

Richard Bradley and Liam Kofi Bright

Draft. Completed version available in: *The Philosophers' Magazine* 90:32-35 (2020)

Governments across the world have responded to the Covid-19 pandemic with measures that are unprecedented in peace time in terms of the degree to which they seek to reshape the behaviour of individuals and organisations. These include both public health measures, such as restrictions on movement and on social interaction, aimed at slowing the spread of the virus, and economic measures aimed at mitigating the potentially catastrophic impact of these policies on the economy. Policy makers have been drawing heavily on scientific expertise and advice to fashion the right policy; a fact which has also played no small role in the way that the policies have been justified to the public. This advice is in turn supported by sophisticated scientific modelling and especially epidemiological models of the spread of the virus in populations. Such modelling is however fraught with difficulty. While basic causal relationships are reasonably well-understood the sorts of details required for accurate prediction are not. Policy makers must therefore rely on science that is uncertain. What implications does this have for the way that policy choice should be made?

Let's start by looking at an idealized, but widespread, view about how policy decision should be made and the role of science in it. What we want to do is choose the policy option (from the set of those available to us) that has the best outcome. But, of course, we almost always don't know for sure what the outcome of any policy will be. So standard theories of rational decision making say that we should choose the policy with the greatest expected benefit or utility. To determine which this is, we need to know two things: the value of each of the different possible outcomes of a policy choice and, for each policy, the probability of each such outcome given the implementation of that policy.

Determining these two required factors for the purpose of assessing pandemic responses – the values and conditional probabilities of the outcomes – is very difficult however. We are primarily interested here in the consequences of a policy for people's lives and livelihoods. Feasibility constraints will require these two to be traded-off to some degree: policies that save more lives by suppressing the rate of infection cause significant damage to the economy and hence to many people's livelihoods, while those avoid economic disruption do so at the cost of more lives lost. Economists have a number of tools for making this trade-off, that in practice involve attaching a monetary value to people's lives (of a certain duration and quality) derived from individuals' preferences for trade-offs between their safety and their wealth. There is much that one can question about how this is done, but let us set those questions aside and focus on the second factor – the probabilities of outcomes.

It is here that science, and in particular epidemiological models of the pandemic, plays a crucial role by supplying predictions about how many people will be infected, how many will be hospitalised and how many will die under various stylised policy scenarios. In the UK for instance the government's adoption of social distancing measures was very strongly influenced by a model of the pandemic produced by the Imperial College Covid-19 team lead by Niall Ferguson. But many other models have been developed, based on different hypotheses about the main causal variables driving the pandemic and the relationships between them (such as between the infection rate and sociability), or using different estimates of the values of crucial parameters (such as what the fatality rate is amongst the infected) or state of the population (such as how many people are already infected). And these different models give quite different predictions about the numbers that will be infected, hospitalized or die. (Here's a simulator that will allow you see what difference choices of

parameter values make not only to these variables but to calculations of benefits of different policies <https://sites.google.com/site/marcfleurbaey/Home/covid>).

Underlying these differences is the simple fact that there is a good deal of uncertainty about all of these elements. Take estimates of the infection-fatality ratio, for instance. These have varied enormously between countries with a recorded ratio of over 10% in Italy and just over 1% in Germany for instance, a fact which probably reflects differences in the amount of testing for infection being conducted as much as anything else. Indeed, without knowing the percentage of the population that is infected it is very difficult to say what the true figure is, and on this question answers vary enormously (some small-scale tests for antibodies conducted in California, for instance, yielded infection rates more than 50 times higher than recorded ones). A different kind of uncertainty surrounds the question of what factors to model. The initial Imperial model for instance didn't include the effect of the swamping of health systems on fatalities due to causes other than the covid-19, nor the endogenous effect on social distancing of its spread (e.g. due to people's fear). Other models incorporate one of these, few incorporate both.

With time estimates will improve, as will the modelling that draws on them. But in the meantime, it is critical that the amount of uncertainty contained in the predictions they make is adequately captured so that policy makers know what they are dealing with. Uncertainty about inputs to the models (e.g. estimates about numbers currently infected) can be captured by making probabilistic predictions of outcomes e.g. by giving predictions of the form "With probability x , more than 100 000 people will die, with probability $x+y$ more than 10 000 will, ..." (something that is standard practice in other domains, such as forecasting natural catastrophes or stock market movements, which face similar uncertainty). But we still need to account for the other uncertainties, in particular those regarding the models themselves. Failure to do so risks inducing unjustified confidence in the models' predictions. For instance, it can lead policy makers to think that one course of action can be expected to yield more benefits than another, when the real situation is that it whether or not this is true depends on very particular parameter values or assumptions about causal relationships.

One way of capturing the remaining uncertainty that has found increasing support of late is to do so by specifying not just a single probability distribution over the outcomes of interest, but a family of them. If we think of each member of the family as being the distribution that we get from particular choices of parameter values and modelling assumptions, then we can see that size of the family will give a measure of just how much uncertainty we face about the consequences of our policy choice. And by looking at range of associated estimates of the expected benefits of a policy one gets a measure of how robust an assessment of a policy's usefulness is to scientific uncertainty.

It might seem that little is achieved, other than complexity, by presenting policy makers with ranges of estimates rather than a single precise one. But this is incorrect. This new way of thinking about our uncertainty allows policy makers to choose actions that can be expected to yield benefits across a very wide range of choices in preference to those that do not, even if amongst the latter are those that are optimal for some very specific choices. Or even to choose very cautiously, by choosing actions that have acceptable expected consequences under all possible choices. Suppose for instance that you have to choose between making two investments, each costing £1000, one of which is guaranteed to yield some small benefit (say £2000) and the other yielding either nothing or £21,000 depending on whether you complete it successfully or not. If the probability of success in the latter case is greater than 5%, then its expected return is higher than the former. Suppose your best estimate of the probability is 6%, but regard any estimate between 1% and 15% are reasonable. Then although the first action is optimal, it is not robustly so – on some estimates you are expected to lose your investment. On the other hand, the second action is robustly beneficial, though far from optimal on your best estimate.

Making it possible for policy makers to calibrate their choices of policy to the amount of uncertainty around the model-based predictions that they draw on for policy assessment, does not of course answer the question of just how much robustness they should seek or how cautious they should be. But it forces this question into the open, where it can be answered in a way that is appropriate to the nature of what is at stake, i.e. by democratic debate amongst those affected by the policies, all of us, and not just implicitly decided by modellers.