Eliminating Dollars from Cost-Benefit Comparisons—Risk-Risk and Health-Health Analysis

The previous chapter on cost-effectiveness analysis examined methods by which analysts try to avoid assigning dollar values to health and safety benefits. In this chapter, we examine methods by which analysts compare program costs and benefits without monetizing either benefits or costs. We first examine risk-risk analysis which was the first tool put forward for such comparisons. Risk-risk analysis is useful only in making choices when options are very restricted. It usually cannot rank programs or indicate whether net benefits are positive or negative. The other technique we examine, health-health analysis, is restricted to cases of mortality. It does, however, maintain some of the desirable characteristics of methods that monetize benefits and costs: it can rank programs and measure net benefits. This method explicitly relies on income to identify health costs of programs. The influence of income and circumstance is an integral part of the analysis.

Risk-Risk Analysis Is Not Influenced by Resource Scarcity

Health policy analysts have long recognized that many policies designed to lower particular public health risks unintentionally raise other risks. Lave (1981) argued that analysts could gauge the net health benefits of intervention by comparing the risks that government programs might reduce with the risks that these programs create. He named such a comparison risk-risk analysis. A risk-risk analysis enumerates the risks that are reduced and risks that are inadvertently increased. Both the desirable and undesirable risk changes are denominated in physical, and not dollar, terms.

Lave used the example of scrubber construction to illustrate how risk-risk analysis could be used to compare program benefits and costs. Health risks derived from pollution emitted from coal-fired electric generating plants might be reduced by installing scrubbers. But, as construction is a relatively risky occupation, building scrubbers is likely to raise the probability of injury for those involved in construction. Lave suggested that costs and benefits of government policy could be examined comparing the health benefits derived from improved ambient air quality with injuries incurred by scrubber construction.

As a practical matter, the health benefits and health costs included in risk-risk analysis are usually quite different. For example, suppose a water chlorination program were being evaluated on a risk-risk basis. Chlorinating water reduces exposure to a wide class of bacterial pathogens. Health benefits of chlorination consist of reduced incidence of many infectious diseases, including typhoid fever and cholera. Health costs consist of a higher risk of cancer through increased chlorine exposure. A risk-risk analysis would tally the reduction in the incidence of infectious diseases as health benefits and the increased incidence of cancer as health costs.

The notion of opportunity cost implicit in risk-risk analysis is a very small portion of opportunity cost in conventional cost-benefit analysis. In risk-risk analysis, the cost of reducing infectious diseases through chlorination is future cancer cases. The cost of carrying out the chlorination program and the cost of allowing preventable infectious diseases to persist are not tallied. Thus, risk-risk analysis does not offer distinctions between expensive programs that offer few benefits and programs that dramatically reduce health risks at little expense. Resource scarcity does not much influence benefits and costs tallied in risk-risk analysis.

As benefits and costs are usually tallied in different units, neither of which is dollar-denominated, risk-risk analysis offers no estimate of net benefits. Even if the limited opportunity cost notion implicit in risk-risk analysis were sufficient for decisionmaking, the lack of a common unit of account for benefits and costs poses problems for decisionmakers. In effect, the decisionmaker must assign prices to both benefits and costs. When a single program is at issue, say, deciding how many cancers can be tolerated to reduce waterborne diseases, the decision may be daunting. Where there are multiple programs at issue, each offering different health benefits or health

Figure 1. Causal chain linking program choices to health

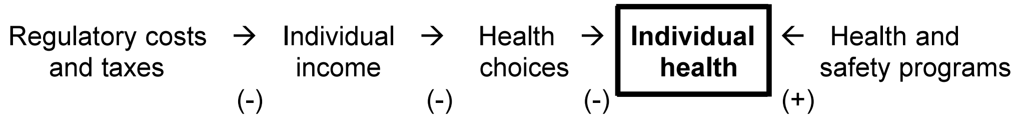

costs, the demands placed on the decisionmaker expand without bound.

