Introduction to Benefit-Cost Analysis

What is benefit-cost analysis?

There are many ways that the advantages and disadvantages of alternative projects or policies can be assessed. Private investors will assess projects on the basis of the rate of return they expect to realize in relation to the risk involved. Environmental advocates will assess projects and policies on the basis of potential bio-physical consequences and risks. Business and labour groups will look to the economic stimulus a project or policy may have—the amount of employment and income expected to be generated. Different groups have different perspectives and interests, and accordingly will assess projects or policies in different ways.

Benefit-cost analysis does not take the perspective of any one particular group. Nor does it impose the standards or values that any one group would argue should be applied in the development of public policy. The purpose of benefit-cost analysis is to identify and compare the advantages and disadvantages of alternative projects and policies from the point of view of society as a whole. It is intended to:

- take into account the impacts on all those who are affected, or at least all those affected within the jurisdiction for which the analysis is being done;[1] and
- measure the magnitude or significance of those impacts on the basis of the values that the affected individuals hold, more specifically, the trade-offs they would willingly make to realize the positive impacts or accept the negative.

Thus, in a benefit-cost analysis, benefits include all impacts or consequences that at least some people value positively—that they either willingly pay for, or in principle would be willing to pay for in order to get. For example, benefits could include an increase in the output or availability of a good or service; an improvement in health or social conditions; a reduction in risk to the environment. They could also include a reduction in the cost of acquiring or accessing goods and services, or an increase in income potential afforded by new skills, education and employment opportunities.

Costs include all impacts or consequences that at least some people value negatively—impacts or consequences for which people either are compensated or would have to be compensated to willingly accept. For example, costs could include the supply of labour and other inputs; a reduction in the availability or increase in the cost of a good or service; a diminution of environmental quality or social conditions; an increased risk of injury or premature death.

In short, benefits are what people are willing to pay for and costs are what people must be compensated for to be willing to accept. At its most fundamental level, benefit-cost analysis sums and compares the total amount that people are willing to pay for the positive consequences of a project or policy in relation to what people have to be compensated to offset the negative.

[1] More specifically, benefit-cost analysis is based on implications and values for all those whose interests are of concern (those with standing in the policy debate). One would not consider, for example, the preferences of thieves in a benefit-cost analysis of a crime prevention program.

Willingness to pay and compensation demanded are both measured, to the extent possible, in dollar terms. But that does not mean benefit-cost analysis is only concerned with impacts or consequences that have market prices or commercial value. Many of the most important consequences considered in benefit-cost analyses—for example, impacts on the environment, health, life expectancy, travel times—are not bought and sold at market-determined prices nor are they valued solely for commercial purposes.

The amount of dollars a person is willing to pay indicates the purchasing power over other goods a person is willing to give up, in other words the total amount of other goods a person is willing to trade off in order to acquire the benefit. Similarly, the amount of dollars a person must be compensated in effect indicates the amount of other goods a person would require as a trade-off for providing labour, supplying other inputs or accepting some negative consequence or cost. The dollars themselves are just a convenient unit of measurement. They are used to indicate the relative value or significance people place on different consequences—relative values that can be summed and compared to calculate whether and to what extent alternative projects or policies may yield positive net benefits.

A project or policy that exhibits positive net benefits (or the greatest net benefits among a group of competing alternatives) is deemed to be economically efficient. The sum of the value that people place on the positive consequences exceeds the value or compensation required to offset the negative. In that sense, the project or policy offers a net increase in value for society as a whole. Of course, not everyone will necessarily be better off. Society as a whole in benefit-cost analysis is simply the sum of the impacts on all affected individuals. The individuals that benefit may not be (and often are not) the same as those who incur costs. The bottom line net benefit is effectively silent on the equity or distributional aspects of the project.

