PRINCIPLES OF COST-BENEFIT ANALYSIS

Robin Boadway

Professor, Department of Economics, Queen's University, Kingston, Canada

Abstract

This paper summarizes the procedure for the economic evaluation of government projects and policy reforms. It begins with the social welfare function underpinnings of cost-benefit analysis including the role of distributive weights and the choice of numeraire. It then turns to the conduct of a social cost-benefit analysis using the net present value criterion. This includes the shadow pricing of market products and inputs affected by the project, indirect welfare effects, the opportunity cost of project finance, the evaluation of non-marketed inputs and outputs, and the opportunity cost of risk. Issues involved in selecting a discount rate are discussed, especially those arising from imperfect capital markets. Finally, since many public projects have long-term consequences, the principles that might be used to take account of effects of projects on future generations are outlined. Techniques for accounting for these effects, such as generational accounting, are summarized and its shortcomings highlighted.

I.   INTRODUCTION : SOME FOUNDATIONS

Cost-benefit analysis in its broadest sense is the process of ranking policy options from an economic point of view, taking account of both the benefits of the policy and its costs. The policies range from an investment project that is small enough so that a partial equilibrium approach will suffice, to a broader fiscal policy change, such as a tax, subsidy or regulation, that will have general equilibrium repercussions on several markets. In taking an economic point of view, we are concentrating on net benefits of the policy as they affect the well being of individual households of the economy, typically as judged by their own preferences. We are eschewing political feasibility considerations as well as non-economic objectives, such as non-discrimination, liberty, and so on.

That does not mean that social values are not involved. Indeed, it is impossible to rank policy options without making some important value judgments, but the ones that we make are those

typically involved in applied welfare economics, and it is worth being explicit about them at the outset. There are three main ones. The first is the precept of *individualism*: an individual's welfare should be based on the individual's own preferences. The second is the *Pareto principle*: if one individual is made better off and none are made worse off as a result of implementing the policy, social welfare is taken to increase. The third is the principle referred to as *welfarism*: social orderings of alternative policies depend only on the welfare of individuals and not on extraneous considerations (like freedom, non-discrimination, etc.). These principles lead to a set of social preferences over allocations of resources that can be summarized in a *Bergstrom–Samuelson Social Welfare Function* (SWF) of the form $W(u_1(\cdot), u_2(\cdot), \cdots, u_n(\cdot))$, which is increasing in all arguments, and where $u_i(\cdot)$ is the utility of household i as a function of a resource allocation.

The SWF represents a preference ordering over individual utilities. As such, it is an ordinal concept, so any increasing function of $W(u_1, u_2, \cdots, u_n)$ will serve just as well. However, the SWF requires that individual utilities be measurable and that they be comparable among households. There is a sizeable literature, summarized in Boadway and Bruce (1984), on the concepts of measurability and comparability of individual utilities functions, and how the extent of measurability and comparability influences the form of the utility function. We need not be detained by that literature. Instead, we shall adopt for illustrative purposes a particular and commonly used form for the SWF that incorporates some reasonable assumptions about measurability and comparability, and also allows us to capture social attitudes towards equity in a single parameter. The SWF will be assumed to take the following additive form:

$$W(u_1, u_2, \cdots, u_n) = \sum_i w(u_i) = \sum_i u_i^{1-\rho} / (1-\rho) \tag{1}$$

where $w(u_i)$ is the social utility of individual i.[1]

The SWF in (1) has a number of notable properties. The social preference ordering is symmetric and anonymous: the same social utility function $w(\cdot)$ applies to all individuals, and it gives rise to preference orderings that are symmetric around the equal utility points. Thus, in the two person case, social indifference curves in (u_1, u_2)-space are symmetric around the 45-degree line. The parameter ρ captures society's aversion to inequality. In particular, ρ is the elasticity of marginal social utility:

$$\rho = -\frac{w''(u)}{w'(u)u} \tag{2}$$

We shall refer to ρ as the *coefficient of aversion to inequality*, analogous to the coefficient of relative risk aversion in an uncertainty setting. The higher is ρ, the more aversion there is to inequality in utilities. It may be reasonable to assume non-negative aversion to inequality, in

---

[1] The SWF in (1) is linear homogeneous. Such a form is possible under the assumption that individual utility functions are comparable and measurable up to a ratio scale. That is, proportional changes in utility are comparable across individuals. Technically, individual utility functions can be subject to multiplicative transformations (e.g., u_i can be transformed to $v_i = k_i u_i$) without affecting the social ordering. See Boadway and Bruce (1984), Chapter 5.

which case $\rho \geq 0$. This is equivalent to assuming that the SWF $W(\cdot)$ is quasi-concave. The two limiting cases are the utilitarian SWF, where $\rho = 0$ (so social welfare is just the sum of utilities), and the maximin (or Rawlsian) SWF, where $\rho = \infty$ (where social welfare is the utility of the least well-off individual).