Further, it would be very unusual for risk-risk analysis to offer a ranking of programs. To rank programs with risk-risk analysis, all programs that reduce a particular risk would have to induce the same set of health effects. Chlorination, for example, could be compared only with other water treatments that both control waterborne pathogens and are carcinogens. And even under such a constraint, it would be very unusual for a straightforward ranking to appear. Suppose a set of water treatments were carcinogenic and each controlled the same single pathogen. Then, analysts could tally the number of desirable health outcomes (say, a number of cholera cases reduced) and the number of induced adverse health outcomes (a number of cancers) attached to each program. If there were a program that displayed greater reductions of cholera and fewer induced cancer risks than all others, analysts could point to that program as most desirable, under a goal of reducing risks regardless of other costs. If instead, programs that most reduce risks also induce the greatest number of adverse outcomes, decisionmaking must be more complex, requiring someone to trade off cholera cases against cancers. That is, a count of desirable and undesirable health outcomes does not by itself suggest a way of trading off one for the other.

It is likely that many programs control or induce multiple health outcomes (like the large set of infectious diseases actually controlled by chlorination). In this case, an obvious best program is extremely unlikely to appear. As in the water treatment example, decisionmakers can avoid making tradeoffs only if there is a program that offers greater reductions in all infectious diseases than any other program, and that program induces fewer cancers. If the programs offer varying levels of control of each of the infectious diseases, no best choice is obvious just from the tally of diseases prevented and cancers induced. Under these conditions, analysts can construct a ranking only if they know how to trade off cholera cases against typhoid fever cases, as well as how to trade off infectious diseases cases against cancers.

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 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 Incorporates Resource Scarcity

Relatively new developments in economics suggest a role for analyses that do not monetize benefits or costs (Lutter and Morrall, 1994). There is a new technique by which analysts can estimate non-monetized benefits and costs that is consistent with the notion that resources are scarce.

The logic of this technique lies in 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.

The causal chain between financial costs imposed by government programs and unintended or induced adverse health effects is shown in figure 1. Negative effects proceed from left to right (eventually influencing individual health). Intended positive effects of risk reduction proceed from right to left, directly influencing individual health.

Reading from left to right, figure 1 shows that taxes reduce individual disposable incomes and constrain each individual's ability to purchase safety. A reduction in individual purchases of health-promoting goods and services will lead to increased mortality and morbidity. Reading from right to left, figure 1 indicates the direct benefits of government-sponsored health and safety programs.

Lutter and Morrall (1994) describe the small set of regulatory and judicial decisions regarding workplace safety that have been influenced by the logic in figure 1. More recently, Gramm and Dudley (1997) analyzed proposed EPA ground-level ozone standards, arguing that the economic cost of complying with the standards would result in a net increase in deaths:

> ...EPA's partial cost estimate implies an increase in mortality in the range of 50 to 700 deaths each year. If our estimate of the full costs is accurate, the financial costs of this rule could result in more than 7,000 deaths per year. (p. 18)

Lutter and Morrall argue that analysts could compare a count of fatalities averted by public-sector programs with a count of fatalities induced by regulatory costs. They named such a comparison "health-health analysis." Such analyses retain some of the desirable characteristics of conventional cost-benefit analysis. Because benefits and costs are measured in the same unit (lives), net benefits can be calculated. If we maintain the notion that positive net benefits are indicated when benefits are numerically greater than costs, net benefits are positive in the case when government health and safety programs save more people than they inadvertently kill.

Keeney provides some illustrative calculations, showing how changes in income could be used to examine regulatory costs and benefits. He postulates a negative exponential shape for the function relating income to mortality based on the observations that the poor do not live as long as the rich and that there is a limit beyond which no amount of health expenditure will reduce the mortality probability. Keeney relies on existing statistical studies measuring income and mortality, and demonstrates a relation between income and health. He calibrates his postulated functional form as in fig. 2.[32]

