How Economists Calculate The Costs And Benefits Of COVID-19 Lockdowns

Chris Conover

March 30, 2020

There is a huge public debate whether the economic costs of actions designed to arrest the spread of COVID-19 are worth the potential health benefits achieved.

Literally trillions of dollars in lost economic output and uncounted lives hang in the balance. No rational discussion of this weighty issue is possible without first having a hard-nosed discussion of the dollar value of saving the lives of COVID-19 patients.

This post will focus on the well-established methods that health economists have devised to answer this question.

Why Set a Price on Human Life if it is "Priceless"?

Attaching a price to human life will strike some readers as uncomfortable, if not offensive. Indeed, Governor Andrew Cuomo recently tweeted "you cannot put a value on a human life." And he is right in important and fundamental ways. But when it comes to public policy, it quickly becomes impossible to think that way.

For example, if we said that it was worth spending a trillion dollars per person to prevent unnecessary deaths, we'd run out of money in a day.

Consequently, when making regulations, policymakers use what's called the Value of a Statistical Life (VSL) to set an upper bound on how much costs regulations can impose on Americans in order to save lives. If policymakers assigned an infinite economic value to each life, we would spare no expense (and be fearless in imposing any inconvenience) if we simply could demonstrate it would save a single life. In reality, we allow Americans to take many risks—driving 70 mph on interstate highways, for example—because the societal burden of eliminating those risks cannot be justified even though we know lives are unavoidably lost as a consequence.

How Do Health Economists Calculate the Value of Preventing Deaths?

There are two fundamental approaches to this problem: one uses the value of a statistical life (VSL) to determine the appropriate level of health, safety and environmental regulation. The second uses the value of quality-adjusted life years (QALYs) to evaluate the cost-effectiveness of clinical medicine, such as diagnostic procedures, surgeries or medications.

VSL. The standard approach to estimating the value of a life for purposes of health-related regulation starts with the Value of a Statistical Life. VSL values generally are obtained by examining how people actually behave (e.g., job choices involving mortality risk, decisions to wear seatbelts) in the workplace or in making consumer purchases that could affect health (e.g., smoke detectors).

For example, if 5,000 workers are paid $2,000 more a year to work in a risky job that will result in one added worker dying (relative to a non-risky job alternative), then we infer that the VSL for these workers is $10 million. That is, as a group, they are willing to be paid $10 million a year to gamble on a slightly elevated risk that statistically is likely to kill one of them (we do not know which one).

It is difficult to imagine any worker agreeing to be paid $2,000 more in exchange for knowing for certain they would die. But that is not what is happening here. Each worker is instead being paid extra to take a 1 in 5,000 gamble of dying in the year ahead. This is not much different than risks we all willingly take every day, whether it be driving a little over the speed limit or even just walking across the street (reportedly a 1 in 300 million risk of death).

But the values obtained in VSL studies obviously will vary based on worker characteristics: 20-somethings might have a much higher willingness to gamble than older workers, for example. Even if that were not true, it would seem odd to attach the identical value of life to someone age 20 who has perhaps 60 more years of life expectancy at stake than to someone age 65 who might only have one quarter of that time left on earth. Therefore, it is not uncommon to convert a VSL into the Value of a Statistical Life Year (VSLY) by dividing the VSL by the average remaining years of life expectancy among the group being studied. The VSLY can then be used to derive a VSL that varies systematically by age, assigning a much larger value to children than elderly, for example.

QALYs. One criticism of VSL/VSLYs is that while they can adequately account for quantity of life, they do not take into account quality of life. Most people would not value a year of life spent bed-ridden as equivalent to a year of life in excellent health. QALYs are a way of allowing the value of life to vary based on subjective assessments of the quality of life. Over decades, a small cottage industry has developed around methods to assess how Americans view different states of health and what value (measured in what are called utilities) to attach to them. So if 1 represents ideal health and 0 equals death, these studies elicit preferences about what value to attach to various states of health. As one example, patients with end-stage renal disease were found in one study to value having to undergo regular dialysis at 0.75, i.e., a 25 percent reduction in the value of life relative to being in optimal health.

In the U.S., the value of one QALY (i.e., full year of optimal health) is conventionally valued at anywhere from $100,000 to $150,000. The Institute for Clinical and Economic Review (ICER), which does cost-effectiveness assessments for clinical procedures (e.g., preventive services, surgeries), diagnostic equipment and medications routinely reports results using QALY values of $50,000, $100,000, $150,000 and $200,000 as a form of sensitivity analysis (this type of fine-grained analysis might allow an insurer to routinely cover anything that passes the

$50,000/QALY threshold, but opt to impose some sort of cost-sharing on covered benefits for procedures that only can meet the $200,000/QALY threshold).

