Refraining from Assigning Values to Life and Health—Cost-Effectiveness Analysis

In conventional cost-benefit analysis, all benefits are assigned a dollar value. But, as noted in the introduction, an analyst or policymaker may be uncomfortable with assigning dollar values when benefits are human health and safety. In this case, cost-effectiveness analysis may look attractive. With cost-effectiveness analysis, analysts do not assign a dollar value to health benefits. Instead, benefits are simply a count of the adverse outcomes averted. Benefits are left in physical terms and not monetized.

Cost-effectiveness analysis is a comparison of costs with the number of physical benefits. The ratio of dollar costs to physical benefits is the cost per physical benefit. The program with the lowest cost per benefit is the most cost-effective. When comparisons are made between programs having identical types of benefits, cost-effectiveness analysis yields a cardinal ordering of alternatives. Numerical summary measures show which intervention is most and least cost-effective. For example, to study the effects of intervention strategies on heart disease, analysts can compare costs and the number of strokes averted by different programs. The program with the lowest cost-per-stroke-averted is the stroke-prevention program that is most cost-effective in the sense that dollars-per-stroke is minimized. The difference between minimum cost-effectiveness and cost-effectiveness of a particular program measures the sacrifice in efficiency associated with the program.

In this section, we show that conclusions drawn from cost-effectiveness analyses share many undesirable characteristics with conclusions drawn from conventional cost-benefit analyses, using COI or WTP to estimate benefits, but do not share many of the desirable characteristics. In particular, when analysts base their policy recommendations on cost-effectiveness analysis, their analyses will usually be subject to the influence of income and circumstance, just as with cost-benefit analysis. However, unlike cost-benefit analysis, they will give up the ability to rank diverse programs and to show whether net benefits are positive or negative. Only by luck will program rankings match those derived from cost-benefit analysis.

Background

Cost-effectiveness analysis has been used extensively to evaluate the desirability of medical technologies (Hildred and Watkins, 1996). In medical decision-making, cost-effectiveness analyses are carried out more frequently than conventional cost-benefit analyses. Elixhauser et al. (1993) indicate that about two-thirds of analyses of health-related technologies, services, and programs are cost-effectiveness analyses.

Mushkin (1979, pp. 19-20) recommends cost-effectiveness analysis over cost-benefit analysis when it is difficult to assign a price to health benefits. Garber et al. (1996, p. 28) state that it is the difficulty of carrying out cost-benefit analyses (largely the task of assigning values to health outcomes) and the discomfort associated with assigning monetary values to health outcomes that has led to rejection of cost-benefit analysis. Haddix and Shaffer (1996, p. 104) report that the public health community initially embraced cost-effectiveness over cost-benefit analysis because the former was less burdensome and less complex to execute.

There are very specific conditions under which cost-effectiveness analysis is especially useful. If an irrevocable decision has been made to take an action to prevent an adverse outcome, but no decision has been made about technique or method of action, cost-effectiveness analysis can help minimize costs. Folland, Goodman, and Stano (1993, pp. 638-9) note that cost-effectiveness analysis has been frequently applied by the Department of Defense. There, objectives can be quantified in terms like the ability to deploy forces, and analysts often are assigned goals of finding the most efficient means of achieving specified objectives. A private sector health care administrator or a government official responsible for some aspect of public health may face similar situations. A health care administrator may be compelled to offer a particular health care service. A government official may be compelled to put in place a program to reduce some particular health risk. Still, each may have the latitude to choose among programs that accomplish the mandated goal. Each can strive to receive more benefits per dollar rather than less. Cost-effectiveness analysis may assist in making such choices.

Inability To Rank Diverse Programs and To Determine Whether Net Benefits are Positive or Negative

Cost-effectiveness cannot be used to rank options when program outcomes differ. For example, cost-effectiveness does not offer a way to compare the costs and benefits of a program that provides kidney dialysis with those provided by a nutrition program. A cost-effectiveness analysis of each would yield two calculations: the cost for treating individuals whose kidneys have failed and the cost of averting a particular diet-related illness. The benefits of blood purification cannot be compared with the benefits of a more healthful diet. When health outcomes differ, program benefits are measured in different units of account and are not comparable. This situation differs from WTP or COI where both benefits and costs of all options are denominated in dollars, and programs with diverse types of health benefits can be ranked.

In addition, cost-effectiveness estimates do not, by themselves, suggest whether any of the examined programs meet the test of efficiency. Because cost-effectiveness analysis measures costs and benefits in different units, no concept of net benefits emerges. Consider, for example, two programs that could avert *Salmonella* infections. Suppose one program costs $1,000 per infection averted and the other $2,000. The cost-effectiveness calculations indicate that the first program offers the greater benefit per dollar. But, it does not reveal whether the program is actually worth the price.