The value of the coefficient of aversion to inequality incorporates a value judgment that reasonable people may differ on, and this poses enormous problems for cost-benefit analysis since virtually all policy alternatives have different relative effects on different individuals. It is not possible to rank such alternatives unless one makes an explicit comparison of utilities of different households. Unless one alternative is Pareto superior to another, rankings depend upon the SWF that one has in mind. In fact, there are really two levels of value judgment involved in arriving at (1). Given individual preference orderings, one first has to measure the utility associated with each outcome, that is, formulate the individual utility function $u_i(\cdot)$, and then one has to decide on the social utility function $w(u)$ to apply to individual utility. As we shall see a common procedure for measuring welfare levels in applied welfare economics is to use a *money metric* measure of utility, that is, to measure welfare levels using the values of expenditures required to achieve different welfare levels given a set of reference prices. Formally, define the real income of household i, denoted y_i, by the expenditure function at a set of reference prices p^r, or:

$$y_i = e(p^r, u_i) \equiv \min_x \{ p^r x | u(x) \geq u_i \} \tag{3}$$

where x is a vector of commodities. Then, the utility function for household i can be written $u_i(y)$, where u'' reflects the rate at which the marginal utility of income declines. If the household utility function also takes the constant elasticity form, $u(y) = y^\alpha / (1-\alpha)$, the social utility function $w(u(y))$ can be transformed into:

$$v(y) = y^{1-\sigma} / (1-\sigma) \tag{4}$$

where $\sigma = \rho + \alpha + \rho\alpha$. Thus, σ now becomes the coefficient of aversion to inequality in real income. Note, however, whereas the same social utility function $w(u)$ may be used to convert individual utility into social utility, it may well be the case that converting real income to individual utility may require an individual-specific utility function $u_i(y)$. For example, different households may have different abilities to convert real income to utility if they differ in needs or disabling circumstances. Moreover, they may differ in preferences as well, which causes even more vexing problems for constructing comparable utility functions.

These conceptual problems cause inevitable difficulties for cost-benefit analysis. There are two standard ways of proceeding, assuming there is no consensus about the appropriate SWF to use. One procedure is to rank policy alternatives using a SWF with an explicit assumption about aversion to inequality, and to conduct sensitivity analysis with respect to the weights. In this way, the decision-maker can be given a menu of rankings depending on the weights used and be

responsible for choosing among them. For this purpose, it is typically convenient to roll the judgments about measurability of individual utility and aversion to inequality in utility into one, and apply the following SWF:

$$W(y_1, y_2, \cdots, y_n) = \sum_i y_i^{1-\sigma} / (1-\sigma) \tag{5}$$

The other procedure, most commonly used, is to ignore redistributive concerns and simple aggregate real income changes yen by yen regardless of to whom they accrue. This is the procedure advocated by Harberger (1971a), and is effectively equivalent to assuming the coefficient of aversion to income inequality σ in (5) is zero. Alternatively, if the SWF in (1) is used, the utilitarian form is assumed (so $\rho = 0$), and the individual utility function is assumed to be quasilinear in some consumption good (e.g., $u(x) = x_0 + g(x_1, x_2, \cdots, x_n)$), so the marginal utility of income is constant. Though this procedure is most frequently used in practice, the general theoretical case for it is disputed in the literature unless one is prepared to assume $\sigma = 0$.[2] The argument in favor relies on being able to separate efficiency and equity considerations in cost–benefit analysis.

Three sorts of arguments can be used for ignoring redistributive concerns. The first is related to the classic separation of efficiency and redistribution functions of government proposed by Musgrave (1959). The government has wide–ranging policy instruments available for redistribution, including the progressive tax system, transfer programs, social insurance and in–kind transfers. If it is using them effectively, equity consequences of other policies should be ignored, it being presumed that account will be taken of them elsewhere. Although this is a seductive argument, its limitations are evident. For one, the argument does not apply for policies that have explicit redistributive intent. For another, even if redistributive policies are set optimally, because these policies are distortionary, they cannot succeed in achieving the first–best social optimum. In these second–best circumstances, one cannot really ignore the redistributive consequences of policies except in special circumstances. The famous Diamond–Mirrlees (1971) production efficiency theorem indicates one such set of circumstances. If taxes are set optimally, there should be production efficiency in the economy, implying that the government should evaluate its own production plans according to the criterion of the value of output measured in producer prices. A similar argument can be adduced in evaluating projects involving traded goods: their shadow prices should be world prices uncorrected for equity concerns (Boadway, 1976; Drèze and Stern, 1987).

