Assigning Values to Life: Comparing Methods for Valuing Health Risks. By Fred Kuchler and Elise Golan. Food and Rural Economics Division, Economic Research Service, U.S. Department of Agriculture. Agricultural Economic Report No. 784.

Abstract

We examine five approaches economists and health policy analysts have developed for evaluating policy affecting health and safety: cost-of-illness, willingness-to-pay, cost-effectiveness analysis, risk-risk analysis, and health-health analysis. We examine the theoretical basis and empirical application of each approach and investigate the influence that assumptions embedded in each approach have on policy guidance. We reach four principal conclusions. First, the approaches are not interchangeable: they measure different things. Even estimates using the same approach are often not comparable because, in practice, there is little consistency in the application of any of the approaches. Second, the usefulness of each approach depends on the unit of account. The philosophical decision to eschew the monetization of health costs or benefits constrains the ability of the approach to rank policy options and to gauge the social desirability of policy. Third, all of the approaches except risk-risk analysis and one variation of cost-effectiveness analysis incorporate the effects of income and circumstance. As a result, policy guidance could be influenced by the distribution of income. Fourth, the theory and practice of willingness-to-pay estimation are in opposition. While it is now common practice for regulatory agencies to adopt the willingness-to-pay approach for estimating health and safety benefits, they do so by assuming away the importance of individual preferences. We build on these four conclusions to suggest the appropriate use of each approach.

Acknowledgments

We wish to thank Jean Buzby, Margriet Caswell, Roger Conway, Steve Crutchfield, Tanya Roberts, Peter Vardon, and Richard Williams for their comments.

Washington, DC 20036-5831 November 1999

Contents

Executive Summary

The resources to eliminate even a small portion of the hazards to life and health do not exist and policymakers must choose which hazard-control programs to fund. The most efficient use of resources is best gauged by comparing the costs and benefits of potential programs. However, where markets do not exist, or are incomplete, it is difficult to accurately gauge the costs and benefits of program choices. This report examines the major approaches economists and health policy analysts have developed for evaluating the benefits and costs of policy affecting health and safety: cost-of-illness, willingness-to-pay, cost-effectiveness analysis, risk-risk analysis, and health-health analysis.

One of the first problems analysts must confront when deciding how to gauge the costs and benefits of potential policy is what unit of measurement to use: dollars or physical units. The primary advantage of using a dollar scale is that it provides both a measure of net benefits and a complete ranking of expenditure alternatives. The philosophical decision to eschew the monetization of health costs or benefits constrains the ability of the approach to rank policy options and to gauge the social desirability of policy. Only the cost-of-illness and willingness-to-pay approaches use dollars to measure both costs and benefits.

Another problem analysts face when measuring the costs and benefits of health and safety intervention is determining what should be counted as a cost. Economists think of costs as consequences of choice and scarcity. Where goods are not scarce (more than what everyone might want) or where alternatives do not exist, choices have no cost. This observation leads to two important conclusions relevant to measuring costs and benefits. First, the range of choices and the perspective of the choicemaker will both have an impact on "cost." Second, "costs" could differ depending on whether they are measured before or after the policy change in question.

The different notions of cost and the various ways in which health and safety benefits and costs can be measured lead to substantial differences among the five major approaches. As a result, the approaches are not interchangeable and there are circumstances where one would be more appropriate than another.

The ***cost-of-illness*** (COI) approach tallies the dollars spent on medical expenses and the dollars of employment compensation that are forgone as a result of illnesses, accidents, or premature deaths. COI estimates frequently have served as a measure of the monetized benefits of government programs that promote health and reduce the number of premature deaths, illnesses, or injuries (the value of program benefits are the costs that are not incurred). Evaluation:

- Its theoretical legitimacy hinges on the assumption that national income is a valid measure of societal welfare. Many economists have challenged this assumption and most reject national income as a welfare measure.

- It equates the value of a life with forgone wages. Thus, higher paid members of society will be assigned higher values of life.

- It is not always a good measure of disease severity. Cost-of-illness estimates are influenced by a number of factors other than disease severity, including the current distribution of income, education, and employment skills, technological constraints to disease treatment, sick-leave policies, and health insurance systems (both private and public). As a result, cost-of-illness estimates often move in the opposite direction from disease severity measures.

- It is often characterized as the most practical of the valuation methods: data on direct medical cost and human capital costs are seemingly easy to collect. However, direct medical expenses are often difficult to assess accurately because of the intricacies of insurance arrangements; and human capital costs are equally difficult to ascertain because of the various forms of compensation that are available to employees.

- It is not a reliable lower bound estimate of willingness to pay.

- Despite its weakness as a measure of welfare or disease severity, it does provide a measure of the economic impact of illness. It provides an accounting of the dollars spent on medical expenses and the wage dollars that are forgone as a result of illness or premature death. Such an accounting is useful to economists and policymakers interested in gauging the magnitude of the economic flows resulting from government programs that improve public health.

The ***willingness-to-pay*** (WTP) approach measures the resources individuals are willing and able to give up for a reduction in the probability of encountering a hazard that will compromise their health. It assigns dollar values to life and health. Evaluation:

- It reflects individual preferences for risk reduction where the demand for risk reduction is derived from *ex ante*, or expected health benefits. These quantities exist only *ex ante*, at the moment of choice. They are not equivalent to realized damages.

