ASSESSING THE WELLBEING IMPACTS OF THE COVID-19 PANDEMIC AND THREE POLICY TYPES: SUPPRESSION, CONTROL, AND UNCONTROLLED SPREAD

Matthew Adler, Duke University
Richard Bradley, London School of Economics and Political Science
Maddalena Ferranna, Harvard School of Public Health
Marc Fleurbaey, Princeton University and Paris School of Economics
James Hammitt, Harvard University
Alex Voorhoeve, London School of Economics and Political Science

Abstract

The COVID-19 crisis has forced a difficult trade-off between limiting the health impacts of the virus and maintaining economic activity. Welfare economics offers tools to conceptualize this trade-off so that policy-makers and the public can see clearly what is at stake. We review four such tools: the Value of Statistical Life (VSL); the Value of Statistical Life Years (VSLYs); Quality-Adjusted Life-Years (QALYs); and social welfare analysis, and argue that the latter are superior. We also discuss how to choose policies that differentially affect people's wellbeing. We argue in favor of evaluating policies using a Social Welfare Function (SWF), which evaluates the possible distributions of wellbeing across individuals that may result from a policy. Such a function, we argue, should regard increases in the wellbeing of the less well-off as especially valuable. We then use a model to illustrate how such a framework can help evaluate two broad policy types in response to the pandemic: eradication of the virus, and more lenient control of the spread. Our model reveals how such evaluations depend on many empirical facts but also on key value judgments about the relative importance of health and on the extent of special concern for the worse off. The purpose of this brief is not to make precise recommendations, as conditions vary widely across countries and over time, but to provide a methodology.

The Challenge

The COVID-19 crisis puts all governments in a difficult position. In the absence of extensive testing capacities, they have to resort to near-universal lockdown and social distancing measures which exact a severe economic toll. While developed countries have the ability to provide temporary support to avert business collapse and worker hardship due to lack of liquidity, most developing countries do not. The choice between lives and livelihoods is starker for the latter.

However, by relying on swift measures of quarantining travelers, testing, and labor-intensive contact tracing before the number of infections rose to unmanageable levels, a few countries and states have managed to ward off the first wave of the pandemic. And some, as in South-East Asia, managed to do so even though their first cases appeared in January, before most other countries took serious measures. They could also count on the cooperation of populations accustomed to public health campaigns and protections against infectious diseases. Many developing countries have benefited from a longer warning period, and might be able to emulate these strategies.

Even in countries where the virus has spread widely, it is technically possible to control the pandemic. Lockdown measures observed in many countries reduce the reproduction rate of the pandemic to numbers that would guarantee its local eradication in a few months if the lockdown were sustained. It is also possible to keep the pandemic under control by a stop-and-go policy of periodic lockdown episodes of a few weeks until a vaccine is found. Two pressing questions with regards to these policies of suppression and control are: How should we conceptualize the benefits and burdens they produce as compared to a policy of uncontrolled spread? And, once we have these benefits and burdens clearly in view: how do we balance them in order to determine which policy is superior?

Our Proposal

We propose that governments rely on transparent evaluation methods in order to assess the wellbeing impacts of the pandemic and of the policy response. There is considerable uncertainty around key parameters of the pandemic as well as the reaction of the economy to exceptional measures. Nevertheless, it is possible to determine an order of magnitude for the key elements of policy decisions. No government can focus exclusively on epidemiological considerations or exclusively on economic considerations; the wellbeing of a population depends on both health and wealth. Finding the right balance requires relying on sound ethical principles and careful estimation of possible scenarios. Obtaining the understanding and cooperation of the population, which may be crucial for successful implementation, requires clarity about the objectives and the value judgments underlying the chosen policy.

Evaluation Methods

How do we determine what economic cost is worth bearing in order to reduce the number of deaths and other harms due to COVID-19? A well-known but often criticized method of policy evaluation is benefit-cost analysis (BCA), which works by converting all the effects of a policy into monetary equivalents and then summing up these equivalents. A different method, social welfare analysis, proceeds by determining the effects of a policy on individual wellbeing and then applying an aggregation formula to them to evaluate the overall effects of a policy. We briefly survey these methods and emphasize the advantages of the latter.