A cost-benefit analysis using WTP or COI as benefits could reveal whether the most cost-effective program costs more than it is worth. Where benefits are monetized, calculating net benefits is easily accomplished (by subtracting costs from benefits), and the sign and magnitude of that calculated value indicates how desirable the program is. But cost-effectiveness analysis does not provide any means of judging whether a program offers positive net benefits. Finding the most cost-effective program is simply a different activity from evaluating benefits using WTP or COI.

Three Variants of Cost-Effectiveness Analysis

Different decisionmakers have different goals and methods for carrying out cost-effectiveness analyses. These differences are not always obvious to those who use cost-effectiveness calculations to make comparisons among programs, and consistency across studies is a problem for cost-effectiveness analysis and for interpretation of results. Gold et al. (1996) note the variety of decisionmakers implicit in cost-effectiveness analyses.

> The imprecision attached to the term "cost-effective" stems also from the variety of masters the concept serves. Purchasers of health care use the term to convey a careful assessment of the relative value of different health care services; producers of health care technologies and programs use the idea to support marketing claims; advocates for particular illnesses or constituencies use the term to garner resource investments. (p. xvii)

The characteristics of cost-effectiveness analysis will differ markedly depending on who the decisionmaker is and what the objectives are. Where analysts address societal problems using cost-effectiveness analysis, they tally a wider class of dollar costs than when choicemakers have a narrower focus, say that of health care institutions and third-party payers (Torrance et al., 1996, pp. 60-61). Activities that count as costs will differ, and the way in which health benefits are tabulated will differ, depending on which goal analysts (at least implicitly) embrace.

At least three variants of cost-effectiveness analysis can be distinguished, depending on the decisionmaker and goals. The first variant of cost-effectiveness analysis entails the ratio of program costs to a count of health benefits. In this type of analysis, costs are outlays for program administration. Benefits are typically the number of adverse outcomes averted, like cancers averted or premature fatalities avoided. From the perspective of satisfying individual preferences or of maximizing aggregate income, such measures are likely to be incomplete and of limited value in making public health decisions. The information would be useful, however, for a financial officer attempting to satisfy a policy goal at minimum budget exposure.

The Superfund site calculations carried out by Viscusi and Hamilton (1996) are examples of the first variant of cost-effectiveness analysis. They calculated cost-effectiveness by dividing *clean-up costs* by cancers averted. Their methods were sufficient to draw attention to the extraordinary costs per benefit that had been incurred. But such calculations do not suggest a means of allocating funds among public health and safety programs that would satisfy individual preferences or maximize aggregate income.

Selecting the most cost-effective program, measured with the first variant of cost-effectiveness, leads to choices that differ from those made by self-interested individuals. Cost-effectiveness tabulates health benefits as a count. This method is egalitarian in that everyone's benefits count equally, regardless of income. However, individual preferences for health benefits vary, and treating everyone alike eliminates the influence of preferences as well as income. Clearly, this is not an individual welfare measure.

The first cost-effectiveness variant is also unlikely to satisfy the objective of maximizing aggregate income. Cost-of-illness is a better tool to satisfy that goal. That is, selecting projects that maximize the difference between cost-of-illness that might be averted and mitigation costs, maximizes aggregate income. This calculation of net benefits is composed of dollar-denominated additions to income and dollar-denominated subtractions from income. Cost-effectiveness fails to maximize income because government expenditures directed toward mitigating a health hazard are only a part of dollar-denominated net benefits.

The only decisionmaker for whom this cost-effectiveness matters is one that attempts to avert the maximum number of adverse health outcomes at minimum budget exposure. Budget exposure is important to questions of government accounting and finance. But it is clearly not equivalent to aggregate income or to individual welfare.

The second variant of cost-effectiveness analysis involves replacing program costs with net costs, where net costs are the direct program costs minus the reductions in cost-of-illness resulting from each program. Haddix and Shaffer (1996) detail methods for the second cost-effectiveness variant, measuring societal costs and benefits (pp. 109-127). For questions with a societal perspective, they argue for comparing net costs with a count of adverse health outcomes averted.[27]

$$Cost\ effectiveness = \frac{Net\ cost}{Total\ adverse\ health\ outcomes\ averted} \qquad (5)$$