- It reflects the observation that individual preferences are unique and individual demands for risk reduction vary. However, because health and safety are normal goods, some of the variance in willingness-to-pay estimates will be explained by income differences rather than preferences. So, just as in cost-of-illness analysis, income and circumstance could play a role in determining the size of willingness-to-pay estimates. When benefits are calculated as willingness-to-pay, policies may be guided away from programs that save poorer lives and toward programs that save more affluent lives.

- With WTP, individual preferences are aggregated and the Kaldor-Hicks potential compensation criterion is used to determine the social desirability of proposed policy. This approach entails an efficiency-first, equity-second rationale. The Kaldor-Hicks criterion has been criticized on the grounds that as long as compensation remains potential, social welfare is not maximized.

This criticism could be particularly serious in the case of policy concerning mortality. If efficient policy results in deaths, equity cannot be redressed through *ex post* redistribution schemes: it is impossible to redistribute between those who are alive and those who are dead.

- Empirical estimates of it have proved sensitive to the characteristics of the study population, the level of risk, and the type of risk. Willingness-to-pay results from one study are therefore not necessarily applicable across studies. Nevertheless, in practice, regulatory agencies that have adopted the willingness-to-pay approach have generally adopted a single value for lives saved where the value has been derived from compensating wage studies. Agencies apply their selected value to every health risk, regardless of the population likely to receive program benefits, the type of risk that might be mitigated, or the level of risk mitigated.

- Its valuations represent a consistent and faithful application of the principles of applied welfare economics. WTP measures provide the best estimate of individual welfare available to economists. While there is little reason to challenge WTP from a theoretical perspective, estimation raises practical problems because it depends on individual and idiosyncratic utility functions. With additional studies analysts may be able to estimate the demand for risk reduction throughout the population for a variety of different risks.

When analysts use ***cost-effectiveness analysis*** they attempt to measure benefits without assigning dollar values to life and health. 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. Evaluation:

- When analysts know anything about preferences for risk reduction and about the distribution of health and safety program benefits, rankings derived from cost-effectiveness analysis are likely to diverge from those derived from the willingness-to-pay approach.

- As it measures costs and benefits in different units of account, it is not intended to yield a net benefits estimate. Thus, the measures do not show whether any program is worthwhile. It is up to the decisionmaker to decide whether any program is worth the price.

- Only programs with identical health outcomes can be ranked using cost-effectiveness estimates.

- Results usually cannot be compared across cost-effectiveness studies as three distinct types of calculations are called "cost-effectiveness." Each type of calculation satisfies different goals. Two variants of cost-effectiveness analysis are subject to the influence of income and circumstance. Policy guidance in these variations will be influenced in the same way as it is with cost-of-illness and willingness-to-pay.

- Cost-effectiveness analysis may help minimize costs when an irrevocable decision has been made to take an action, but no decision has been made about technique or method.

- Of the three variants of cost-effectiveness analysis in common use, the simplest, the ratio of program costs to a count of health benefits, may be the most useful (this is the variant that is not influenced by income). 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. However, 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.

With risk-risk analysis and health-health analysis, analysts compare program costs and benefits without monetizing either benefits or costs. A ***risk-risk analysis*** enumerates the risks that are reduced and risks that are inadvertently increased by government health and safety policy. Both the desirable and undesirable risk changes are denominated in physical terms, though each could be denominated in different physical units. Evaluation:

- Estimates of benefits and costs are not influenced by resource scarcity and net benefits are not calculated. As a result, risk-risk analysis does not distinguish between expensive programs that offer few benefits and programs that dramatically reduce health risks at little expense. It is up to the decisionmaker to decide whether the benefits are worth the costs.

- It can only rank programs for which benefit and costs are measured in the same physical units.

- Risk-risk analysis is most useful in cases of all-or-nothing decisions. That is, only one program is offered and the decisionmaker must decide either to go forward with the program or to accept the status quo. When there are more options, risk-risk analysis shifts most of the burden of analysis to the decisionmaker.

Health-health analysis evaluates policies by comparing a count of deaths prevented with a count of deaths induced by transferring income from individuals to the government in order to finance government health and safety programs. This approach is built on two observations. First, risk reduction is a normal good, purchases of which increase with increasing income and decline when income falls. Second, government programs, even those that directly serve public health, have to be financed. Money for those programs has to come from individuals, and, thus paying for programs reduces individuals' ability to purchase risk reduction privately. Evaluation:

- It tallies benefits and costs in the same unit of account (lives), meaning that analysts can rank programs and calculate net benefits.

- Because income and mortality rates vary inversely, income effects in health-health analysis have an influence opposite to that of cost-of-illness or willingness-to-pay. Policies are guided toward programs that save poorer lives and away from those that save more affluent lives.

- It is restricted in its usefulness because it is applicable only to policy influencing mortality, not morbidity.

- Health-health analysis is an appropriate technique for comparing costs and benefits when analysts want to highlight both policy efficiency (net benefits) and the distribution of health (the extent to which one subpopulation might benefit at the expense of another). However, until the relationships between income and morbidity are better understood, health-health analysis can address questions only where benefits are denominated in the number of lives saved. Further, because analysts who use health-health analysis must translate dollars (income) into health, it may be easier to simply use standard cost-benefit analysis.