The monetary measure of the value of saving lives most widely used in benefit-cost analyses is that of the Value of a Statistical Life (VSL). When the individuals whose lives would be saved are not known, each of the people at risk would be willing to pay some money to increase their chance of survival. VSLs describe the monetary values that individuals attribute to a reduction in their own mortality risk. To be precise, these monetary values are derived from the rate at which people are willing to trade off small changes in their income against small changes in their risk of death. This, in turn, is estimated from individuals' reported preferences or from those that they reveal in workplace and consumption behavior, such as the choices they make amongst jobs involving different levels of risk or their purchases of risk-reducing equipment. For example, if someone would accept a pay cut of $1,000 per year to reduce their annual risk of mortality by 0.1% (but would not accept a larger pay cut), then we say that the monetary value of their statistical life is $1,000,000. Note that this is not the same as saying that they would be willing to accept $1,000,000 in return for certain death or would pay this amount to guarantee their survival. Rather, it means that each of 1,000 people, identical in all relevant ways, would, considering their self-interest alone, be willing to pay an equal share of a $1,000,000 cost for something that reduces the number of them expected to die in the year by one.

A person's VSL can depend on their characteristics such as their age, their income and wealth, or the overall level of risk they face. This can have unacceptable consequences for BCA. In particular, the fact someone who is well-off is likely to place a higher monetary value on risk reduction than someone who is less well-off implies that if individual-specific VSLs are used in a benefit-cost analysis of policies, the interests of the well-off will count for more than those of the less well-off (because the monetary value they place on reducing their risk of death is higher).

By using a single VSL, such as the population average, rather than individual-specific ones, this problem is avoided. But others are then created. In particular, it seems reasonable to treat people in different age-groups differently when assessing policies. Death is certain for all mortals and is commonly considered a more serious loss from the societal and ethical perspective when it occurs earlier in life. Reasoning in terms of life years preserved rather than lives saved appears to better take account of this widespread sentiment. The skewed age distribution of COVID-19 fatalities makes this problem especially pressing.

A common solution is to use a different measure for policy evaluation: the Value of a Statistical Life Year (VSLY). The VSLY is obtained by dividing the average VSL of the population by the average life expectancy remaining (an individual's current life expectancy remaining is the number of years she can be expected to live, if she doesn't die now). The value of saving the life of someone in any particular age cohort is then given by the product of the VSLY and the life expectancy remaining for the cohort. This yields a value of life saving that varies by age.

One criticism of both VSL and VSLY is that they do not take into account quality of life. Many people would not regard a year of life spent bed-ridden as equivalent to a year of life in excellent health, for instance (but see below). Quality Adjusted Life Years (QALYs) are a way of allowing the value attributed to a life saved to depend on both its remaining length and its quality. The value of living in an impaired health state—say, with diminished lung function due to COVID-19—is derived from people's preferences. These preferences may be elicited in a number of ways. Individuals may simply be asked to assign a numerical value to life in a particular health state in comparison to both death and life in full health. Alternatively, they may be asked how they would balance a longer life in an impaired state against a shorter, healthier life, or asked what risk of death they would be willing to run in order to be fully cured of their impaired health. These preference-based assessments can be questioned (Hausman 2015). For example, there is evidence that healthy people are poor predictors of what life would be like in states of impaired health (Dolan and Kahneman 2008, Walasek et al. 2019). Nonetheless, even rough indicators of the quality of life in impaired health states can be better than measures that neglect quality altogether. This is particularly clear in the pandemic, in which it is important to take into account the often substantial effect of contracting the illness on those who do not die from it. For this reason, the use of QALYs in public health decision making is widespread.

It is also controversial. One key concern is that when it comes to life extension, the use of QALYs regards as more valuable the life years gained by people who would, if saved, be in good health than the life years gained by people who would, if saved, live with disabilities or in poor health because extending the lives of the former would generate higher health-related quality of life (National Council on Disability 2019). In our view, this objection is best addressed not by rejecting the use of QALYs, but by assigning special value to improvements in the quality (and length) of life of those who are worse off (John, Millum, and Wasserman 2017).