The second argument is that if the aggregate change in real income is positive ($\sum y_i > 0$), this indicates that those who gain from the policy should be able to compensate those who lose and still be better off. There are three problems with this hypothetical compensation argument. The first is that the compensation will not actually be paid, even if it hypothetically could be. In

---

[2] A comprehensive summary of the arguments against using this procedure may be found in Blackorby and Donaldson (1990).

these circumstances, ignoring equity consequences seems particularly arbitrary. Second, even if one accepts the hypothetical compensation criterion as a valid one, the aggregate change in real income turns out not to be a perfect indicator: positive aggregate real income changes need not be sufficient to satisfy the hypothetical compensation criterion (Boadway, 1974). Third, the hypothetical compensation criterion can take various forms depending on what one assumes to be the mechanism, and none of them yields a complete ordering of outcomes.

The third argument is that if there are many policies being undertaken over time, the redistributive consequences of them should roughly cancel out, so ignoring them for any one policy evaluation is innocuous. This is especially true if, following argument one, there are policies in place to address redistributive concerns. This argument seems to be no more convincing than the others.

Perhaps all these arguments are simply rationalizations for the more practical one. That is that it is very difficult to take account of the redistributive consequences of policies. Most often, only aggregate data that do not distinguish among individuals of different real income levels are available, so measuring aggregate real income changes is the only feasible procedure. This might be supplemented by some imperfect indication of the distributive consequences of the policy. Much of our discussion will focus on the use of real income measures of costs and benefits without redistributive weights explicitly incorporated. The general principles can best be illustrated this way.

Whatever view we take of the redistributive consequences of various policies, a basic ingredient is what we have referred to as the real income of individuals in the economy. To repeat, this is a money metric of utility (preference) levels of the household under different policies. Conceptually, in the simple case of utility in two goods, a money metric measure of utility change is the distance between two indifference curves measured in monetary terms. Of course, there are an indefinite number of ways of measuring the distance between two indifference curves, so the convention followed is to measure the distance using a budget line with given relative prices. Thus, our money metric utility measure is just the expenditure function: the value of expenditures required to reach a given indifference curve at some reference set of prices. It can be interpreted as the *willingness to pay* for the bundle of goods in a given situation. (This is a useful interpretation when it comes to evaluating non-marketed commodities as discussed later.) Of course, even this is not without ambiguity since the expenditure function differs according to the reference prices used. Two common examples when considering the effect of moving from policy 'zero' to policy 'one' are the compensating variation (CV) and equivalent variation (EV). The CV uses prices in the new situation to obtain a monetary measure of welfare change:

$$\mathrm{CV} = e\left(p^1, u^1\right) - e\left(p^1, u^0\right)$$

where superscripts 1 and 0 refer to the prices and utilities after and before the change

respectively. Thus, p^1 represents the vector of prices of both goods demanded and factors supplied in the new situation, and $e(p^1, u^1)$ is the net expenditure required to achieve the utility level of the new situation at the prices p^1. Alternatively, the EV uses the old prices:

$$\text{EV} = e(p^0, u^1) - e(p^0, u^0)$$

In general, CV and EV will differ even though both are equally legitimate measures of changes in utility levels. More generally, any set of reference prices would serve as well. In principle, the choice of reference prices can have consequences for relative utility changes across households, for example, if they have different preferences for goods. But, in practice, given measurement problems that are likely to exist, any set of reference prices would serve as well.

   Before leaving this section on the foundations, it is worth briefly turning to the practical issue of how one might incorporate distributive weights into cost–benefit analysis given the data constraints that are likely to apply. Distributive weights should take account of the fact that one yen of real income is of different social value to different individuals (e.g., Boadway, 1976; Drèze and Stern, 1987; Ahmad and Stern, 1991). Suppose we take as given the coefficient of aversion to real income inequality, σ, in the SWF given in (5). Then, changes in real income accruing to individual i should be weighted by the marginal social utility of income β_i defined by:

$$\beta_i \equiv \frac{\partial W}{\partial y_i} = y_i^{-\sigma} \tag{6}$$

Note that (6) simplifies to $\beta_i = 1$ for the case where $\sigma = 0$, so no welfare weights are used. Apart from the issue of determining the correct value of σ to use, implementing (6) for evaluating policies is difficult since it will not generally be possible to attribute costs and benefits to individuals according to their real income levels. Often, only market–wide data will be available. There may be some exceptions. For example, wage payments for individual workers might be available in which case they could be weighted using β_i, but generally that will not be possible for other cost items.

