

The limits of a model

Without mathematical models, we would understand almost nothing of physics, from the motions of the planets to the behaviour of conductors, insulators and superconductors, fluid flows and turbulence. Across all science, modelling is our most powerful tool, as models let us focus on the few details that matter most, leaving many others aside. Models also help reveal the typically far-from-intuitive consequences when multiple causal factors act in combination.

The power of modelling has, of course, been vastly multiplied by computation. Yet the ease of such modelling also brings the temptation to mistake the model for reality, especially with models coupled to seductive graphics and video displays. A profound challenge for future science will be psychological — finding ways to ensure that scientists remain bound to legitimate evidence and logic even as visual display technology comes to make the output of even poor models seem highly persuasive.

The issue has been well illustrated by the COVID-19 pandemic. Around the world, as lockdowns and social-distancing measures slowed the initial spread, authorities wanting to re-open their societies have asked researchers to make projections of the likely future of the epidemic in different regions. Unfortunately, some models have produced enticingly optimistic scenarios that were anything but realistic.

For a time, for example, the Institute for Health Metrics and Evaluation (IHME; <https://go.nature.com/3bPm35o>) offered specific projections for all US states on when it would be safe to reopen, based on predictions for when the death rate would drop below 1 per million. Surprisingly — and suspiciously — these dates all fell only a month or so in the future, leading to reassuring downward sloping curves showing the epidemic ending quite soon. Too good to be true? Indeed.

As it turns out, these rosy projections simply reflected a choice in the modelling approach. The IHME model doesn't actually simulate the dynamics of the epidemic spreading, but fits a curve to the recent disease data. Moreover, it demands that the best fit be approximately Gaussian, with the up and down of the death rate being (roughly) a bell curve. This assumption alone guarantees the near-future disappearance of the epidemic. It would do so even if the recent data showed nothing but exponential growth. Combined with comforting curves, this kind of modelling risks creating misperceptions.

Similarly, other misinterpretations have plagued model-inspired discussions of the pandemic. Another common topic as nations try to resume some social and economic activity is the crucial importance of the reproduction number R as a metric for judging how far distancing measures can be safely relaxed. This is the number in epidemic models reflecting how many further infections will arise, on average, from a single new infection. As of early May, estimates of R in most European nations put it crudely around 0.8, while in the US it still appears to be just above 1. Plans for ending lockdowns stipulate that care should be taken to ensure R stays less than 1 — and generally assume all will be ok if it is.



The emphasis in modelling should be in accepting the likelihood of errors, and systematically searching to find out where models are likely to be wrong.

Yet this is somewhat naive, as R only describes what happens on average, and fluctuations about the average can have important consequences, as computer scientist Cris Moore recently pointed out (<https://go.nature.com/2WHgZzy>). The effective value of R will vary by location, and tends to be high where people don't or can't distance themselves very well. So, one should not be surprised by bursts of continued epidemic growth among particular groups or communities. Moreover, even with $R < 1$, a single new infection will sometimes give rise to hundreds of further infections just by chance. If $R = 0.8$, for example, then 1 new infection should lead on average to a total of 5 further infections. Yet computer simulations show that about 1% of these infections will lead to more than 50 others. Of 100 small towns experiencing 1 new infection, 1 will likely see such an outbreak. If the local hospital can handle only 10 cases, that's a crisis.

Other recent studies highlight just how difficult it is to use models to predict epidemic trajectories, especially given data limitations. One common modelling approach divides a population into groups of susceptible, exposed, infected and recovered

individuals. Such models predict a sigmoid shape of the total number of infections versus time. Using data from any region, this result can be used to make long-term estimates of the total number of infections. In this scenario, Davide Faranda and colleagues (Faranda, D. et al. *Chaos* 30, 051107; 2020) have made estimates of the sensitivity of this approach to the last available data point just before the inflection point of the $I(t)$ curve. In effect, they add some stochastic noise to the virus dynamics, to reflect many uncertainties in how the virus spreads, containment measures in place, and other factors.

They demonstrate that uncertainty in the last data point has a huge effect on the usefulness of long-term predictions, yet this uncertainty doesn't show up in the usual regression errors used to judge accuracy. Mean square error estimates can look excellent, giving false confidence in the forecast. For example, even a 20% error in the value of the last data point can lead to changes of several orders of magnitude in predictions of the final number of infections. One way to protect against such errors, they suggest, is to simply exclude the last data point and check the stability of the estimates. Or, add noise to the last data point so as to produce an ensemble of estimates, revealing just how much scatter there is in the prediction.

Another problem, discussed in a [Perspective](#) from Pasquale Cirillo and Nassim Taleb elsewhere in this issue, is the 'fat-tailed' nature of the distribution of epidemics. Unlike processes described by Gaussian statistics, it's not meaningful even to make calculations of quantities such as the expected number of deaths. And that's even without data errors.

The emphasis in modelling should be in systematically searching to find out where and why models are likely to be wrong or misleading. That's the only way to protect against thinking the model works like 'the real thing'. Erica Thompson and Leonard Smith of the London School of Economics have written about the need to escape from 'model land' — a dangerous territory into which every modeller may tumble (Thompson, E. L. & Smith, L. A. *Economics* 13, 2019-40; 2019). Theirs is a light-hearted article, yet serious, and the problems they discuss very widespread. □

Mark Buchanan

Published online: 8 June 2020
<https://doi.org/10.1038/s41567-020-0934-5>