Chapter 3

Considering cost-effectiveness: the moral perspective

Achieving good value with scarce resources is a substantial moral issue for global health. This claim may be surprising to some, since conversations on the ethics of global health often focus on moral concerns about justice, fairness, and freedom. But outcomes and consequences are also of central moral importance in setting priorities. In this chapter we explore the moral relevance of cost-effectiveness, a major tool for capturing the relationship between resources and outcomes, by illustrating what is lost in moral terms for global health when cost-effectiveness is ignored.

The cost-effectiveness landscape in global health

The moral importance of cost-effectiveness is illustrated by its substantial variation between different interventions. A simplified example will help clarify how this becomes a moral consideration. Suppose we have a $40,000 budget that we can spend as we wish on the health condition of blindness in an African country. One option is to provide blind people with guide dogs to help them overcome their disability. This costs approximately $40,000 per person due to the training required for the dog and its recipient.[i] Another option is to pay for surgeries to reverse the effects of trachoma. This costs less than $20 per patient cured, with the cost of the surgery itself less than $10.[ii] Many other options exist, but for simplicity, let us just consider these two.

We could thus use the entire budget to provide a single guide dog, helping one person overcome the challenges of blindness, or we could use it to cure more than 2,000 people of blindness. By allocating all the resources to purchase a guide dog for one person we are ignoring the legitimate claims of 2,000 other people. Some would reframe the choice by arguing that the second option is more than 2,000 times better than the first, or, even more starkly, that the first option squanders about 99.95 percent of the value that could have been produced.

This example illustrates the basic point, but it is also unrealistic in at least two ways. First, we often have an expansive spectrum of options—with greater variation in scope and complexity than just the two previously considered. Second, and more importantly, the class of interventions under consideration is often broad enough that it is difficult to make direct "apples to apples" comparisons between the effects of two interventions.

Health economists have an answer to the second issue. They use measures that transform all health benefits into a single metric, thus allowing for direct comparisons between interventions that are aimed at different threats to health and that affect different health outcomes. The standard measure in global health is the disability-adjusted life year (DALY), which gauges the disvalue of health conditions by the number of years of life lost due to the condition plus the number of years lived with disability multiplied by a number representing the severity of the disability. For example, a condition that caused one to die 5 years prematurely and to live the past 10 years with deafness would be valued as 5 + (10 × 33.3 percent) = 8.33 DALYs.

There are a number of methodological complications and choices for calculating DALYs, which give rise to a number of different versions of DALYs and the closely related units—quality-adjusted life years. Chief among these choices is the question of the size of

i. Guide Dogs of America estimates $19,000 for the training of the dog. When the cost of training the recipient to use the dog is included, the cost doubles to $38,000. Other guide dog providers give similar estimates, for example, Seeing Eye estimates a total of $50,000 per person/dog partnership, while Guiding Eyes for the Blind estimates a total of $40,000.

ii. Cook et al. (2006) estimate their figure to be $7.14 per surgery, with a 77 percent cure rate.

The first draft of this chapter was written by Toby Ord, Balliol College, Oxford University.

the weightings representing how bad it is on average to suffer from a particular disability. There are also considerations on discount rates and age weightings.

These methodological choices have ethical implications. People disagree on the appropriateness of particular disability weights, or about the method for eliciting these weights, or about discounting health benefits, or weighting benefits depending on the age of the recipients, or whether other issues such as equal moral consideration need to be factored in. Members of the Working Group and the author of this chapter have many of the same concerns and agree that DALYs should be considered only as a rough measure of the disvalue of different conditions.

In most cases, however, different reasonable choices on these parameters change the number of DALYs due to a condition by a few percent or by as much as a factor of two. By contrast, the difference in cost-effectiveness between interventions is often a factor of 100 or more. Thus, even a rough measure of DALYs saved can supply information of critical moral importance in informing key comparisons.