Estimates of the VSL and VSLY vary considerably between countries. Part of this variation is due to differences in income per capita. For example, Robinson et al. (2019) recommend calculating the average VSL for a country by multiplying the country's per capita income by a factor proportional to the square root of per capita income (but no smaller than 20), and a VSLY by dividing the VSL by the average remaining years of life. For the USA, the typical VSL is around $10,000,000, and the VSLY a little over $300,000 (see also Kniesner and Viscusi forthcoming); for a country with per capita income of $10,000, a VSL of around $670,000 and a VSLY of around $20,000 (international dollars) can be suggested.

A similar variation is observed in the monetary costs that public actors regard as reasonable to incur to gain one QALY. One approach is to estimate the value of a QALY by dividing the VSL by the average remaining QALYs (Hirth et al. 2000). This produces values modestly larger than the VSLY. By contrast, the World Health Organization (WHO) has suggested that interventions that generate a QALY for less than 1 times per capita income are good value for money and that interventions that generate a QALY for up to 3 times per capita income may be worth the cost (Bertram et al. 2016). In line with this formula, for the USA, the Institute for Clinical and Economic Review suggests values between \$100,000 and \$150,000 for one QALY (ICER 2018). In contrast, the British National Institute for Health and Care Excellence applies figures in the \$25,000 to \$40,000 range for one QALY (NICE 2008). The fact that these countries have different per capita incomes only partially explains these differences. As we show below, the ranking of policies to deal with the pandemic based on benefit-cost analysis may well depend on which values are adopted. It is therefore critical that attention be given to the justification of any particular choice.

The alternative approach we emphasize here, social welfare analysis, proceeds by measuring the joint health and economic impact of policies on individual wellbeing and then aggregating individual wellbeing gains and losses to yield an overall measure of how beneficial a policy is. This method has a singular advantage over BCA. Unlike population-average VSLs, the individual-specific wellbeing values that social welfare analysis uses are sensitive to individuals' characteristics, such as their age and income. And while the individual-specific VSLs, VSLYs and values of QALYs of the well-off are inflated, relative to those of less well-off individuals, by the fact that money has relatively lower marginal value for the well-off, this is not true of wellbeing values. So, the aforementioned bias in favor of the well-off that the use of individual-specific monetary values introduces into BCA does not plague social welfare analysis.

Social welfare analysis begins with a measure of the wellbeing levels associated with different possible lives, represented by bundles of the goods that matter to individuals: income, health, longevity, and so on. Many methods are used to obtain these measures (for a review, see Adler and Fleurbaey 2016). One draws on reported levels of subjective wellbeing or life-time satisfaction scores to identify the determinants of wellbeing (Clarke et al. 2018). Another derives a wellbeing measure from individuals' preferences between probability distributions (lotteries) over alternative possible lives (Adler 2019). Yet another relies on income, corrected for the value of non-market aspects of life, such as longevity, on the basis of population preferences over these aspects (Blanchet and Fleurbaey 2013). Finally, the capability approach measures opportunities in various aspects of life (Sen 1999).

The effect of a policy that reduces the risk of contracting COVID-19 by some percentage and at some financial cost can be modelled by the shift in the population distribution of wellbeing levels generated by the bundles of longevity, health, and income mentioned above. Since this shift captures not just the impact on individuals' longevity and health of the policy but also how these factors, together with income, co-determine changes in individuals' wellbeing, it provides all the information required for a comprehensive analysis of the overall effect of implementing the policy.

To determine this overall effect, social welfare analysis proceeds by aggregating the set of individual wellbeing values achieved by implementing the policy. It does so by means of a social welfare function (SWF) that assigns to each distribution of individual wellbeing a measure of social value. A commonly used SWF is the utilitarian one, which assigns to each set of individual wellbeing values the total (or the average) of the values. This way of aggregating individuals' wellbeing is insensitive to inequalities of wellbeing in the population—in other words, it is indifferent to whether a given increment in wellbeing accrues to a well off person or a badly off person. In this respect, it is in tension with the common conception that a given improvement in the wellbeing of the worse off matters more than the same improvement to the well off, because it comes to those who are in greatest need or because such improvements reduce inequality (Adler 2019; Voorhoeve 2019). This problem can be addressed by using distribution-sensitive SWFs that prefer policies that produce wellbeing gains for those with low wellbeing over policies that produce the same wellbeing gains for those who are better off.

The choice of SWF is fundamentally an ethical one, as it requires balancing the wellbeing interests of different individuals. This is a second important advantage of social welfare analysis: it allows for these ethical choices to be made more explicitly and transparently than with BCA, which largely ignores the distribution of benefits and burdens among individuals.