Figure 2
Negative exponential income-mortality function

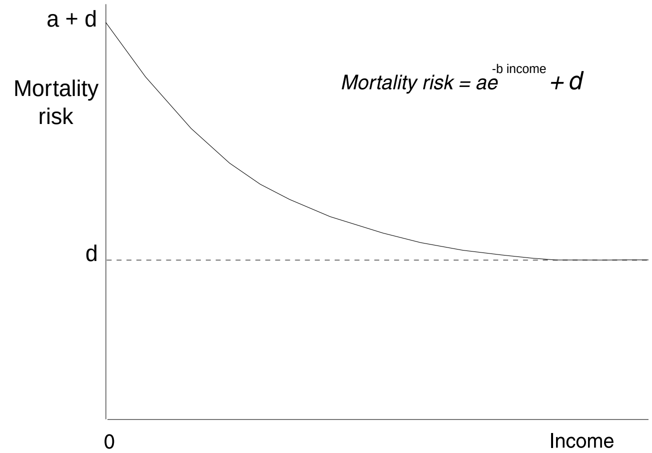

Figure 2 is an exponentially decreasing function relating income to the probability of death. The parameters a, b, and d are assumed positive. The probability of death is highest, a + d, when income is zero. Parameters a and b indicate the rate at which the probability of death diminishes with increases in income. With this function, that rate decreases as income increases. Increases in income are unable to reduce the probability of death below the level indicated by d.[33]

Lutter and Morrall (1994) show that the relation between income and mortality could be derived from a model of individual utility maximization.[34] Their model reveals a particular theoretical relation between income and mortality: the income loss nec-

[32] Lutter and Morrall graphically presented income and mortality panel data from 101 countries. Their visual data presentation is striking confirmation of Keeney's observation.

[33] Figure 2 has a negative slope, indicating that many risk decisions could be affected by income changes and analysts do not know exactly which health effects are most important. There are cases where particular health effects could be linked to income, and in those cases the relation between income and mortality can be more complex than figure 1 indicates. For example, Ruhm (1996) empirically shows the relation among income, alcohol consumption, and automobile fatalities. As alcohol is a normal good, income reductions lead to reduced consumption and a reduced number of intoxicated drivers. Fewer alcohol-related automobile fatalities occur when income levels fall.

[34] They did not specify a particular form for the utility function. Thus, their model does not yield a particular functional form for the linkage between income and mortality.

essary to induce one premature fatality is proportional to the willingness to pay to avoid a premature statistical death.

Instead of looking for particular risks that inadvertently arise, as in risk-risk analysis, Lutter and Morrall argue that analysts should look at how programs influence individual behavior toward risk. They argue that individuals are responsible for managing risks they face. Every choice a person makes requires choosing an acceptable level of risk. Individuals make risk decisions when they choose their medical care, the neighborhoods in which they live and work, safety features built into cars and appliances, foods they eat, and a host of other goods. These risk choices affect health and safety, and like other health- and safety-enhancing goods and services, these choices tend to be influenced by income (see discussion in "Cost-of-Illness Approach"). The risk levels individuals voluntarily accept depend on how much risk reduction they can afford. When incomes rise, individuals generally purchase greater assurance of safety. When incomes fall, individuals can afford less risk reduction and life becomes more risky.[35]

Lutter and Morrall explain the relation between taxes and health risk by observing that tax and regulatory policies influence disposable income, and through income, these policies influence the way individuals manage the risks they face. That is, policies influence individuals' ability to pay for risk reduction. Thus, any government action financed by additional taxes or any government program imposing compliance costs, will be accompanied by a predictable increase in adverse health outcomes. As a result, an action intended to protect public health may reduce some risks while inducing others.

For the many people in robust good health, it would take an extraordinary income loss to reduce their health, far beyond the tax price they might face from a single new government program. It would be very unlikely for a typical tax price to influence a particular person's health. But, these observations do not diminish the importance of the relation between government expenditures and health. When we consider costs incurred in national or global markets, even price increases that appear small (or wage increases that fail, by even the slightest margin, to keep pace with inflation) may have measurable risk consequences. Consider an action that raises some consumer prices and takes only a few dollars from an individual's purchasing power. Forecasts for the survival of that particular person would not be much affected by such a small change in opportunities for risk reduction. Perhaps fatality risks might rise by a factor of one-in-a-million because of some trivial sounding change in behavior. However, with 260 million people facing similar reductions in purchasing power, and each making some trivial sounding adjustments in behavior, we could anticipate 260 deaths. (See Chapman and Hariharan (1996, p. 53) for statistical evidence showing the relation between marginal income changes and mortality.) The numerous sources of adverse health outcomes might not be identifiable, but their aggregate result would be real deaths and illnesses.

An income-mortality function of the shape Keeney described means that the health consequences of new taxes and regulatory costs may differ among consumer sub-classes. Keeney showed that all income changes are not alike: changes that vary according to demographic patterns may change aggregate mortality even when income changes are strict transfers. For example, the health consequences of income losses imposed on the relatively wealthy may be much smaller than those from losses imposed on consumers of modest means, and a transfer from one group to the other may leave aggregate income unchanged but change average mortality. Following Keeney's postulated exponentially decreasing function relating income to mortality, the incidence of program-induced mortality is regressive, with equiproportionate impacts on income causing more than proportionate adverse impacts on the poor.