   A less demanding procedure that has been proposed by Feldstein (1972b) in the context of public sector pricing is to assign social weights to commodities according to estimates of the proportions in which they are consumed by individuals with different real incomes. The procedure is as follows. From basic consumer theory, the change in individual i's real income of from a change in a the price of commodity j is given (using Hotelling's lemma) by:

$$\frac{\partial y_i}{\partial p_j} = \frac{\partial e_i(\cdot)}{\partial p_j} = -x_j^i$$

where x_j^i is the demand for commodity j by individual i. Then, the change in real income from a change in all commodity prices becomes:

$$dy_i = -\sum_j x_j^i dp_j$$

Using this and (6), the change in social welfare from a change in consumer prices may be written:

$$dW = \sum_i \beta_i dy_i = -\sum_i \sum_j \beta_i x_j^i dp_j = -\sum_j R_j X_j dp_j \tag{7}$$

where $X_j = \sum_i x_j^i$ is aggregate demand for commodity j and:

$$R_j \equiv \sum_i \frac{\beta_i x_j^i}{X_j}$$

is defined as the distribution characteristic of good j. It is a weighted average of β_i's, where the weights are the proportions of commodity i consumed by individual j. Given that β_i falls with real income, R_j will be higher for commodities with lower income elasticities of demand since these will be disproportionately consumer by lower income persons. More generally, costs and benefits of a policy can be weighted by their distributive characteristics in evaluating a policy (see Boadway, 1976). This is clearly a procedure that is less demanding empirically than calculating real income changes for each individual and weighting them by β_i.
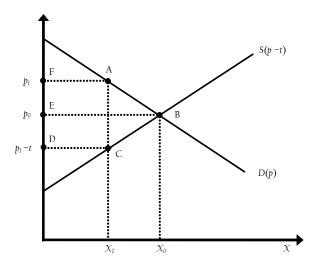
Consider, for example, the following illustrative application of the welfare effects of an excise tax, drawn from Harberger (1978). Given that the tax is on a single good, a partial equilibrium can be taken. Denote the aggregate output for the good simply by X, whose market is depicted in Figure 1. The imposition of a tax at the per unit rate t causes market output to fall from X_0 to X_1, and consumer price to rise from p_0 to p_1. If distributive effects were ignored, the standard measure of welfare loss would be the triangular area *ABC*, which is the difference between the losses in consumer and producer surplus (*FABE* and *EBCD*) and the gain in government revenue *FACD*. If distributive weights are attached to these losses and gains, the change in social welfare will be:

$$\Delta W = -R_D FABE - R_S EBCD + R_G FACD$$

where R_D, R_S, and R_G are the distributive characteristics associated with consumers surplus, producers surplus and government revenue, respectively. In the case of the first two, these reflect the shares of individuals of different income groups in the demand and supply of the good in question. The distributive weight on government revenue reflects the shares of government revenue raised from different income groups. In principle, this kind of methodology can be applied to any policy change.

Finally, before moving to the specifics of cost−benefit analysis of policies, it is worth drawing special attention to the problems associated with evaluating the net benefits of projects, as opposed to fiscal policies. A natural question that arises is why we cannot simply use financial profitability measures such as are used in the private sector, and must instead rely on cost−benefit analysis. In other words, what are the sources of differences between social profitability and private profitability?