Reviewing Policy Options: An Illustration

To illustrate these approaches, a model simulating the pandemic as well as lockdown and testing policies has been developed and adapted to several countries (USA, UK, France, Belgium, and Guinea).[1] The model takes account of inequalities in income and life expectancy across social groups, and allows for various assumptions about the distribution of the economic cost and the fatality burden among these groups. Such assumptions relate to policy choices about social protection and income support, as well as access to health care. This model is not designed to make predictions or precise policy recommendations, but rather to highlight the parameters relevant to sound decision making.

The illustration here focuses on the comparison between an *suppression* policy which implements a lockdown order ("contact reduction" of a certain percentage) for as long as it takes to make the virus almost completely disappear, on the one hand, and a *control* policy which limits the lockdown to periods following weeks where fatalities are over a certain threshold. The latter policy may also eventually cause the virus to nearly completely disappear, but over a longer period of time and with a succession of shorter lockdown episodes rather than a single one. Supplementary policies (testing, mask-wearing, income support) can also be taken into account and will modify the evaluation. The simulations we present include testing and mask-wearing through the "contagion reduction" that they entail. It is assumed that only contact reduction has a substantial economic cost, because testing and producing protective equipment are relatively low cost and do not slow down economic activity—even if displacing resources toward them has an opportunity cost. In this section, the discovery of a treatment or a vaccine is assumed to come too late to affect the evaluation of these policies. How relaxing this assumption would affect the analysis is briefly described at the end.

Fig. 1a illustrates the policy problem. The model starts with a single infection in week 1, which then spreads over the population. The policy (suppression or control) starts on a particular week (from week 10 to week 23) and is continued until the pandemic virtually disappears. Either policy entails an economic cost (due to lockdown)[2] but also saves lives, compared to the absence of intervention. The graphs do not show the outcomes over time, but the final outcome (economic cost and fatalities) as a function of *when* the policy is initiated. When the policy starts on week 23, it is almost as if no

[1] The model is contained in an Excel spreadsheet and can be downloaded freely from https://sites.google.com/site/marcfleurbaey/Home/covid. Users can change all parameters and assumptions and determine the timing and intensity of contact reduction and testing policies. Refined versions of the model (in Python code) are available upon request.

[2] In these simulations, it is assumed that contact reduction by 70% reduces economic income during the lockdown period by 35% (half the percentage).

intervention took place, because the first wave of the pandemic has almost fully passed. Therefore, our analysis also covers the case of uncontrolled spread. The scissors pattern for both policies shows that an *early* policy saves more lives (though a very early control policy is counterproductive in this respect because it hinders the build-up of collective immunity) but has greater economic cost because it requires longer lockdown episodes.[3]

Figure 1: The policy problem

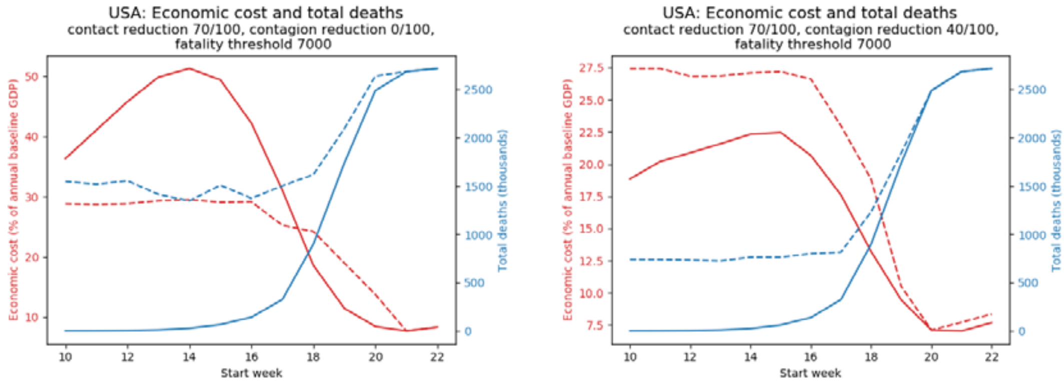

(a) (b)

How to read the graphs: solid curves describe the outcomes of the suppression policy, dashed curves the control policy; economic cost in red (left axis), total fatalities in blue (right axis). Note that the left axis has a different scale in (a) and (b).