The two concerns raised in the hypothetical blindness case can be addressed by looking at a real-world example of funding the prevention or treatment of HIV/AIDS. Let us consider five intervention types: surgical treatment for Kaposi's sarcoma (an AIDS defining illness), antiretroviral therapy to fight the virus in infected people, prevention of transmission of HIV from mother to child during pregnancy, condom distribution to prevent transmission more generally, and education for high-risk groups such as sex workers. It is initially very unclear which of these interventions would be best to fund, and one might assume that they are roughly equal in importance. However, the most comprehensive compendium on cost-effectiveness in global health, *Disease Control Priorities in Developing Countries,* 2nd edition (2006), hereafter *DCP2*, lists their estimated cost-effectiveness as follows (figure 3.1).

Note the wide discrepancies between the effectiveness obtained for the same amount of money. Treatment for Kaposi's sarcoma cannot be seen on the chart at this scale, but that says more about the other interventions being good than about this treatment being bad: treatment for Kaposi's sarcoma is often considered cost-effective in high-income countries. But antiretroviral therapy is estimated to be 50 times as effective as treatment for Kaposi's sarcoma, prevention of transmission during pregnancy is 5 times as effective, condom distribution is about twice as effective, and education for high-risk groups is again about twice as effective. In total the best of these interventions is estimated to be 1,400 times as cost-effective as the least good, or more than 1,400 times better than it would need to be in order to be funded in rich countries.

Figure 3.1
Cost-effectiveness of HIV-related health services

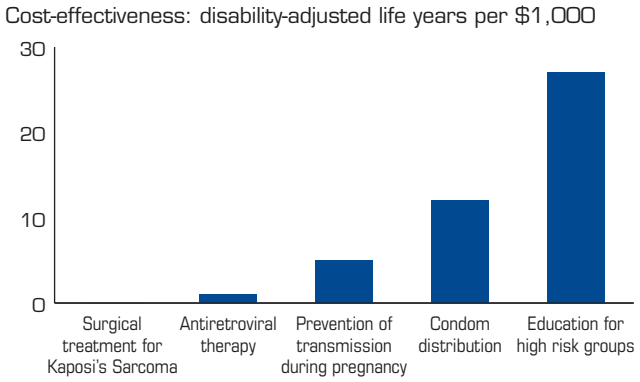

Source: Jamison et al. (2006).

This discrepancy becomes even larger if we make comparisons between interventions targeted at different types of illness. *DCP2* includes cost-effectiveness estimates for 108 health interventions, arranged from least effective to most effective (figure 3.2).

This larger sample of interventions is even more disparate in cost-effectiveness. The least effective intervention analyzed is still the treatment for Kaposi's sarcoma, but some other interventions are up to 10 times more cost-effective than education for high-risk groups. In total the interventions are spread over more than four orders of magnitude, ranging from 0.02 to 300 DALYs per $1,000, with a median of 5 DALYs. Thus, moving money from the least effective intervention to the most effective would theoretically produce about 15,000 times the benefit, and even moving it from the median intervention to the most effective would produce about 60 times the benefit.

It can also be seen that due to the skewed distribution, the most effective interventions produce a disproportionate amount of the benefits. According to the *DCP2* data, if we funded all of these interventions equally, 80 percent of the benefits would be produced by the top 20 percent of the interventions. It must be noted that these

Figure 3.2
Cost effectiveness of 108 health interventions evaluated in DCP2

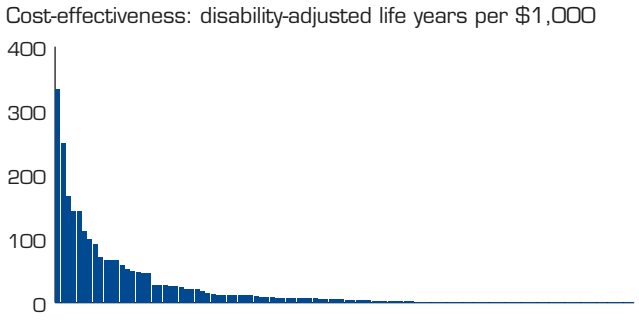

Source: Jamison et al. (2006).