Figure 1

The most fundamental differences arise because of the fact that in a second–best setting–when private markets are distorted–market prices typically deviate from social values. These distortions arise from several sources. One is that some markets may exhibit non–competitive behavior either on the supply or the demand sides. This arises because of the fact that scale economies lead to a relatively small number of market participants. In these circumstances, a distortion arises between the demand price (reflecting the marginal benefit) and the marginal cost. Another is that a commodity may emit an externality that accrues to third parties and that, if not priced, causes marginal social benefits to deviate from marginal social costs. Third, in some markets agents involved in transactions may be imperfectly informed about relevant characteristics or hidden actions of those on the other side of the market. Well–known examples include firms' inability to observe either the productivity or the work effort of their employees, and consumers' inability to observe the quality of products before they purchase them. This gives rise to inefficient market outcomes. Finally, government policies, such as taxes, subsidies and regulations, themselves introduce distortions in market economies. If these distortions must be taken as given by the analyst (even though in some cases, policies could be taken that might mitigate them), they will affect the evaluation of the net benefits of the project. As we shall see, one way to take market distortions into account is to devise 'shadow prices' to measure the social value of commodities sold on distorted markets rather than using the prices set on markets. This may be a difficult task because in many cases, the information needed to calculate shadow prices (such as the magnitude of externalities) may not be readily available from observed market data.

Related to this problem of evaluating project inputs and outputs when there markets are distorted is the fact that projects may have indirect effects on distorted markets elsewhere in

the economy. As emphasized by Harberger (1971a), if a market has a positive tax distortion, the amount transacted will be too low in equilibrium. Equivalently, the last unit sold will have a greater value to purchasers than to sellers. In these circumstances, if a project indirectly increases the output of this product because of complementarity relationships, there will be a net benefit on that account. This should be treated as a benefit of the project.

A third reason why private profitability might not adequately reflect social values is that some of the outputs produced or inputs used in a project may have no explicit market price attached to them. Examples include the benefits of new information produced by research and development, the value of time saved on a public transportation facility, the value of improvements in health and longevity, and the value of environmental amenities. Moreover, the government may decide to make the project's output available free of charge, as might be the case with a recreation of transportation facility. Evaluating the benefits of any of these things involves resorting to techniques to elicit the willingness to pay for the benefits by households or firms.

Fourth, and related to the last case, projects undertaken by the public sector may not be financially self-sufficient but must rely on public funds to cover their costs. In a second-best economy, the costs of raising public revenue will involve a marginal deadweight loss from the distortions of the tax system. This will be the case whether tax or debt finance is involved. Note that this cost is over and above the value of resources used in the project. The two are often aggregated into a *marginal cost of public funds* (MCPF) that reflects the full cost of raising an additional yen of public revenues.[3] The excess burden of project financing must be included as a social cost of the project. Private profitability calculations ignore this cost.

Finally, projects are typically intertemporal in nature, so the benefits and costs occur over a number of periods. In arriving at a measure of the net benefit of the project, a discount rate must be used. The discount rate used on public projects–the *social discount rate*–may differ from the discount rate used for private projects. This will occur because of either distortions on capital markets (e.g., capital income taxes) or because of externalities associated with saving for future consumption, perhaps for heirs. Of course, to the extent that benefits or costs accrue to individuals of different generations or birth cohorts, additional equity issues would arise. How should we weigh real income changes of future generations relative to those of current generations? Should we discount them, or give them the same weight as in the static SWF of (5).

All of these issues will be taken into account in what follows. We begin with a discussion of the decision rule for aggregating project or policy costs and benefits, and then turn to the evaluation of the various elements of costs and benefits.

---

[3] On the MCPF, see Browning (1976), Wildasin (1984), Usher (1986) and Sandmo (1998).

II. THE DECISION RULE

The ranking of projects involves aggregating the social benefits and costs into a single measure. In what follows, we neglect distributional weights and proceed by summing up all real income changes for individuals in the economy. In the absence of distributional weights, a yen's worth of real income is worth the same to all individuals, which simplifies our analysis considerably. We can aggregate net benefits and costs without regard to whom they accrue. In an intertemporal context, the appropriate decision rule is the *present value criterion*. In the simplest case where the social discount rate is constant over time and denoted *r*, the present value of a project or policy is written:

$$\text{PV} = \sum_{t=0}^{T} \frac{B_t - C_t}{(1+r)^t} \tag{8}$$

where B_t and C_t are the benefits and costs of the project in period *t*, which goes from 0 until the termination date of the project *T*. If this PV is positive, the project is socially profitable. More generally, policies can be ranked according to their PV, and that determines which one(s) should be undertaken.[4]

Various choice situations are possible. If the choice is simply whether to undertake a given project, then the project should be undertaken if its PV is positive. If the policymaker is faced with a choice between two or more mutually exclusive projects or sets of projects, the one with the highest PV should be undertaken. These may be given types of projects of different scales, of different starting times, or of different lengths of life. A slightly more complicated case is when there is a fixed capital budget available to allocate to a chosen set of projects. In this case, the best option is the set of projects that satisfies the capital budget and has the highest summation of PVs. This may well involve excluding projects with higher PVs if, for example, they are relatively large in size. Care must be taken to include any welfare consequences of unused funds in evaluating among options. In particular, unused funds that revert to general revenues will avoid excess burdens of taxation that would otherwise arise. As well, care must be taken to account both for capital budgets that cover several periods, and for future capital expenditures that might be needed as the project continues over time.

