Appendix: Benefit-cost analysis, discounting, and value-of-life

Here, I'll provide more information about the benefit-cost methodology that underpins the research for this book, especially on discounting and the value of a statistical life.

As I outlined in Section I: The Overarching Argument, the idea of benefit-cost analysis is simple: Identify *all* the benefits and *all* the costs in dollars, divide the two numbers, and the result tells us how much benefit each dollar will deliver. But for both costs and benefits, the calculation can be quite involved. So, to better understand how benefit-cost analysis works, let's dig deeper into a particular case: The use of tablets to teach children according to their learning levels, an educational policy from Chapter 5.

Here, as with many policies in this book, finding the cost is relatively straightforward because it generally must be paid upfront. For tablets, this cost includes training teachers to use the technology. The tablet, of course, has its own cost, but it will be used over many years by multiple children, so this needs to be divided between them. It needs to be stored overnight securely so that it isn't stolen or damaged, which requires locked storage. The tablet requires electricity, which means installing solar panels or another source of energy if there isn't grid power. In low-income countries, more classrooms also need to be built to allow the use of tablets indoors and to accommodate more children. All these costs have to be assessed and put into one currency, typically US dollars, and into a common format, here the cost per student per year. It takes some effort, but in principle, it is straightforward to calculate.

It's trickier to estimate the benefits because they will arrive across many decades. The main upside of a better-taught child is that she will be more productive in adulthood. Generally, economists express the value of better education in terms of a percentage increase in the wages that a student will earn because of greater productivity. But what is the wage going to be in 10 years when today's primary-school children start working? And what will it be in half a century when they retire?

This is not a question of calculating inflation. While the nominal dollar amount (or any other currency) will have inflated over the decades, economists use inflation-adjusted dollars. So, in all of these analyses, inflation plays no part.

As time passes, countries generally get richer on a per-person basis, and that means salaries go up. But by how much over the next half-century? This is important to establish to work out the salary boost from better education.

Wages over the 21st century

The United Nations has a set of long-term economic predictions called the Shared Socioeconomic Pathways, which maps out population and GDP for the rest of the century for regions and countries. Across all 12 papers, we use the central, so-called middle-of-the-road scenario, "SSP2", which also fits with the long-term economic predictions of the OECD and other economic evidence. This scenario gives us a likely development of GDP per person in the poorer half of the world across the century, which is standardized across all the research for this book.

In the rich world, labor income is approximately equal to GDP per capita. It sounds obvious, but it is actually somewhat surprising because part of the GDP goes to capital, but old and young people don't have labor income, so the two facts end up canceling out.

In the poorer world, because there are many more young people and less capital, typical labor income can be up to twice the GDP per person. We trace out a path of likely income growth for individual countries and regions across time, starting at their historical share of GDP and ending with income being equal to GDP per person when they become sufficiently rich.

This gives us the income of the pupil across the rest of her working life. If she learns more in school, the economic evidence shows she will see her salary increase by a certain fraction across her work life, which we can now also estimate.

Because these same assumptions are used across all the studies, comparisons between cost-benefit rates are going to be more accurate, even if actual future growth is a bit different than projected, with all errors biased in the same direction.

But this still doesn't deal with the fundamental fact that while almost all the costs take place now, the benefits are a slow drip through the next half-century. We must make these comparable by calculating the present-day value of both the costs and the benefits. To do that, economists use *discounting*.

Discounting the future

Even adjusted for inflation, a promise of a hundred dollars in a year is worth less than getting a hundred dollars today. So, any policy benefits that arrive next year have to be adjusted with a discount rate to account for that drop in value.

This might seem counterintuitive: Whether today or in a year, it's the same amount of money. But there are four factors to consider. First, waiting introduces risk. If you accept my hundred-dollar bill now, you're sure to get it. In a year, I might have forgotten our deal, we might have fallen out, I could have gone bankrupt, or either of us could have died.

Second, waiting a year means losing out on investing the hundred dollars now so that you'll have more in the bank by this time next year.

Third, if you live in a fast-growing economy, you will likely be richer next year. If making $100 for a job feels right today, by next year, $100 will likely feel too little.

Fourth, we're all just impatient. We would simply rather have $100 of opportunities *now* rather than next year.

A discount rate tries to take into account all these facts and tell us how much $100 in a year is worth now. In rich countries, the discount rate is typically somewhat lower, and in poorer countries, it is somewhat higher.

One recent study found the average discount rate across rich countries was 4.2%. In general, research shows rates between 3% and 5%. In contrast, poorer countries' discount rates range from 6–12%, with an average of 8.9%.

This means that $100 one year from now in a rich country is worth about $96 now (at 4%), whereas it is worth just $92.60 in poorer countries (at 8%).

This does not just reflect an interest rate but also the other considerations I outlined above. In a poorer country, there are more risks—for example, it is harder to enforce contracts. The economy typically grows faster in a poorer country, so next year, you'll be richer by a larger margin (albeit off a lower base).