Figure 3.3
Preventable deaths from immunization-preventable illness, diarrhea, malaria, smallpox, and war

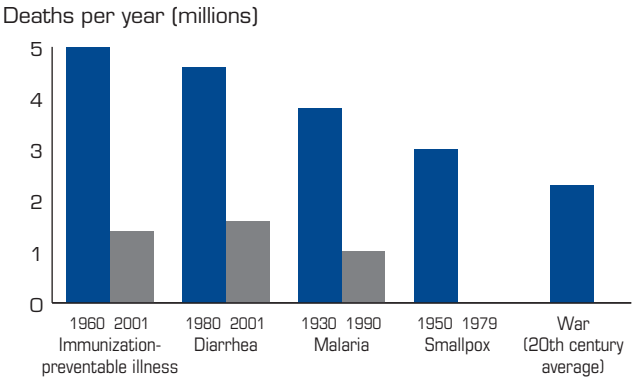

Note: Estimates for the death toll from all acts of war and genocide in the 20th century vary from about 160 million to 240 million, differing in which deaths they include (Leitenberg 2006).

Source: Jha et al. (2004).

are merely estimates of cost-effectiveness and there may be less variance between the real, underlying cost-effectiveness values. However, even if the most effective interventions are one-tenth as effective as these figures suggest and the least effective are 10 times better than they appear, there would still be a factor of 150 between them.

Moreover, there have been health interventions that are even more effective than any of those studied in the *DCP2*. For example, consider the progress that has been made on saving lives lost to immunization-preventable illness, diarrhea, malaria, and smallpox (figure 3.3).

In all cases these interventions have led to at least 2.5 million fewer deaths per year. To aid the reader in comprehending the scale of these achievements, a final bar in the graphic shows the average number of deaths per year due to war and genocide together over the 20th century (2.3 million).

Moreover, these gains have been achieved very cheaply. For instance, with smallpox, the total cost of eradication was about $400 million.[1] Since more than 100 million lives have been saved so far, this has come to less than $4 per life saved—significantly superior to all interventions in the *DCP2*. Moreover, the eradication also saved large amounts of money. About $70 million was being spent across developing countries per year in routine vaccination and treatment for smallpox, and more than $1,000 million was lost per year in reduced productivity.[2] Even just in the United States, smallpox vaccination and vigilance cost $150 million per year before eradication.[3] The eradication program thus saved an enormous number of lives per year, while saving money for both donors and recipients, paying back its entire costs every few months. It serves as an excellent proof of just how cost-effective global health can be.

The main effect of understanding the moral need of consideration of cost-effectiveness is spending our budgets to produce greater health benefits, saving many more lives and preventing or treating more disabling conditions. However, it also shows a very interesting fact about global health funding. If we can save 1,000 lives with one intervention and 10,000 with another at an equal price, then merely moving our funding from the first intervention to the second saves 9,000 lives. As such, moving funding from one intervention to a more cost-effective one can produce almost as much benefit as adding an equal amount of additional funding to the more cost-effective intervention. This is counterintuitive since it is not the case when one option is merely 10 percent or 30 percent better than another. However, when one option is 10 times or 100 times better, as is often the case in global health, redirecting funding is so important that it can be equivalent to new funding directed toward the superior intervention. In times of global austerity and shrinking budgets, it is good to know how much more can be done within current resources.