The contrast between Fig. 1a and 1b shows the complementarity of contagion reduction with contact reduction. In the absence of contagion reduction (1a), the control policy is less costly but it also saves fewer lives. With ambitious contagion reduction (1b), in contrast, control is dominated by suppression, because the latter saves more lives at a smaller cost. This occurs because thanks to contagion reduction, the suppression of the virus can be achieved in a shorter time.

The benefit-cost analysis can be illustrated by the contrast between the VSL and the VSLY approaches, in the COVID-19 context where the victims are mostly the elderly (at least in developed countries). This is shown on Fig. 2, which displays hypothetical simulations for the total cost (adding up the value of lives lost to the crisis and the economic cost).

[3] As shown on the graph, a very early suppression policy can have a smaller cost, though, because a quick elimination of the virus is then possible before it spreads widely over the population.

Figure 2: The VSL and VSLY approaches

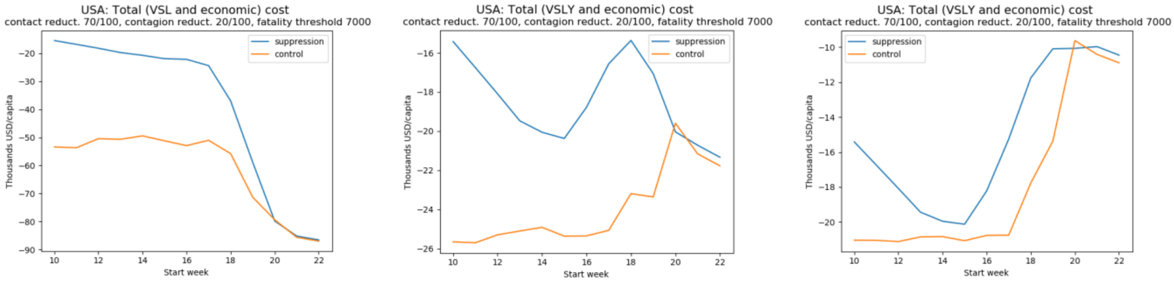

(a) (b) (c)

How to read the graphs: The vertical axis displays losses as negative numbers so that the higher on the axis, the better. This makes the graphs directly comparable to Fig. 3 below.

In Fig. 2a-b, both approaches favor the eradication policy over the control policy, the VSL approach is favorable to early eradication whereas the VSLY approach would also condone letting the virus spread before eradication is implemented. This is because elderly victims of the virus do not lose many years of life. These two graphs rely on a VSL equal to 150 times the GDP per capita and a VSLY equal to 3 times the GDP per capita. Fig. 2c relies on a lower VSLY equal to the GDP per capita, and there one sees that, in spite of the staggering death toll, delaying policy adoption is acceptable and that very late adoption of either policy is optimal. This illustrates how crucial the ethical parameter for the value of life is.

The SWF approach is also implemented in the model. Individual wellbeing is computed for a whole life, not just for a year, and depends on income and longevity. It is calibrated in a way that guarantees that the willingness to pay of the average individual for a life year is equal to the same VSLY used in the BCA method (3 times GDP per capita). In this way, in the absence of priority for the worse off, the VSLY approach and the social wellbeing approach deliver similar assessments. But when a degree of priority for the worse off is introduced, the evaluations can differ markedly. The social wellbeing approach is then uniquely sensitive to three considerations. First, given the assumptions about the value of longevity, the worse off include the victims of the virus, because their loss of longevity is a very substantial wellbeing loss. Concern for the worse off therefore assigns greater significance to health outcomes than economic outcomes. However, there are several elements to take into account. Many of the victims have attained an old age, so that, compared to younger survivors, they have the advantage of having avoided other fatality risks to reach that age, but also the disadvantage of having lived earlier, in less affluent economic times. The very worst off tend to be among the middle-aged victims, whose premature deaths are great losses.[4]

Second, inequalities in life expectancy and fatality rates across social groups reinforce the concern for health, because the worse off in income incur a double penalty through a greater health toll. Unlike

[4] In the version of the model with which the figures have been produced, these considerations are incorporated in a simplified way. The victims are assumed to have had greater life expectancy than the average level in the population (because most of them have already reached an old age), whereas the (younger survivors) are assumed to face the prospect of a greater lifetime income. But no detailed depiction of inequalities in income and fatality rates among age groups is made. A refined version of the model in construction will have this feature.