Jones-Lee (1994) describes net costs as capital costs minus other benefits. Haddix and Shaffer treat net costs of carrying out a program as program costs minus the change in the cost-of-illness. That is, if a program reduces risks and COI falls, the reduction in COI is attributed to the program. In that case, the real costs imposed on society by financing the program are not as large as the direct expenses for the program. To find the real cost to society, Haddix and Shaffer suggest that analysts subtract the reduction in COI attributed to the program from the direct financing costs of the program.[28]

$$Net\ cost = Cost_{Program} - Reduction\ in\ cost\ of\ illness \qquad (6)$$

[27] Haddix and Shaffer define the ratio of net cost to total adverse health outcomes averted as "average CE ratio," noting that incremental costs of various levels of programs can also be calculated. The average and incremental estimates allow decisionmakers to find an optimum level for the most cost-effective program. The importance of the incremental calculation can be seen in an often-cited cost-effectiveness study by Neuhauser and Lewicki (1975). They studied a colon cancer screening protocol. The protocol consisted of six sequential stool tests for occult blood with follow-up testing for positive results. Neuhauser and Lewicki showed that incremental cost per detected cancer increased an order of magnitude with each sequential test, reaching over $47 million for the last test (not accounting for inflation). The importance of program scale was not so well revealed by the average cost-effectiveness.

[28] Clearly, both Jones-Lee and Haddix and Shaffer describe programs financed by the public sector. An entirely different conception of costs, including regulatory compliance costs, must accompany an analysis of regulatory changes intended to protect public health.

As COI consists of direct and indirect costs, analysts have to calculate changes in both types of costs.

$$\begin{aligned}&\textit{Reduction in cost of illness} = \qquad (7)\\ &\quad \textit{Direct medical expenses averted}\\ &\quad + \textit{Value of productivity losses averted}\end{aligned}$$

Haddix and Shaffer define direct costs as the costs of diagnosis and treatment associated with cases of the health problem averted, as well as the cost of unintended side effects of treatments. For indirect costs they recommend calculating productivity losses (human capital costs). In this variant of cost-effectiveness analysis, analysts assign values to people equal to their productivity. As a result, these cost-effectiveness estimates functionally depend on income and circumstance. Thus, everything else equal, a program that benefits only men will be more cost-effective than a program benefitting women because men's wages are generally higher than women's wages.[29]

Garber et al. (1996, p. 51 footnote 11) note that calculating cost effectiveness by subtracting the change in COI from program costs is identical to a cost-benefit analysis in which COI serves as program benefits. From a purely mechanical perspective, exactly the same quantities are calculated as would be in a cost-benefit analysis in which the change in the COI represents benefits. With cost-benefit analysis, net benefits could be calculated as the change in COI minus the program costs. Net benefits would be, in this case, exactly what Haddix and Shaffer describe as net costs, after accounting for sign differences. The principal difference between the two calculations is that cost-effectiveness analysis divides net cost (or net benefits) by a count of adverse outcomes averted. This last calculation means that cost-effectiveness is a per capita measure rather than a total. In this case, cost-effectiveness analysis appears neither less complex nor less burdensome to carry out than a cost-benefit analysis.

Like COI, the implicit goal of the second variant of cost-effectiveness analysis is national or aggregate income maximization. This variant of cost-effectiveness analysis would be useful to an aggregate income-maximizing central planner.

The third variant of cost-effectiveness analysis is described by Garber et al. (1996). They suggest that program costs ought to include individual opportunity costs, such as the value individuals place on time lost to morbidity. With this variant, instead of counting adverse outcomes averted, analysts tabulate annual program-induced health changes over an individual's lifetime. Each of the annual changes is expressed on a zero-to-one scale, with zero representing no change in health over a year and one representing an added year of life in robust health. Intermediate health increments are valued in the 0-1 interval, converting all benefits to a common unit of account, the Quality-Adjusted Life Year (QALY).[30] For example, a new vaccine for a quickly fatal illness would prevent some premature deaths. The benefits of making the vaccine widely available could be calculated as the discounted sum of life years added. On the other hand, a palliative treatment might not add any years of life, but make years of illness more pleasant. The years of improved well-being could be evaluated as the fractional equivalent of a year of good health. The discounted sum of these fractions can be counted as QALYs. Adding all individual's QALYs yields a measure of program benefits. This variant of cost-effectiveness analysis is often called cost-utility analysis

QALYs translate all health consequences into a common unit of account for health benefits. Dividing program costs by QALYs yields a price per QALY. Thus, quite diverse programs can be ranked. However, as benefits and costs are in different units of account, no net benefit concept emerges. Like the other cost-effectiveness variants, this version does not suggest whether any program is worthwhile.

