Lives v livelihoods, part 1: how can we measure the value of a life?

Policies that suppress or control the COVID-19 pandemic prevent illness and save lives, but exact an economic toll. How should we balance lives and livelihoods to determine which policy is best? In the first of two posts, ***Matthew Adler (Duke University/LSE), Richard Bradley (LSE), Maddalena Ferranna (Princeton), Marc Fleurbaey (Princeton and Paris School of Economics), James Hammitt (Harvard)*** *and* ***Alex Voorhoeve (LSE)*** *compare the benefit-cost and social welfare approaches to doing so.*

What is benefit-cost analysis?

Benefit-cost analysis evaluates a policy in terms of the sum of the monetary equivalents of its outcome. The most widely used monetary measure of the value of saving lives is the Value of a Statistical Life (VSL). This is 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. For example, suppose someone would accept a pay cut of $10,000 per year, but no more, to reduce their annual risk of mortality by 0.1%. Then the monetary value of their statistical life is $10,000/0.001 = $10,000,000. (What this means is that, if we had a population of 1,000 such people, then, in the aggregate, they would be willing to pay their fair share of $10,000,000 to reduce their individual risks of death by 0.1%, thereby lowering the expected number of deaths in their population by 1.)

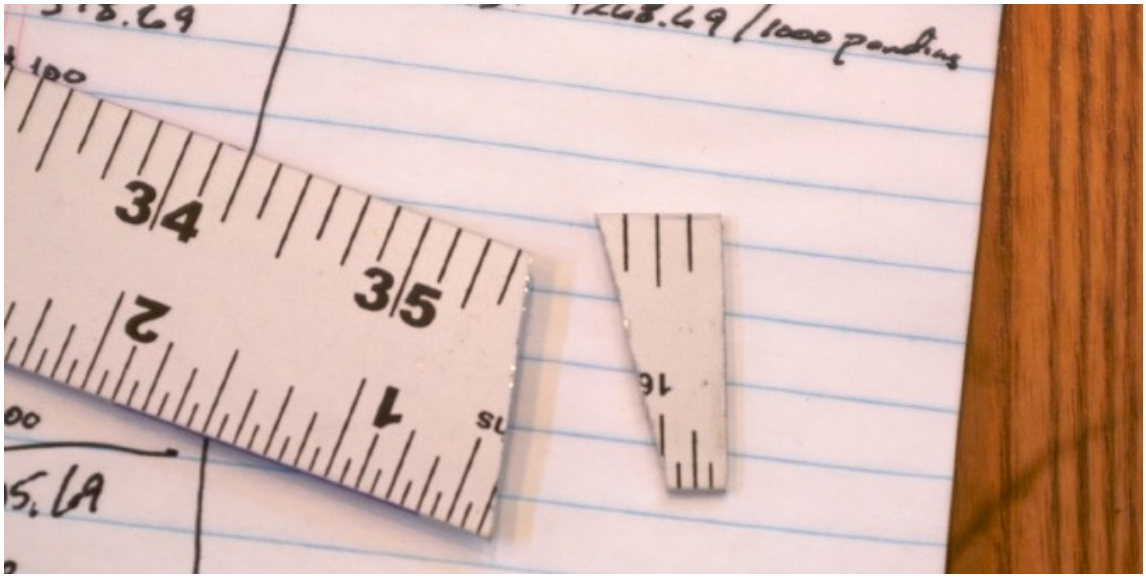

Photo: JB via a CC BY NC 2.0 licence

The value of a statistical life derived from an individual's preferences will depend on their income and wealth, which can have unacceptable consequences for benefit-cost analysis. In particular, the fact that 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 the interests of the well-off will count for more. By using a single value of a statistical life, such as the population average, this problem is avoided.

A new one arises, however, because such an average assigns the same value to every person's life saved, independently of their age. But death is generally a more serious loss when it occurs earlier in life. Reasoning in terms of life *years* preserved rather than *lives* saved is therefore more sensible. This can be done by using the Value of a Statistical Life Year (VSLY) measure. This is obtained by dividing the average value of a statistical life of the population by the average life expectancy remaining. The value of saving the life of someone in any particular age cohort is then given by the product of the value of a statistical life year and the life expectancy remaining for the cohort.

Date originally posted: 2020-08-06

Permalink: https://blogs.lse.ac.uk/covid19/2020/08/06/lives-v-livelihoods-part-1-how-can-we-measure-the-value-of-a-life/

Blog homepage: https://blogs.lse.ac.uk/covid19/

Naturally, estimates of the value of statistical lives and life years will depend on a country's income per capita. But they also vary considerably between different agents, even for a single country. For example, in the USA, the typical value of a statistical life used by government agencies is around $10,000,000, and for a life year a little over $300,000. However, the World Health Organization has suggested that interventions that generate a year in full health for more than 3 times per capita income are likely not worth the cost. 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 healthy life year, which is between a third and half of the just-mentioned estimates for the value of a statistical life year. Similar variations exist in other countries' assessments. 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. So it is critical to pay attention to the justification of any particular choice.

The two measures adopted in benefit-cost analysis face the same dilemma. If one uses individual-specific values, then the lives and (quality-adjusted) life-years of the well-off are judged to be more valuable than those of the poorest. But if, in order to avoid this inequity, one uses population averages, then the impact of policies on non-average individuals will be assessed in a way that need not correspond to their interests.

For example, a policy which imposed a loss in income just shy of $10,000 on a poor person in order to reduce their chance of death from COVID-19 by 0.1% would appear to produce a net expected benefit to this person if we used a population-average value of a statistical life of $10,000,000. But this person might reasonably judge that, for them, the risk reduction is not worth the income lost.

Social welfare analysis

Social welfare analysis avoids both the inequity of putting an individual-specific monetary valuation on life and the inefficiency of population-average valuations, by focusing directly on each individual's wellbeing, assessed in terms of their health and wealth. Since the wellbeing of everyone, rich or poor, counts equally, there is no bias towards the rich. And since health and wealth are combined, ideally in a way that suits each individual, only policies that promote each person's wellbeing will be favoured, which solves the inefficiency problem.

Policies have different impacts on wellbeing: a policy may protect the old and vulnerable, for example, but impose substantial losses on the young poor who are at risk of unemployment. We therefore need a measure of social wellbeing, called a social welfare function, to weigh these different impacts. A commonly used social welfare function is the utilitarian one, which assigns to each set of individual wellbeing values the average of the values. This way of aggregating individuals' wellbeing is indifferent to whether a given increment in wellbeing accrues to a well-off or badly-off person. But it is commonly argued that it is more important to increase the wellbeing of the worse-off, because the improvement comes to those who are in greatest need or because such improvements reduce inequality. This problem can be addressed by using social welfare functions that give extra weight to improvements in the wellbeing of the worse-off.

The choice of social welfare function is fundamentally an ethical one. An important advantage of social welfare analysis is that it allows for this choice to be made explicitly and transparently.

This post represents the views of the authors and not those of the COVID-19 blog, nor LSE. It draws on a policy brief for the G20 policy and advice network, Think20, and is the first of two posts looking at how to measure the impact of COVID-19 containment policies: the second is here.

Date originally posted: 2020-08-06

Permalink: https://blogs.lse.ac.uk/covid19/2020/08/06/lives-v-livelihoods-part-1-how-can-we-measure-the-value-of-a-life/

Blog homepage: https://blogs.lse.ac.uk/covid19/