BCA, the SWF approach thus clearly identifies the value of policies which give greater access to health care and reduce the correlation between health and income.

Third, inequalities in the economic cost of lockdowns may attenuate the previous considerations if the disadvantaged social groups are more severely affected by the economic slowdown.[5] This is especially relevant for poor countries in which the most disadvantaged may fall into extreme poverty under these circumstances. When priority is given to the worse off, the SWF approach favors strong social protection measures that ensure a more equitable distribution of the economic cost. This point is illustrated in Fig. 3. It displays one particular example of a measure of societal wellbeing, the equivalent income, defined as the level of income per capita which, equally distributed and associated with equal longevity for everyone, would yield the same societal wellbeing as the contemplated (unequal) situation.

Figure 3: The SWF approach with priority to the worse off

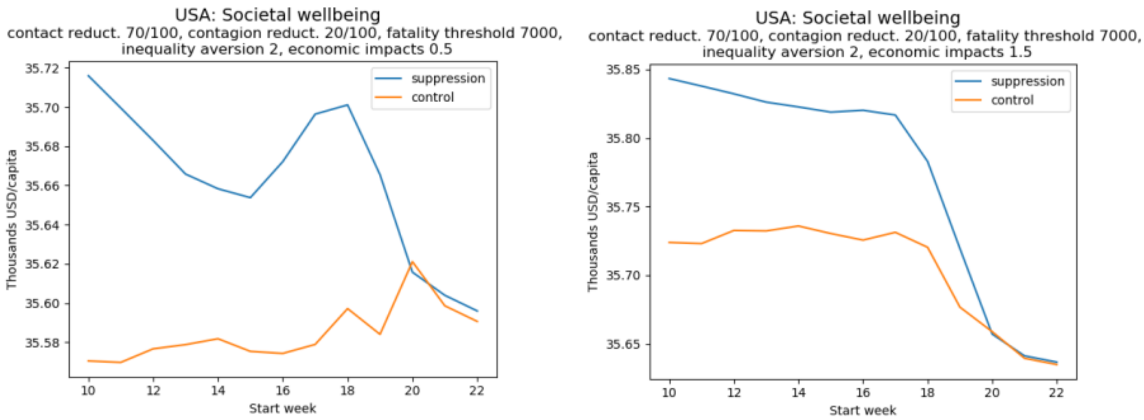

(a) (b)

How to read these graphs: The vertical axis is measured in equivalent income per capita which reflects societal wellbeing, taking account of inequalities. As in Fig. 2, the higher the better.

In Fig. 3, the elasticity of the economic cost to income is either 0.5 (Fig. 3a), meaning that the distribution of the cost is regressive, or 1.5 (Fig. 3b), meaning that the distribution is progressive.[6] The priority for the worse off is substantial in Fig. 3, and implies that improving the wellbeing of an individual who is half as rich is four times as important for the social evaluation.

Fig. 3a illustrates the fact that if the lower income groups in the population bear a disproportionate economic burden, and if this economic burden is heavy due to a lack of contagion reduction, then late may be preferred to mid-time adoption of a suppression policy, because the economic cost is substantial when the policy starts in the middle of the time range. And late adoption of a control policy is preferred to any other time. In contrast, Fig. 3b illustrates the situation in which, with a progressive distribution of the economic, early policy adoption is preferred when eradication is considered, and as far as control policy is concerned, very late policy adoption is not optimal at all.

[5] This particular model does not include the long-term and indirect effects on people's health and other long-term outcomes, but it does include the additional deaths not due to the virus but caused by the disruption of health care (either on the supply side or on the demand side, when patients for other conditions shun health care facilities out of fear).

[6] Inequalities of income are represented by distinguishing quintile groups. An elasticity of 0.5 means that a group with an income level greater by 1% than another group bears a share of the economic cost that is greater by 0.5% (therefore, the cost as a fraction of its baseline income is half a percentage point lower).