Empirical Evidence Linking Income and Mortality

One of the most difficult steps in conventional cost-benefit analysis is monetizing health benefits. The counterpart in health-health analysis is the step that transforms dollar costs into lives lost. Cost-benefit

[35] Of course there are highly risky goods and services such as skiing, sky diving, and mountain climbing that are consumed primarily by the rich.

and health-health analyses share the characteristic that they transform one variable into the unit of account of the other. For health-health analysis, the critical step is estimating the impact of income losses on mortality.

Numerous studies have offered insights into the relation between income and mortality using both macroeconomic data and individual health and income records. Most studies have produced point estimates of the relation between aggregate (or average) income and average mortality. Both Lutter and Morrall (1994) and Viscusi (1994) summarize results of statistical studies examining the relationship between income and mortality. Both establish some commonality among the studies by calculating, like Keeney (1990), an aggregate income loss per statistical death for each study. The estimates range from $1.9-$33.2 million (Nov. 1992 dollars). As Viscusi notes, "...these studies differ widely in the time period analyzed, the sample being addressed, and the other variables taken into account" (p. 8). Viscusi cites a study by Chapman and Hariharan (1994) that yields a middle-of-the-range estimate of a $13.3-million drop in income yielding a statistical death. The Lutter and Morrall list includes several studies with estimates similar to those of Chapman and Hariharan.

An analysis limited to deriving impacts based on relations among averages may conceal distributional effects. More precise information about who incurs dollar losses will result in better estimates of the number of induced fatalities. Without accounting for incidence, estimates could fail to describe some large adverse effect on a particular subpopulation. More recent work has accounted for major demographic differences in estimating the functional relation between income and mortality. Chapman and Hariharan (1996) carried out a longitudinal study of men initially aged 45-59, accounting for initial health status and some genetic factors influencing longevity. They found that the drop in income that induces a death in the lowest income quintile is approximately one-half the cost that induces a death in the highest income quintile. Further, they showed that most of this difference is between the highest and second-highest income quintiles: differences among the bottom four quintiles are relatively small. Differences between the median and lowest quintile range from 12 to 16 percent.

Clearly, two sets of factors influence the impact of income on an individual's health: consumption of health-enhancing goods and services, and the productivity of that consumption. The former depends on income, but may be conditioned by other factors, like education. The latter depends on demographic characteristics, like age. Productivity of health-enhancing goods also depends on existing health stock: the benefits of consuming health-enhancing goods depend on whether a person is healthy or ill. A person whose health is significantly compromised may have more difficulty surviving health insults than a person in robust good health.

To completely characterize the adverse consequences of regulation-induced income losses, analysts will have to account for both consumption choices and the productivity of these choices. The most transparent way to calculate health costs of income loss is to break the relation between health and income in two parts: (1) a behavioral function, relating the demand for health-enhancing goods to prices, income, and socio-demographic factors, and (2) the health production function, relating health status to consumption of health-enhancing goods. Ultimately, the usefulness of the two functions will depend on how detailed they are with respect to social, demographic, and economic variables.

But how detailed will the two functions have to be before they are useful? If all individuals had identical incomes and identical risk preferences, health-health analysis would be a minor variation of cost-effectiveness analysis. But, as noted in the sections on the willingness-to-pay approach and cost-effectiveness analysis, we do not have to look very far to see that profound and systematic differences in risk attitudes exist. That risk attitudes vary throughout the population implies differences in willingness to prevent exposure to hazards. We can therefore expect that the likely number of induced adverse outcomes will depend on whose income is compromised and the magnitude of the loss. Keeney (1990) argues that program-induced mortality varies systematically, largely influenced by income and by gender. Lutter and Morrall (1994) suggest the importance of age and ethnicity. In any case, the demographics of income losses largely determines the count of adverse outcomes. The estimate of induced deaths depends on who bears the costs.

Keeney (1997) analyzed data from the National Longitudinal Mortality Study to estimate the relationship between income and the annual probability of death. He estimated the relation as a negative exponential function for white males, black males, white females, and black females. In effect, these distinctions recognize that income and individual physiological characteristics determine mortality probabilities.