There are some other specific issues that must be dealt with in applying PV formula (8). One has to do with the fact that when the project is expected to terminate at time *T*, there may well be some assets still on hand. Whatever scrap value they have must be treated as a benefit and evaluated appropriately. That evaluation may be problematic, depending on whether the

---

[4] There are alternatives to the PV that use the same basic information. One is the benefit-cost ratio, which is the ration of the PV of benefits to the PV of costs. While it will indicate whether a given project is worth undertaking, it can give a misleading ranking of mutually exclusive projects of different scales. Another is the internal-rate-of-return (IRR) criterion, which calculates the discount rate that makes the PV of benefits equal to the PV of costs. If the IRR exceeds the social discount rate, the project has a positive PV and is worth undertaking. However, the IRR may not be unique, and it may also rank mutually exclusive projects incorrectly if the time profile of benefits and costs differ.

terminal capital can be put to alternative uses. If so, the present value of its subsequent benefits must be estimated. There may also be some negative assets (future liabilities) at the terminal date. For example, there may be some site clearing requirements, if there are hazardous waste materials. The cost of disposing of these must be treated as a cost, assuming that the project is responsible for creating them in the first place.

A very important issue that must be decided is the unit of measurement of costs and benefits to be used in applying (8), or equivalently the numeraire. In principle, the choice of numeraire will not affect the ranking of projects, but it will give rise to numeraire-specific procedures for evaluating benefits and costs. The issue is complicated by the fact that markets are distorted, and some numeraire choices may themselves face distortions. As well, benefits and costs occur over time, and the measuring rod must take that into account. The standard approach is to use consumption in the present period as a numeraire, where consumption can be thought of as a composite of the consumption of all goods and services measured by the real income required to purchase them as discussed earlier. There are several conceptual issues that are involved in evaluating all benefits and costs in terms of current consumption, some of which involve converting benefits and costs into consumption in the current period, and others involving discounting benefits and costs accruing in different periods.

With respect to measuring benefits and costs in terms of consumption in any given period, several issues arise, which are treated in more detail in the following sections. The first is how to evaluate project inputs and outputs in terms of current consumption when these are traded domestically on distorted markets. As we shall see, shadow-pricing techniques can be used for this purpose. A special case is that of labor used on the project, given that labor markets not only face tax distortions, but may be characterized by unemployment and the adjustment costs associated with changing employment. Special techniques that exist for shadow pricing labor must be applied. Related to this in the evaluation of inputs and outputs for which there is no market price. Here shadow pricing involves estimating the willingness-to-pay for these items in terms of current consumption. A further issue is how to evaluate commodities that are traded so that their financial benefits or costs ultimately involve a supply or demand for foreign exchange. Here, the evaluation involves converting foreign exchange amounts to domestic consumption values using a shadow price of foreign exchange that takes account of distortions that might exist in foreign exchange markets. Next, projects or policies typically are not self-financing but have effects on government revenues. If there are distortions involved in raising revenues, a yen in the hands of the government will be worth more than a yen in the hands of consumers: the MCPF, which is the opportunity cost of converting a unit of consumer income to government revenue, will exceed unity. This premium on government revenues must be accounted for in evaluating the benefits and costs of the project in terms of current consumption. Finally, if distributional considerations are an issue, the value of current consumption will differ according to whom it accrues. One approach that can be taken is to treat as the numeraire consumption in

the hands of a given income class, such as the lowest. Then, to the extent that benefits accrue to individuals in other income classes, those must be discounted using the relative values of β_i defined above. Of course, it may well be decided to avoid incorporating distributional concerns directly into the analysis, and reporting distributional effects separately instead.