There are many parameters on which these simulations depend, but which cannot be discussed here in detail. For instance, if contact reduction can be raised from 70% to 80%, early eradication is much quicker and much less costly and does appear preferable even if one adopts a low figure for the value of life, keeping the other parameters as in Fig. 2c. The model incorporates spontaneous contact reduction by individuals when they witness a peak of mortality, or additional fatalities when hospitals are overwhelmed. The results would also change with an early discovery of a treatment, which would lower the fatality rate and enhance the relative value of controlling the pandemic compared to full suppression. The demographic structure of the population and background health conditions may affect the lethality of the virus; in higher-income countries, deaths are concentrated in the elderly but it appears that in developing countries, worse background health makes younger patients more vulnerable than in rich countries. There are important missing elements in these simulations. In particular, they do not assume any spatial heterogeneity between regions in a country, they ignore possible contamination by travelers coming from abroad, and they do not finely distinguish the situations of different age groups. They also ignore possible longer-term economic and health consequences of the crisis.

The upshot of these remarks is that precise prediction with this, or any other model, is not possible at the moment. The point of this policy brief is primarily to propose this methodology as a framework within which the combined effect of different values for relevant empirical and ethical parameters on policy conclusions can be studied. When available, more precise predictions can be introduced in this framework to support decision-making that is both scientifically and ethically well-informed.

Our model shows how sensitive the path of the pandemic is to the various parameters, reflecting how much uncertainty there is, which makes decisions quite hard for policy-makers. What is especially difficult, from a policy point of view, is the following. Looking at the figures presented here, it is tempting to conclude that near-total suppression of the virus is, in most cases, the dominant policy. But the duration of the required lockdown period can stretch over 4 months at 70% contact reduction (3 months at 80% contact reduction), if the start date is around a time corresponding to week 15-16 in our figures, as seems to be the case in many countries.[7] This means that, even when public health authorities are convinced that eradication is the best policy, implementing this policy may be politically difficult and requires very strong support measures for the population suddenly deprived of work and income. If, in the middle of the way towards near-eradication, the authorities revert to a more modest control strategy, then much of the effort has been in vain, because they have only pushed the infection wave into the future (see also Gollier 2020 and Kissler et al. 2020). This is why a clear communication on the strategy and the ethical choices, based on rigorous modelling as proposed here, may be crucial to convince the population of the need to stay the course.

In conclusion, although the current crisis presents a difficult trade-off between lives and livelihoods, especially for governments in fragile states or with frail leadership and a low degree of cooperation in the population, it is possible to lay out the main considerations that should guide policy, including the key normative issues about valuing lives and giving priority to the worse off. Our quantitative analysis illustrates how the relevant empirical and normative elements of sound policy-making can be put together into a rigorous framework, and why the SWF approach which takes account of the distribution of impacts and of background inequalities is more attractive than the BCA approach and, in particular, more consistent with widely shared ethical views regarding life-saving by income and age.

[7] Most countries which have been successful in eradicating the virus (such as Vietnam, New Zealand, or South Korea) have adopted very early lockdown and testing policies and have endured little damage compared to the other countries. China appears to have implemented a very strict lockdown policy (perhaps more than 90% contact reduction in the Wuhan area).

As illustrated here, modelling costs and health impacts may reveal dominant policies which would not have been obvious otherwise. In particular, it is crucial to check if the less politically difficult (in the short run) strategy may turn out to be worse, all things considered. However, there are many aspects to policy that cannot be resolved by dominance considerations, and this is where weighing the relative importance of fatalities and economic costs is required.

References

Adler, Matthew D. (2019). *Measuring Social Welfare: An Introduction.* Oxford University Press.

Adler, Matthew D. and Ferranna, Maddalena and Hammitt, James K. and Treich, Nicolas (2019) "Fair Innings: The Utilitarian and Prioritarian Value of Risk Reduction over a Whole Lifetime" (November 16, 2019). Duke Law School Public Law & Legal Theory Series No. 2019-79.

Adler, Matthew D. and Fleurbaey, Marc (2016). *Oxford Handbook of Wellbeing and Public Policy*. Oxford University Press.

Adler, Matthew D., Hammitt, James K. and Treich, Nicolas (2014). "The Social Value of Mortality Risk Reduction: VSL versus the Social Welfare Function Approach." *Journal of Health Economics* 35: 82-93.