The moral case

In these examples we have seen how incredibly variable cost-effectiveness can be within global health. The least effective intervention in the HIV/AIDS case produces less than 0.1 percent of the value of the most effective, and if we are willing to look at different kinds of disease, this fraction drops to less than 0.01 percent. Ignoring cost-effectiveness thus does not mean losing 10 or 20 percent of the potential value that a health budget could have achieved, but it can easily mean losing 99 percent or more. Even choosing the median intervention can involve losing 85 percent of the potential value. In human and moral terms this can mean hundreds, thousands, or millions of people who will lose their lives due to the failure to take cost-effectiveness into account in allocating health resources. In non-life-saving contexts it means thousands or millions of people who will live with significant disabilities that could be prevented, mitigated, or cured.

In this chapter we make the case that considering cost-effectiveness is a necessary condition for making decisions on global health spending, but on its own, it is not enough. Other moral values similarly need to be factored into a decision-making framework; these include, but are not limited to, fairness, impact on unjust inequalities and systematic disadvantage, and other dimensions of well-being such as respect and self-determination, as well as proximity, financial protection, and the like.

Learning how to factor these other ethical considerations correctly into our decision making is an important and challenging problem. But it is important to recognize that we are failing at one terribly important moral imperative—how to achieve good value with limited resources—at the expense of untold human suffering, and not because we are intentionally forgoing securing more good in the interests of acting virtuously or avoiding violating people's rights.

Challenges addressed

Some people do not see cost-effectiveness as advancing ethical considerations, perhaps because its empirical methodology makes it appear more like a technical or scientific matter. This is misguided. People who decide how to spend health budgets hold the lives and well-being of many other people in their hands. They are literally making life-or-death decisions. Most decisions of this sort take dramatically insufficient account of cost-effectiveness. As a result, thousands or millions of people die who otherwise would have lived. The few are saved at the expense of the many, without any justification or compelling rationale. It is typically done out of ignorance about the significance of the cost-effectiveness landscape rather than out of prejudice, but the effects are equally serious.

Another reason people might be initially suspicious of using cost-effectiveness to guide prioritization is through confusion with cost-benefit analysis. The latter is an economic method for prioritization that involves determining the benefits for each person by how many dollars they would be willing to pay, adding these up, and then dividing by the total costs in order to produce a benefit-cost ratio in units of dollars. This method is ethically suspect as it considers benefits to wealthy people (or groups) to be worth more than comparable benefits to poorer people (or groups), since the wealthy are willing to pay more for a given benefit.

However, the cost-effectiveness discussed in this chapter is very different, a type of analysis known as cost-effectiveness analysis. This method does not convert benefits into dollars but provides a raw measure of the benefits in units such as DALYs per dollar, or lives saved per dollar. Thus this method is not biased toward interventions favored by the wealthy.

Concern may be an inevitable response to cost-effectiveness since it makes a connection between dollars and health (or even life itself). Making tradeoffs between sacred values such as life and nonsacred values such as money strikes many people as morally problematic. However, no such tradeoff is made in cost-effectiveness analysis. Instead there is a budget constraint of some fixed number of dollars and the cost-effectiveness ratios help to illustrate how much benefit could be produced if this money were spent on a given intervention —for example, saving 1,000 lives or saving 10,000 lives.

Conclusions

Ignoring cost-effectiveness altogether in assessing global health risks means losing much if not most of the value that we could potentially create. For this reason alone the cost-effectiveness of interventions needs to be considered on moral grounds. This does not simply mean implementing current interventions in the most cost-effective way possible, for the improvements to be gained within a single intervention are quite small in comparison. It also does not mean just doing retrospective measures of the cost-effectiveness of the interventions you fund as part of program evaluation. It does mean expanding the domain of interventions

under consideration to include all those whose cost-effectiveness has been established and that are currently underfunded. And, above all, it means allocating funds to interventions that correspond to all relevant moral considerations, including, as a core ethical priority, the moral value of producing the good of global health efficiently. Without a commitment to making value for money central to priority setting, we will continue to fail to honor obligations to improve the life prospects of those in need and to act justly.

Notes

1. Fenner et al. (1988).
2. Fenner et al. (1988).
3. Fenner et al. (1988).