[29] Examining the period 1967 through 1984, Berndt (1991) observed that median weekly earnings of females were about 62 percent of those of males. The ratio has been rising, and in 1996 reached 75 percent. The ratio of earnings of blacks to whites was 76 percent in 1996 (calculated from statistics in U.S. Dept. of Labor, Bureau of Labor Statistics, *Employment and Earnings*, Jan. 1997, p. 204).

[30] Gold et al. (1996) examine a variety of ways of revealing these relative utility levels.

This variant of cost effectiveness analysis offers additional information when program benefits are primarily improvements in the quality of life, rather than in its quantity. That is, it offers a quantitative measure, a count, of a subjective quality variable. However, when program benefits are entirely reductions in premature death, program benefits are identical for everyone. Each fatal illness averted represents a fixed number of life years added. As each life year counts as 1.0, regardless of who accrues the life year, the QALY count would be a multiple of the number of illnesses averted. The multiplier would be the number of life years saved by the vaccine. In this case, a cost-effectiveness estimate would be equivalent to a multiple of the first variant of cost effectiveness.

When programs affect only quality of life, leaving expected life spans unchanged, program ranking derived from a QALY count will be guided by income and circumstance just like the second variant, in which the choicemaker is a net national product-maximizing central planner. If the opportunity cost of morbidity is not counted in QALYs, such costs are counted in the numerator, evaluated at current wages. Garber et al. (1996, p. 41) observed that this practice raises the now familiar fairness question.

> To the degree that wages reflect opportunity cost, the time of persons in demographic groups that tend to have lower-paying occupations would be valued less. It remains controversial whether it is ethically acceptable, for example, to value the time of women less than that of men in CEAs [cost-effectiveness analyses], although this is the implication of the theory.

Income and Circumstance Influence Cost-Effectiveness Analysis

The surprising feature of cost-effectiveness analysis is that when analysts attempt to make their calculations relevant to public health decisions, either accounting for individual preferences (variant 3) or accounting for social costs (variant 2), policy guidance will be influenced by income and circumstance. All else equal, programs that offer benefits for the well-to-do will show greater cost-effectiveness than programs offering identical health benefits to the poor. In this regard, policy guidance offered by an analyst using cost-effectiveness is no different from guidance derived from conventional cost-benefit analysis where all benefits are monetized.

Cost-Effectiveness Analysis Does Not Measure Welfare

Using cost-effectiveness analysis, it is possible to have income and circumstance influence program choices without necessarily satisfying individual preferences. A simple example using the second variant of cost-effectiveness analysis shows that the ranking of programs from cost-effectiveness analysis may be entirely different from cost-benefit analysis based on WTP.[31] Suppose there is an environmental contaminant that causes cancer, and that everyone is equally exposed and susceptible to that cancer. Everyone faces exactly the same lifetime probability of contracting that cancer. Treatment costs are identical for everyone. Consider two potential programs that could eliminate exposure to the carcinogen, with one program benefiting only men and the other only women. That is, each program eliminates the risk for one half the population and does nothing for the other half.

The male-female wage differential means the program benefiting men will be calculated to be more cost-effective than the program benefiting women because the calculated indirect costs-of-illness prevented by the program for men would be higher than that for women. As the program costs and direct medical expenses averted are identical, net costs of the program benefiting men would be less than net costs of the program for women. With identical health benefits, the program with lower net costs would be more cost effective.

If we knew nothing about preferences, one might suspect that conventional cost-benefit analysis would point programs in the same direction as the second variant of cost-effectiveness. That men's wages exceed, on average, women's wages implies men have greater ability to pay for cancer risk reduction. If cancer risk reduction were a normal good, men's

[31] Conversely, Phelps and Mushlin (1991) argue that cost-benefit analysis and cost-effectiveness analysis often suggest similar or identical decisions.

greater ability to pay would imply greater demand for risk reduction, or equivalently greater willingness-to-pay for risk reduction.

However, studies of risk perceptions show clear demographic differences that sometimes swamp the influence of income. Flynn et al. (1994) show men are likely to dismiss the importance of a small environmental cancer risk. Men's willingness-to-pay for such risk reduction is therefore likely to be negligible. Women's willingness-to-pay to eliminate such a risk may be positive. Clearly, the ranking that results when projects are ranked by WTP analysis depends on the distribution of benefits and individual preferences. By luck, the ranking could mirror that of cost-effectiveness analysis. However, it would be just by luck for similar rankings to occur.