Some have raised concerns that taking into account quality of life might bias policy decisions unfavorably against people with disabilities or who otherwise are in poor health. That is, any sort of quality of life adjustment is going to automatically assign a higher value to someone who is healthy compared to someone with the identical life expectancy who is somehow incapacitated. For this reason, ICER also reports results using an Equal-Value Life Year Gained (evLYG) method, which assigns an identical value to every year of added life expectancy regardless of the health status of the patients benefiting from a given drug or procedure and without any discounting of future years.

However, the conventional practice, as articulated by the Second Panel on Cost-Effectiveness in Health and Medicine, is to use QALYs rather than raw years of remaining life expectancy and further to discount those added years of life expectancy by 3% annually. With discounting, a year of life gained in ten years would be weighed about 25% less than a year of life gained today [1]. This has the effect of shrinking the difference between children and old people in terms of their calculated value of life. Interested readers can see here a detailed explanation for why discounting life years makes sense. Even though I am a policy analyst by training—not a card-carrying health economist—I think discounting is a highly sensible idea and believe that the value of life estimates derived from using discounted QALYs yield far more defensible numbers than the VSL estimates. See the end of this post for further discussion.

The only residual question is what discount rate to use. ICER and the Second Panel continue to rely on 3% as most reflective of the social discount rate. However, a recent Mercatus paper offers what I view as a compelling case for using 7% rather than 3% as the social discount rate when evaluating public policies. For obvious reasons, the values produced using a 7% discount rate are notably lower than those using 3%, which the estimates for children, for example, being cut in half.

How Quality-Adjusted Life Year Calculations Vary by Age

As shown in Fig. 1, each of these well-established methods produces widely disparate results on the value of a human life at various ages. All are consistent in assigning a much higher value to children than the "oldest-old" adults. But for children age 0 to 19, there literally is a tenfold difference in the value calculated using the VSLY method and the value calculated using QALYs and a 7 percent discount rate. The differences across methods are straight-forward to explain:

The VSLY method understandably produces the highest results because it starts with an implicit value per added year of life of $311,194 (see Methods below). In contrast, all the QALY methods using $100,000 per quality-adjusted life year as the starting point for calculation.

Some might argue that using $100,000 to value each QALY is arbitrary. Yet as a practical matter it is in widespread use by ICER and health economists who do cost-effectiveness analysis in medicine and health. Indeed, Center for Evaluation of Value and Risks in Health at Tufts University maintains an on-line registry of literally thousands of studies making use of this methodology.

Medicare Spends >$121,000 per QALY on Dialysis. As well, consider that the Medicare program in the U.S. provides kidney dialysis at an average annual cost of $91,300 [2]. The average quality of life for a patient needing regular kidney dialysis is 0.75 [3]. Given that these patients would die without dialysis, this is good evidence that Americans are willing to spend at least $121,711 (i.e., $91.300/.75) per QALY keeping their fellow citizens alive.

$100,000/QALY is Probably A Conservative Figure for the U.S. I used $100,000 in my analysis since that made it easy for readers to cut my figures in half if they thought $50,000 more appropriate (for comparison, the English National Health System uses an upper threshold of roughly 30,000 pounds sterling per QALY, or roughly $37,000 in U.S. dollars) or double the amount if they wished to use ICER's upper bound figure of $200,000. Yet even the very generous latter adjustment would not bring the cost/QALY figures up to the same level as VSLs.

The Equal-Value Life Year Gain method produces slightly higher figures than using QALYs with a zero percent discount rate simply because each year of life expectancy counts at its full value rather than being adjusted downward based on average quality of life for patients in a given age category.

A larger discount rate will produce smaller value-of-life estimates simply because a smaller weight is being attached to future years. However, the difference between using a 3% and 7% is most pronounced for children than for the most elderly seniors simply because the discounting formula shrinks the value of years very far in the future (e.g., five decades) much more than if the added years of life expectancy are more proximate (e.g., five years).

What is the Economic Value of Preventing COVID-19 Deaths?

Last week the Centers for Disease Control and Prevention released the most recent statistics regarding the first wave of COVID-19 patients (covering cases reported from February 12 through March 16). While the situation is obviously very fluid, these are the latest "official" statistics regarding the age distribution of patients dying from COVID-19. Of the 44 reported deaths, more than one third (34%) were patients age 85 or older and another 46% were between age 65 and 84.

I used this age distribution to calculate the weighted average economic value of preventing COVID-19 deaths. Not surprisingly, these averages fall much closer to the estimated value of life for 85-year-olds than for children.

My figures—especially the QALY-based numbers—are much lower than estimates recommended by others:

· University of Southern California scholars Mireille Jacobson and Tom Chang rely on a detailed paper by Aldy and Viscusi to suggest that "If the average age of those killed by Covid-19 is 60, value of statistical life estimates put the cost of each death at approximately $5 million."

· University of Michigan economist Betsey Stevenson, a former member of the Council of Economic Advisers under President Obama, suggested in a back-of-the-envelope calculation on Twitter that we use a "reasonable statistical value of a life of $7 million."