For many people in the poorer world, concerns for the future are strongly overshadowed by immediate and urgent concerns about illness, education, housing, or putting food on the table. Getting money today would enable you to stop current problems from getting worse and stave off new ones. If you have to wait and forgo food or medical care, you might suffer serious consequences because of that delay. It would be far more valuable to you to receive $100 *now* than in a year.

Things are different if you live a relatively comfortable life in a rich country where malnutrition, basic education, and deadly, untreated illness are not a concern. Waiting to get money in a year carries costs in terms of lost investment returns, and it's possible it won't really arrive. But the loss is much less serious.

The World Bank's internal guidelines suggest that a discount rate should be about double a country's projected GDP growth per person. The growth estimates for the coming decades of the standardized SSP2 scenario for the poorer half of the world suggest a discount rate of 7.8%. Looking at this together with the $8.9% average discussed above, we chose an 8% discount rate for this project. Again, since all evaluated projects use the same discount rate, this means even if the 'real' discount rate is somewhat different, results will all be biased in the same direction.

Let's come back to the example we were looking at. With the 8% discount rate, the education economists can estimate the present-day value of the future income increases for primary school pupils (just as all the other economists in this project can do for their proposed investments). They simply discount the income increase from each year back to today, using an 8% reduction for each year. This means that the present-day value of the future income stream is *much* lower (about one-tenth) than simply summing all the additional benefits over the next half-century for all the reasons mentioned above.

Incidentally, the researchers also use a bit of discounting on the cost side because the cost of the tablets is spread out over four years, solar panels over ten years, and classrooms over 20 years.

With all of these estimates, discounting allows us to find the present-day cost per child and the present-day benefit per child of the intervention.

The value of a statistical life

One final—and crucial—issue is how to value the avoidance of death and disease. This isn't relevant in education, but it comes up in many other programs we have looked at for this book, like reducing malaria deaths with insecticide-treated mosquito nets. The research shows that about $6,700 spent on more nets and information will avoid one malaria death. The price comes from estimating the purchase cost and distribution costs of the nets (treated with cheap or more expensive insecticides) along with information campaigns needed to achieve their proper use. So, what is the benefit of avoiding one person dying from malaria?

If you ask anyone about the value of human life, the innate response will be that it is incalculably large or even infinite. This is, of course, especially true if it is your own life or the life of someone you know and love. However, even the option to save a complete stranger, say an infant in peril, would elicit an incredibly high willingness to pay.

But saving a specific, named individual is not typically the choice faced by societies, or indeed, in the 12 phenomenal policies described in this book. The $6,700 of mosquito nets will not save *this* boy or *that* woman, but rather it will save one faceless, statistical person somewhere in some village in some nation.

How much are we willing to pay for that? Most people's immediate and well-mannered response would again be to insist on an incalculably large or even infinite amount. However, reality belies that we manifestly don't spend near-infinite amounts of resources to save anonymous lives.

This is not because we are fundamentally uncaring but because there are many different demands on our resources. We need to pay for our own housing, food, and our kids' education, and we want money left for entertainment and the odd luxury. This means we can spend some (but not all) our money on saving lives through charity.

Let us first look at our willingness to save lives in rich countries. Here, societies spend some (but not all) money to help save some (but not all) lives.

The most obvious example is if we look at traffic death prevention, as I briefly mentioned in Chapter 3. Road accidents kill 41,000 people in the USA and 35,000 in the EU every year (with an annual global traffic death toll of 1.3 million people). Yet, even rich countries don't spend everything they can to avoid traffic deaths.

We know that many transport policies dramatically reduce deaths. One good example is highway design, where median barriers can save lives by preventing cars from accidentally crossing the median. The famous Golden Gate Bridge installed median barriers in 2015. Before that, every five years saw 14 head-on collisions and two deaths, while the first five-year period after the installation has seen zero of either.

However, the barriers also had a high cost of $30 million. That cost explains why societies will install these median barriers in high-traffic, dangerous zones where the investment can save many lives over decades but will not construct median barriers everywhere, especially not in low-traffic, low-accident areas. At some point, every society decides that the additional cost of more barriers in ever-less dangerous zones is not worth the slight additional probability of a saved life.

Such a decision implicitly reveals society's willingness to pay to save an extra life. Societies will eagerly pay thousands of dollars to avoid the loss of a single life, but they are unlikely to spend billions to save a single life. Reams of research have looked across many areas of spending dedicated to saving human life. The result is the estimate that for America, the likely cut-off point for being willing to save an extra life is about $10 million.

Interestingly, we, as individuals, also reveal a similar implicit valuation of our own lives. Of course, if we are asked for our willingness to be compensated for our certain death, under ordinary circumstances, no amount of money would be enough. But most people will be willing to cross a road to get, say, a Snickers bar, although even a careful road-crossing involves a non-zero risk of dying. Implicitly, we're willing to suffer a slight risk of death for the benefit of candy.

We see this more formally in employment, where more risky jobs like mining or logging are compensated with additional hazard pay (with endearing honesty called "danger money" in England). The additional hazard pay compared to the increased death risk implicitly reveals how much more money has to be paid to a lot of workers for each extra death caused. It turns out that the amount similarly is around $10 million in the USA.