Program rankings derived from cost-effectiveness analysis will be similar to rankings from WTP studies only when WTP is so constrained that it loses its ability to represent individual preferences, its reason for being. When agencies require analysts to use the same single value for value of life (say, $5 million for all people and all risks), there is little difference between cost-benefit analysis and the first variant of cost-effectiveness analysis. To see this relation, denote this single value as *WTP*. Then, the dollar-denominated benefits would be nothing more than a multiplicative transformation of the count of adverse outcomes averted.

$$\begin{aligned} &\textit{Dollar benefits} = \\ &\qquad \textit{Total health outcomes averted} \; x \, \textit{WTP} \end{aligned} \tag{8}$$

In this case, a cost-benefit analysis would compare program (or compliance) costs with dollar benefits, as in equation 8. This comparison is nearly identical to calculations carried out under the first variant of cost-effectiveness analysis, which compares the count with program (or compliance) costs. As the difference between the two types of analyses is only a constant multiplicative transformation, this variant of cost-effectiveness analysis yields a program ranking identical to the ranking based on *WTP*. When analysts can estimate some of the systematic differences that exist, cost-benefit analyses may provide policy guidance that is quite different from cost-effectiveness.

The Importance of Transparency

A primary failing of cost effectiveness is that it does not convey information about net benefits. Thus, cost-effectiveness calculations can guide policymaking only if the person using the results assigns prices to life and health. Cost-effectiveness might reveal the cost of treating an individual whose kidneys have failed. But the decisionmaker has to decide whether it is worth incurring that cost. He has to compare that cost with what it is worth to him to keep the person alive. In effect, the decisionmaker has to acknowledge some reservation price, the maximum he would be willing to pay to continue the life of the individual whose kidneys have failed. An undertaken program must satisfy the condition that the decisionmaker's reservation price (the highest price he attaches to each health outcome) is greater than the cost of averting the adverse health outcome.

$$\begin{aligned} &\frac{\textit{Net cost}}{\textit{Total health outcomes averted}} < \\ &\qquad \textit{Reservation price} \Rightarrow \textit{Undertake project} \end{aligned} \tag{9}$$

When the decisionmaker chooses projects based on a rule relating his own subjective valuations of life and health to cost-effectiveness estimates, two entirely different sets of values are driving decisions. In calculating cost-effectiveness, analysts using variant two or three have imposed values for life and health equivalent to the value of lost productivity, usually wages paid. The decisionmaker's values are of course unique and not necessarily equivalent to earned income. However, it is the decisionmaker's reservation prices for different health outcomes that set net-benefit levels. The decisionmakers' unique values transform cost-effectiveness calculations, derived valuing lives as income, into net benefits that are either positive or negative.

Analysts' policy guidance may be greater if they know the decisionmaker's reservation prices. If they know the decisionmaker's reservation price for a cancer averted, then analysts can calculate the value to the decisionmaker of the health benefits each cancer prevention program provides. Armed with dollar values of benefits and costs, the analyst can calculate net benefits (using a consistent set of values), and thereby show which program offers greatest net benefits and which offers positive net benefits. Given reservation prices for strokes and cancers prevented, the analyst can compare the relative cost-effectiveness of stroke and cancer prevention programs. Of course, this estimate of net benefits would be in terms of the decisionmaker's values. Such an estimate would differ from conventional cost-benefit analysis because reservation prices are based on decisionmakers' values rather than on the values of those who might benefit. Obviously, the two sets of estimates need not converge.

Conclusion

The best use for cost-effectiveness analysis may be that demonstrated by Viscusi and Hamilton (1996), using the first variant. This variant of cost-effectiveness may serve as a coarse filter, helping to screen out programs that more complex analyses would also show are not worthwhile. Cost-effectiveness analysis could reveal those programs for which benefits are dwarfed by costs. Though useful in some situations, this use of cost-effectiveness has no theoretical appeal. It is not an individual welfare measure and does not fully account for costs avoided by programs. Thus, it cannot provide information useful for satisfying individual preferences or for helping an income-maximizing central planner. Failure on both of these counts means it cannot be considered a substitute for conventional cost-benefit analysis. Cost-effectiveness is less useful than conventional cost-benefit analysis because it cannot rank all activities or address whether any is worthwhile.

The often-stated reason for using cost-effectiveness analysis is that it avoids assigning values to life and health. Clearly, there is no merit to that claim. Cost-effectiveness simply pushes the pricing problem on to the decisionmaker. In addition, analysts may implicitly assign prices when making cost-effectiveness estimates. Those prices will be influenced by income and circumstance, making policy guidance subject to the same factors that cost-effectiveness tries to avoid. Cost-effectiveness can require making all the same estimates as conventional cost-benefit analysis. So the claim that it is easier to accomplish is incorrect.