Bertram, Melanie, Lauer, Jeremy, De Joncheere, Kees, Edejer, Tessa, Hutubessy, Raymond, Kieny, Marie-Paule, Hill, Suzanne (2016). "Cost–Effectiveness Thresholds: Pros and Cons." *Bulletin of the World Health Organization* 94:925-930. doi: http://dx.doi.org/10.2471/BLT.15.164418

Clark, Andrew, Flèche, Sarah, Layard, Richard, Powdthavee, Nattavudh, and Ward, George (2018). *The Origins of Happiness: The Science of Wellbeing over the Life-Course.* Princeton NJ: Princeton University Press.

Dolan, Paul and Kahneman, Daniel (2008) "Interpretations of Wellbeing and Their Implications for the Valuation of Health." *The Economic Journal*, Vol. 118(525): 215-234.

Gollier, Christian (2020). "Cost-benefit analysis of age-specific deconfinement strategies," *Covid Economics* 24: 1-31.

Hirth Richard, Chernew Michael, Miller Edward, Fendrick Mark, Weissert William. Willingness to pay for a quality-adjusted life year: in search of a standard. *Med Decis Making*. 2000;20(3):332-342. doi:10.1177/0272989X0002000310

ICER (Institute for Clinical and Economic Review) (2018). "ICER Value Framework". https://icer-review.org/wp-content/uploads/2018/05/ICER-value-framework-v1-21-18.pdf [accessed 23 May 2020].

John, Tyler, Millum, Joseph, and Wasserman, David (2017). "How to Allocate Scarce Health Resources without Discriminating against People with Disabilities," *Economics and Philosophy* 33 (2): 161-186

Kissler, Stephen, Tedijanto, Christine, Goldstein, Edward, Grad, Yonathan, and Lipsitch, Marc (2020). "Projecting the Transmission Dynamics of SARS-CoV-2 through the Postpandemic Period. *Science* 368 (6493): 860-868. DOI: 10.1126/science.abb5793

Kniesner, Thomas J. and Viscusi, W. Kip (forthcoming) "The Value of a Statistical Life" *Oxford Research Encyclopedia of Economics and Finance*. Vanderbilt Law Research Paper No. 19-15.

Layard, R., A. Clark, J.-E. De Neve, C. Krekel, D. Fancourt, N. Hey, and G. O'Donnell, "When to release the lockdown: A wellbeing framework for analysing costs and benefits," CEP Occasional Paper, 49, 2020.

Love-Koh J., M. Asaria, R. Cookson, S. Griffin (2015). The Social Distribution of Health: Estimating Quality-Adjusted Life Expectancy in England. *Value in Health* 18: 655-662.

NICE (National Institute for Health and Care Excellence). (2008). *Social Value Judgments*, *Principles for the development of NICE guidance 2nd edition*. https://www.nice.org.uk/about/who-we-are/our-principles [accessed 23 May 2020].

Robinson, L., Hammitt, J., & O'Keeffe, L. (2019). Valuing Mortality Risk Reductions in Global Benefit-Cost Analysis. Journal of Benefit-Cost Analysis, 10(S1), 15-50. doi:10.1017/bca.2018.26

Salje H., C. Tran Kiem, N. Lefrancq, N. Courtejoie, P. Bosetti et al.. 2020, Estimating the burden of SARS-CoV-2 in France. pasteur-02548181.

Schelling, T.C., "The Life You Save May Be Your Own," in S.B. Chase (ed.), *Problems in Public Expenditure Analysis*, Brookings, Washington, D.C., 1968.

Sen, Amartya.K. (1999). *Commodities and Capabilities*. Oxford University Press.

Verity R., L. C. Okell, I. Dorigatti, P. Winskill et al. (2020). Estimates of the severity of coronavirus disease 2019: a model-based analysis, The Lancet.

Voorhoeve, Alex (2019). 'Why Health-Related Inequalities Matter and Which Ones Do.' In. O. F. Norheim, E. Emanuel and J. Millum (eds.) *Global Health Priority Setting: Beyond Cost-Effectiveness.* Oxford University Press, pp. 145-62.

Walasek, Lukasz, Brown, Gordon D. A. and Ovens, Gordon D. (2019). "Subjective Wellbeing and Valuation of Future Health States: Discrepancies between Anticipated and Experienced Life Satisfaction." *Journal of Applied Social Psychology* 49 (12): 746-754.