· The economists Martin S. Eichenbaum and Sergio Rebelo of Northwestern University, with Mathias Trabandt of the Free University in Berlin, in a very detailed paper examining the optimal containment policy for this pandemic, use EPA's figure of $9.3 million as their value of life.

· Cato Institute's Alex Nowrasteh uses $10 million in his calculation of the benefits and costs of COVID-19 restrictions.

Note that all of these experts rely on a VSL calculation to ascertain whether the benefits of COVID-19 restrictions exceed their considerable cost. Given that other health, safety and environmental regulations are likewise tested against a VSL standard rather than cost/QALY standard, clearly one could argue in favor of that approach on grounds of sheer consistency. However, I would argue that using a population-wide average VSL standard for assessing COVID-19 policies is inappropriate insofar as such an average does not accurately capture the age distribution of expected COVID-19 decedents.

But for consistency enthusiasts I would argue it's a little odd that in the U.S. we rely on a cost/QALY standard in determining what clinical procedures, diagnostic equipment and pharmaceuticals should be covered by public and private insurance (and in some instances how to appropriately price such covered services) even while letting regulatory agencies price human lives much more generously.

More concretely, a COVID-19 treatment that was found to save lives at a cost of $311,194/QALY would be quickly deemed "cost-ineffective" by those in our health system in charge of allocating scarce resources—whether it be insurance executives/committees charged with determining what specific treatments/medications etc. should be covered by a health plan or those in charge of allocating resources for a fixed population (e.g., prisons or the VA health system). Yet if a hypothetical federal agency were to propose some public health regulation to "flatten the curve" at a cost to the economy that was shown to be less than $311,000 per statistical life year for each of the unknown COVID-19 whose lives were estimated to be spared by such a rule, that would pass the VSL sniff test.

Of course, that fundamental policy inconsistency in how Americans handle the knotty problem of what value to assign life goes far beyond COVID-19. So perhaps it is a discussion for a calmer time when the current pandemic is in our rearview mirror.

In my next post I will play out the implications of my value of life estimates to assess the aggregate amount of economic damage that can be justified through various restrictions on the freedom of Americans to go about business as usual.

Those curious about how I arrived at all the figures used in my charts will find these details in the concluding section below.

Methodology

Lost Life Expectancy (LLE). The latest available breakdown of U.S. COVID-19 patient deaths uses seven age categories: 0-19, 20-44, 45-54, 55-64, 65-74, 75-84 and 85+. I calculated total deaths for each group by multiplying the reported total number of patients times the case fatality rate. I estimated the average age within each age cohort using its midpoint and then used years of remaining life expectancy for such an average age individual as reported in the latest U.S. Life Tables, 2017 (see Table 1). For the 85+ age group I used life expectancy at age 85.

These are period life tables, which underestimate true life expectancy since they assume today's 20 year old will experience the same death rate as today's 65-year old once they reach that age. In reality, death rates generally fall every year across all age groups (the fact that life expectancy happened to decline in the past few years is a historical anomaly not expected to persist). A cohort life table takes this into account (details here). Using Social Security Administration estimates of cohort life expectancy in 2020 (Table V.A4) and period life expectancy for the same year (Table V.A3), I calculated the following ratios for cohort vs. period life expectancy: at birth, males=1.077; females=1.064; at age 65, males=1.049; females=1.039.

I used the ratios for males to inflate the period life expectancy for each of the seven age cohorts (using the at birth ratio to multiply all period life expectancy estimates below age 65 and using the age 65 ratio to adjust all estimates for those age 65 and older).

To derive discounted cohort LLEs, I used Excel's Present Value formula, using the relevant discount rate (3% or 7%) and inserting cohort life expectancy as the Number of Periods parameter.

VSLY. I used the methodology recommended in a recent working paper done by Harvard researchers Lisa A. Robinson, James K. Hammitt and Lucy O'Keeffe for the Bill and Melinda Gates Foundation: Valuing Mortality Risk Reductions in Global Benefit-Cost Analysis. The authors recommend calculating VSL by multiplying a country's Gross National Income (GNI) per capita times 160. To obtain VSLY, they further recommend dividing VSL by the remaining years of life expectancy (undiscounted) for persons at the average age of adults in a country.

- Using World Bank figures, I estimate U.S. GNI in 2020 to be $66,712 [4]; thus VSL=$10.7 million.

- Using Census figures, I estimate the average age of U.S. adults in 2018 as 47.6 years [5].

- Using the latest available period life tables for the U.S. population in 2017, I estimated life expectancy at this age to be 34.3 years [6].

- This yields a VSLY (undiscounted) for 2020 of $311,194.

- I multiplied VSLY by undiscounted cohort LLE to obtain VSL for each age group.

Cost/QALY. For reasons described earlier, I used a standard figure of $100,000 per QALY to convert cohort LLE (i.e., years of remaining life expectancy) into a VSL for each COVID-19 decedent age group.