There was a tradition in economics that assumed one could measure total social welfare—the total level of utility or satisfaction realized by all members of society. Total social welfare would be the sum of the dollar measures of impact weighted by the marginal (increase or decrease in) utility or well-being that each dollar of benefit or cost would have for the affected persons. The weighting would take into account the generally held assumption that a dollar of benefit or cost would have less impact on a person's total utility or well-being the higher the person's income. A $100 benefit to a person earning $100,000 per year typically means a lot less than a $100 benefit to a person earning $10,000 per year. The sum of the weighted benefits and costs would then indicate whether and to what extent total social welfare would increase as a result of any proposed project or policy.

While there are still some proponents of weighted benefit-cost analyses,[2] it is widely recognized that there is no correct way to decide how much more one person's benefit or cost should count relative to another. Nor is there a clear policy implication from a

[2] For a discussion of weighting net benefits to account for income distribution, see H. Campbell and R. Brown, **Benefit-Cost Analysis**, Cambridge University Press, 2003, chapter 11.

finding that a project or policy would increase total social welfare. Even if a progressive weighting scheme had been used in the calculation, the increase in social welfare might still leave disadvantaged groups (or for that matter any particular group) worse off. On that basis alone it may not be appropriate or acceptable to undertake.

It may be of interest to determine how sensitive a finding of positive net benefits is to different possible weighting schemes for the net beneficiaries and losers of a project or policy initiative. However, it is important to recognize what benefit-cost analysis does and does not do. It provides a measure of the economic efficiency of different alternatives. In economic jargon, it determines whether and to what extent different alternatives offer a ***potential Pareto improvement***. A potential Pareto improvement is a change where the net benefits of the winners are sufficiently large that the winners could in theory compensate the losers and still be better off. [3] Of course, that compensation need not be paid. A potential Pareto improvement doesn't indicate that no one will be worse off—compensation or other distributive measures accompanying the project or policy will determine that.

Why Benefit-Cost Analyses Are Undertaken?

In free market economies, individual households and firms are assumed to undertake their own benefit-cost assessments of the purchases, sales or investments they may make. When deciding whether to buy a particular good or service, households will effectively compare how much they are willing to pay (the value they place on the good or service) with the cost. When firms consider whether to produce a particular good or service, they will compare the price they expect to receive with the compensation they require to offset the costs they will incur.

These individual assessments will govern the demand and supply of different goods and services in the economy and ultimately the amount of each good and service that is produced. A principal strength of the free market economies is that resources are allocated in response to people's values or preferences for different goods and services, and the costs (compensation required) for their production. Under ideal market conditions, the economy will produce those goods and services that people value most in relation to their costs—in other words, that generate the greatest net benefits.[4]

The question is, if free market economies can generate the greatest net benefits—why should benefit-cost analyses be undertaken? Why not let the market decide what projects

[3] Pareto was a famous Italian economist who developed evaluation criteria that don't require the weighting and aggregation of different individual benefits and costs. A *Pareto improvement* is a change that makes at least one person better off and no one worse off. A project or policy that offered a Pareto improvement would be a circumstance where there would be an efficiency gain without any adverse distributional effect (in that no one is worse off). However, opportunities for Pareto improvements are likely to be rare. In any event, overall net benefits do not indicate whether or not there is a Pareto improvement; rather they indicate whether there is a potential for a Pareto improvement if the winners did compensate the losers and the payment of compensation didn't affect the valuation of the benefits and costs, which it could in fact do.

[4] More carefully stated, the economy will produce an amount and mix of goods and services that are *Pareto efficient*—from which no Pareto improvement can be made.

should be undertaken—what goods and services should be produced—and restrict government policy to ensuring that markets work effectively?

There are those who argue precisely that—that there is no need for government intervention and therefore no need for benefit-cost or other assessments of projects or policies that government might consider authorizing or undertaking. However, most recognize that there is a role for government and there is a consequent need for the benefit-cost assessment of alternative projects and policies.

Optimal market outcomes require ideal market conditions. Firstly, there have to be markets for all of the goods and services that people value. And, the markets have to be perfectly competitive, so that individual buyers or sellers cannot manipulate the market price. Finally, there must be no taxes, subsidies, regulations or other factors distorting the market price.