These two data points suggest that if we could save one life in the USA, it would not be unreasonable to value that at about $10 million. Indeed, that is what most rich world benefit-cost analyses will use for avoided deaths.

It bears repeating that our willingness to spend some but not all resources to save lives is not an indication that we are monstrously self-absorbed or uncaring. It simply reflects the reality that we live in a world with limited resources and infinite spending opportunities.

In the poorer half of the world, typically, there are many more risks, and resources are much scarcer. This means that societies and individuals have much more constrained opportunities to pay to avoid additional deaths. The sights you might see in developing country environments of dangerously overcrowded trains or trucks packed with people emphasize this trade-off. Any individual traveler will, of course, strongly prefer to arrive alive, but buying a safer trip can easily be prohibitively expensive, obliterating the chance to pay for other compelling outlays.

This contrast is most blatantly obvious when we note that rich countries spent $5,600 in 2019 on healthcare for each inhabitant, while low- and lower-middle-income countries spent just $86 per person, or 65 times less. Clearly, poorer nations are indicating that given their current resources and other responsibilities, they are not able to have as high a cutoff point for saving lives as rich nations. This suggests that the value put on saving a life in the poorer world is lower.

A common first reaction is that the rich world ought to step in. However, there is no serious political conversation about transferring any substantial part of rich country spending on healthcare to the poorer world. The entire healthcare spending for the poorer half of the world is about $360 billion annually. High-income countries spend 21 times more ($7.5 trillion) on their much smaller population. Few people in the rich world would be willing to consider an equalization in health spending, resulting in $5.6 trillion in annual transfers from the rich to the poor (let alone the $24 trillion annual transfer to get everyone at the rich world's health level).

There is a small number of studies in low- and lower-middle-income countries that attempt to estimate the statistical value of a life, both through government expenditure and through hazard pay for more dangerous work. In our work, we use the analysis from a highly cited paper that uses the best available evidence to link GDP per person to the likely value of statistical life. In 2023, across low- and lower-middle-income countries, this value is 54 times the GDP per person, or $128,000. We also take into account that as economic growth makes people richer, this value will increase over time.

We set this value of $128,000 across all studies, meaning that the results are all comparable. We also set that value across all nations in the poorer half of the world. It should be noted that this disregards the reality that within this group of nations, some poorer countries could not afford to value an additional life as highly as our analyses does, while some richer countries could afford to save more lives.

There is one more technicality that needs to be addressed. Some benefit-cost analyses simply use the value we have identified above as a value for *all* saved lives. However, this ends up meaning that saving an infant or a young parent who will likely go on to live for many decades is valued as equivalent to saving a very old person who has just a few years or even a few months left. Analyses show that most people consider this unreasonable.

Instead, this book estimates the value of each saved life-year as worth $4,300. That means saving an old person from chronic heart disease for, say, six years is worth 6 x $4,300. With an average life expectancy of 69, saving a one-year-old baby with a rotavirus vaccine is worth 68 x $4,300.

Across all studies, for 2023, we use the value of one avoided life-year or the so-called Value of Statistical Life Year (VSLY) set at $4,300.

The usefulness of consistent benefit-cost analysis

As is evident in the individual chapters and even more so in the academic papers, executing a benefit-cost analysis is conceptually simple but involves a great many considerations. Here, we've touched on some of the most significant factors.

Ideally, such benefit-cost analyses can show us how well our next dollar can be spent. Certainly, because we use the same assumptions across all the studies, we can be quite confident that these studies are comparable.

Nevertheless, it would be folly to believe that these complicated models, often contingent on far-off futures and significant changes in policy, give the absolutely correct answer in a simple $33-back-on-the-dollar number. There is no way all costs and all benefits have been correctly incorporated across decades. They are our best estimates, not truths.

Since benefits are often harder to calculate than costs, benefits are typically estimated in one or a few ways. This means the benefit-cost ratio could easily be underestimated. For instance, within education, economists typically estimate the benefits solely through higher productivity, which we then estimate with a fractional increase in income over working years. But surely, more education also has other benefits in addition to increasing productivity. Better-educated children may well grow up to be more free-thinking, better at problem-solving in their daily lives, and perhaps better citizens. These benefits are not captured in the benefit-cost ratio.

Yet, benefit-cost analyses give us a unique opportunity to systematically identify phenomenally good policies. Certainly, compared to the way current policy is decided (remember the SDG targets, many of which were adopted simply because they *sound* good), such an approach is a step up. Moreover, it is likely that while the benefit-cost ratio is not absolutely correct, it is a much better approximation of where we can do the most good.

Finally, since this book only includes the most phenomenal policies, even if we have underestimated the benefits, this simply makes the argument for these policies stronger.

It may still seem harsh to discount the future with a fixed percentage to measure the value of more learning. It clearly comes across as crass to estimate the value of an additional saved life-year in plain dollars. Yet, not doing so doesn't make our hard choices go away. It simply makes them less transparent.

So, while these calculations may feel awkward, they are academically grounded and will help us be more honest about our hard trade-offs. Certainly, it is not wrong to put a price on life and a price on the future if it helps us to save more lives and make the future better.