economic growth (in simpler words, the NHS needs "£X" next year to treat everyone for normal problems; with less money in the economy, finding those £X is problematic). The modelling done by Ferguson et al. does not include such considerations in its scope.

References

Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand, Ferguson et al., 16 March 2020

Review of Ferguson et al 'Impact of non-pharmaceutical interventions ...', Shen, Taleb et al., 17 March 2020