Once in-period benefits and costs have been estimated, they must be converted to current consumption values. The appropriate discount rate is the consumption discount rate, that is the rate at which individuals actually discount future versus present consumption. A number of issues arise here as well. An obvious one is the choice of discount rate. A natural candidate might be the after-tax interest rate faced by individuals in the economy. However, if there are externalities associated with saving (e.g., altruism toward future generations, or externalities associated with capital accumulation), these ought to be taken into account. Similarly, if capital market constraints restrict the ability of individuals to borrow and lend, or if there is uncertainty about the future that cannot be shed on capital markets, market interest rates may not reflect true consumption discount rates. Another issue is that projects may affect the level of investment elsewhere in the economy. If, for example, some investments are crowded out, the opportunity cost of that will be the present value of forgone consumption that would have been generated by the investment, which in the presence of capital market distortions, will be greater than one per yen of forgone investment. The crowding out of private investment may come about from the way a project is financed, for example, whether by tax or debt finance. Finally, to the extent that the consumption of future generations is affected, the question arises as to what weight should be attributed to it when evaluated against current consumption. In other words, what distributive weight should be attached to the consumption of future generations. This will be important with respect to long-lived projects as well as with respect to fiscal policy changes that have intergenerational impacts, such as major tax reforms or reforms of social transfers.

Another issue that arises in aggregating benefits and costs over time is the manner of accounting for price level inflation. Here, what is important is consistency. Two alternative-and equivalent-procedures exist. One is to use nominal prices in each period, that is, the prices that obtain on markets. In this case, the discount rate that must be used is the nominal interest rate, which incorporates the effect of inflation. The other is to use constant-yen prices, which are obtained by deflating nominal prices by a price index that reflect the rate of inflation. Following either procedure will generate the same results. To see this, consider the PV formula given by (8). Suppose that benefits and costs are measured in nominal terms, and that the interest rate r is a nominal one. If the inflation rate is π per period (assumed constant here), converting nominal benefits and costs to real ones in terms of period 0 prices involves the following transformations:

$$b_t = \frac{B_t}{(1+\pi)^t}, \quad c_t = \frac{C_t}{(1+\pi)^t}$$

As well, if we denote the real interest rate by *i*, it must satisfy:

$$(1+i) = (1+r)(1+\pi)$$

Then, it can be seen that the PV formula (8) can equivalently be expressed in real terms as follows:

$$PV = \sum_{t=0}^{T} \frac{c_t - c_t}{(1+i)^t} \tag{8'}$$

Of course, either approach involves predicting values into the future, whether they be constant- or current-yen values. All cost-benefit analyses face this problem.

Some accounting issues also arise in the costing of capital inputs because of their durable nature. Their costs are incurred up front, but their use is spread over the future. Standard accounting principles suggest two alternative approaches to account for capital costs, cash-flow accounting and accrual accounting. Under cash-flow accounting, all outlays and inflows are costed as they occur. Capital expenditures are fully expensed when they are made. These include all investment expenditures – additions to a project's capital stock, replacement and depreciation spending, and scrap value salvaged when the project terminates. Costs of financing and ongoing depreciation are not included since that would be double counting, except to the extent that depreciated capital is replaced.

Accrual accounting costs capital as it is used rather than when it is acquired. Two costs are included, depreciation and financing costs. Depreciation includes the amount of capital value that is used up each period, and comprises obsolescence, wear and tear, and changes in the relatively price of the capital. Financing costs reflect the forgone interest associated with holding capital rather than investing the funds into financial assets.

Cash and accrual accounting for capital costs are alternative ways of presenting the same information, and ought not to be mixed: otherwise, double-counting or under-counting will occur. If properly done, the present value of the accrual costs of an investment should equal its cash flow. Accrual accounting is by its nature more difficult since it requires knowledge of depreciation, which cannot readily be observed from market prices. As well, the principles of shadow pricing discussed below are much less transparent under accrual accounting. Hence, cash accounting is typically used for cost-benefit analysis, contrary to the practice in the private sector. Accrual accounting is the norm for private firms since it conveys more information to shareholders of the current profitability of the firm.

The above discussion summarizes the issues that arise when current period consumption is chosen as the numeraire. Shadow prices for inputs purchases and outputs sold must take account of distortions on domestic markets; net foreign exchange earned from transactions involving tradables must be converted to consumption equivalents using a shadow price of foreign exchange; changes in government revenue must be converted to consumption equivalents using a MCPF; and future net benefits, including those affected by crowded-out investment,

must be discounted using a consumption discount rate. The result is the PV of a project measured in terms of consumption or real income of domestic individuals, with or without distributional considerations having been incorporated.