There are numerous reasons why these ideal market conditions don't prevail—why there are market failures preventing the economy from producing the mix of goods and services that people value most in relation to their cost. Some of the major ones are:

- Externalities:- Positive externalities are beneficial impacts on third parties for which the third parties don't pay; negative externalities are adverse effects on third parties for which the third parties are not compensated. Immunization, education, safety programs are all examples where there can be significant positive externalities. Pollution and traffic congestion are classic examples where there can be negative externalities. Market economies will undersupply goods that offer positive externalities and oversupply those that generate negative externalities.
- Public goods:- In economic terminology, public goods refer to those goods that individuals collectively consume (for example, the preservation of endangered species, parkland, certain types of infrastructure). Pure public goods are non-rivalrous (one person's consumption does not detract from the amount available for others) and non-excludable (individuals can enjoy or benefit from them without paying a fee). The purchase of such goods by any one individual would in effect create positive externalities for others. Market economies will not supply an optimum amount of public goods.
- Price distortions:- These arise whenever prices aren't set competitively, don't balance demand and supply, or are subject to taxes and subsidies. Whatever the cause (monopoly market structure, government or regulatory intervention, unemployment, institutional arrangements), the distortions will result in inefficient price signals and consequently sub-optimal levels of production.

In all of these cases, there is a need for government intervention to achieve an optimal level and mix of goods and services. Benefit-cost analyses are used to assess whether different proposed projects or policies will in fact contribute to that goal.

What Role Can Benefit-Cost Analysis Legitimately Play?

The more zealous proponents of benefit-cost analysis would argue that by identifying the most economically efficient alternative—the one that best reflects individual preferences—benefit-cost analyses determine what *should* be done. There is, however, no moral or other justification for such assertions any more than there is a moral justification for 'optimal' market outcomes free of all market failures.

Optimal market outcomes are optimal only in the sense that goods and services are efficiently provided in response to individual values and preferences. They don't, however, ensure there is equity in the distribution of income governing what is produced, and for whom. Nor do they question the nature, basis or quality of the preferences that people have.

Benefit-cost analysis, like the market itself, ignores the income dependence of preferences and it is silent on the income distributional effects of the alternatives being assessed. It also ignores the context dependence and validity in any philosophical sense of people's manifest preferences. The choices people make or preferences they reveal depend on social, media and other influences They can be misinformed and can change. What is considered optimal at any point in time may be based on ephemeral and arguably shallow factors.

In short, there are significant social and philosophical limitations to benefit-cost analysis. There are as well very significant empirical limitations. Even if one accepts the equity and validity of the preferences that one is trying to satisfy, it can be very difficult to determine exactly what those preferences are. Benefit-cost analysis uses dollar estimates of the trade-offs people would willingly make to measure relative preferences for different benefits and costs. In many instances those trade-offs cannot be readily inferred from observable behaviour or reliably estimated from survey or other means.

For these reasons, it is a mistake to suggest that benefit-cost analysis can determine what *should* be done. Benefit-cost analyses are important. But the role they serve is not to resolve public policy debates; rather their role is to inform those debates—to challenge advocates and decision-makers to take account of the values and preferences of all of the affected parties—what society as a whole, for better or worse, would choose.

Disaggregated multiple account evaluation frameworks explicitly recognize this role of benefit-cost analysis. These frameworks present the different types of benefits and costs in different accounts, clearly distinguishing how different interests (e.g. project developers, taxpayers, consumers, the environment, workers) are affected. Moreover, not all of the consequences need to be summarized in dollar terms. Trade-offs can be identified and critical values—what an impact would have to be worth for a particular alternative to be favoured—can be calculated. The goal is to inform in a clear and consistent way the different implications to the different affected parties. It isn't simply to provide, as in traditional benefit-cost analysis, a bottom line monetary estimate of overall net effect.