Kuchler et al. (forthcoming) examined a proposed seasonal harvesting ban on Gulf of Mexico oysters intended to prevent exposure to the bacteria *Vibrio vulnificus*. Oyster-related deaths have been traced to consumption of infected raw oysters harvested from the Gulf of Mexico during warm weather months. Epidemiologists identified the at-risk population as adult raw-oyster consumers suffering cirrhosis or immune-compromising diseases. Kuchler et al. estimated dollar costs imposed on the oyster industry, its harvesting and processing component, and the Gulf economy. As most oyster harvesters are white males and oysters shuckers are typically black females, Keeney's estimated functions were used to distribute costs among white males and black females, reducing income levels in each demographic category. Kuchler et al. calculated that the seasonal harvesting ban would, in an expected value sense, annually induce three deaths from each category and two to six more across the Gulf region (where no particular demographic cost incidence information is available). These estimates can be compared with an estimate of 17 deaths prevented.

As illustrated above, it is important to establish the influence of economic and demographic characteristics on health outcomes. As a result, empirical applications of health-health analysis must also carefully identify the distribution of the costs and benefits of health policy. In practice, when a solution to a public health problem is proposed, analysts can often characterize the demographics of those who suffer the regulatory costs. Further, analysts are likely to know some details of incidence of the public health problem. Epidemiological evidence usually reveals the demographic characteristics of the group that might benefit from public action. For example, *Salmonella*-caused human diseases are often more serious for children than for adults (CAST, 1994).

Health-Health Analysis—A Second Best Approach?

Health-health analysis shares some characteristics with conventional cost-benefit analysis. Like cost-benefit analysis, health-health analysis uses a single unit of account to measure costs and benefits. Therefore, like cost-benefit analysis, health-health analysis provides a measure of net benefits. For example, suppose the goal of a regulatory agency is to maximize the number of deaths prevented. Such an agency would view a policy that prevented 17 deaths but caused 12 deaths as offering positive net benefits. The policy would be inferior (in the sense of net deaths prevented) to one that prevented 15 deaths but (because of smaller taxes or compliance costs) induced 2 deaths. Either policy would be superior to one that prevented 20 deaths but induced 18 deaths.

Income and circumstance play a powerful role in both health-health analysis and conventional cost-benefit analysis. In conventional cost-benefit analysis, these variables influence the theory and practice of benefits estimation, thereby influencing policy guidance. In health-health analysis, income effects are integral components, although incorporated through costs rather than benefits. While conventional cost-benefit analysis might show relatively larger benefits when benefits accrue to wealthy individuals, health-health analysis might show relatively larger costs when costs accrue to poorer individuals. That is, health-health analysis is more likely to guide policies away from programs that impose costs on the poor.

Though health-health analysis shares many characteristics with conventional cost-benefit analysis, it is not a perfect substitute. Lutter and Morrall note that health-health analysis is a second-best test "relative to BCA [cost-benefit analysis] because it excludes from consideration those costs unrelated to health and safety risk. If used alone, rather than as the first step in a benefit-cost assessment, the health-health test is more lenient than BCA."

Of course, passing a weak test may not be informative. Analysts would not be able to say whether the passing grade occurred because benefits really exceed costs or whether there simply is not enough information available about costs. However, it is not necessarily the case that the more stringent test is always preferred. Failing a lenient test is powerful evidence that costs really do exceed benefits.

Health-health analysis is an appropriate technique for comparing costs and benefits under limited circumstances. 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), health-health analysis is appropriate. However, until relations between income and morbidity are 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 is surprising that they do not simply use standard cost-benefit analysis. When analysts can assign prices and can discuss dollar-denominated costs, conventional cost-benefit analysis provides a straightforward market test for government programs. In choosing to use health-health analysis, there must be some reason why analysts cannot or choose not to assign values to life. As discussed in "An Introduction to the Methodologies," one reason may be that costs and benefits that are denominated in lives convey a different type of information than those denominated in dollars. A decisionmaker confronted with a benefit-cost ratio of 5 dollars to 4 would have an easier decision than one confronted with a benefit-cost ratio of 5 lives for 4. Dollar-denominated transfers are unlikely to raise the questions that health transfers do.