There are other options for the numeraire. One that has been prominent in the development literature is foreign exchange in the hand of the government. This was the choice proposed by Little and Mirrlees (1968), and developed further by Ray (1984). In this case, the value of all inputs and outputs are converted into equivalent values of foreign exchange using conversion factors that are analogous to the above shadow prices. Traded commodities can be evaluated at their foreign exchange costs (so-called *world prices*). Non-traded commodities are then converted into equivalent values in terms of foreign exchange using the inverse of a shadow price of foreign exchange. The use of this numeraire entails the use of analogous conversion factors as above, but the conversion is typically done in the reverse direction.   Thus, any changes in output or use of non-traded commodities must be converted into foreign exchange using effectively the reciprocal of the shadow price of foreign exchange. Changes in domestic consumption changes are discounted relative to foreign exchange in the hand of the government because of the MCPF factor already mentioned. It was argued that this MCPF exceeds unity not just because of tax distortions, but also because in a developing country context, government funds might be an important source of finance for investment, which has higher value than consumption because of capital market distortions. The Little-Mirrlees approach also recommended incorporating distributional weights into the cost-benefit evaluation, it being the case that policy instruments were not as readily available in developing countries for addressing equity concerns directly. But in the end, nothing of real substance is involved in choosing between current consumption and foreign exchange in the hand of the government as the numeraire. If used consistently, both should yield the same project rankings.

In what follows, we follow the convention of taking current consumption as the numeraire. We take up some of the issues involved in measuring costs and benefits in terms of current consumption in more detail.

III.   VALUING MARKETABLE INPUTS AND OUTPUTS

Cost-benefit analysis involves putting social values on the net benefits of projects or policies in distorted economies. Before turning to the specific of particular types of costs and benefits, it is useful to put things into context by outlining a general expression for welfare change measures in a distorted economy. The expression we develop was proposed by Harberger (1971a), so we refer to it as *Harberger's measure of welfare change*. It is a measure that ignores the distributive effects of policies, so we can base our discussion on an economy consisting of a representative household.

Let the representative consumer have a utility function $u(x_0, x_1, x_2, \cdots, x_n)$ where x_i is the

consumer's purchase of commodity i (which will be negative for factors supplied). The consumer's price for commodity i is p_i, where $p_0 = 1$ for the numeraire good x_0. Differentiating the utility function and using the first–order conditions for the consumer's utility maximization problem, $u_i / u_0 = p_i / p_0 = p_i$, we obtain an expression for the change in utility from a change in consumption measured in terms of the numeraire good:

$$dw = \frac{du}{u_0} = \sum_i \frac{u_i dx_i}{u_o} = \sum_i p_i dx_i \tag{9}$$

Next, suppose consumption is supplied from either the private market, y_i, or from public sector projects, z_i, so that market clearing implies $x_i = y_i + z_i$. Private sector production is implicitly determined by a transformation function $f(y_0, y_1, y_2, \cdots y_n) = 0$. Assuming there is production efficiency, the producer price for commodity i in terms of the numeraire good can be expressed as $f_i / f_0 = p_i - t_i$, where t_i is the distortion in market i. We can think of this as a tax distortion, but it could also be an externality or a distortion due to monopoly power. Differentiating the transformation function, and using the expressions for relative producer prices, we obtain:

$$\sum_i (p_i - t_i) dy_i = 0 \tag{10}$$

Then, combining (9) and (10) and using the market clearing condition, we obtain:

$$dw = \sum_i t_i dx_i + \sum_i (p_i - t_i) dz_i \tag{11}$$

This is the general expression for the change in welfare arising from any small change in consumer demands and project production. The former are evaluated at the size of the distortion, while project inputs and outputs are evaluated at their producer prices. It is useful for what follows to recognize that projects will only operate in some markets. Let k index the markets on which a project under consideration operates and j index all other markets. Then, (11) may be written:

$$dw = \sum_k t_k dx_k + \sum_k (p_k - t_k) dz_k + \sum_j t_j dx_j \tag{12}$$

Thus, if we are evaluating a project that involves a change in the z_k's, (12) indicates that we must take account of the direct opportunity cost of the project's inputs and outputs, indirect changes in welfare on the markets on which those inputs and outputs are traded and welfare changes in other markets. In the following we first consider first the welfare effects arising from changes in outputs on the markets in which the project operates–the first two terms in (12)–and then the indirect effects–the third term. We then look at issues that arise in particular markets.

III.1. *DIRECT WELFARE EFFECTS ON MARKET INPUTS AND OUTPUTS*

We consider here the welfare effects associated with the purchase of a commodity on a

distorted market, where the welfare effect includes both the value of the resources transferred to the public sector and changes in the cost of the distortion itself. The technique used is to construct a shadow price that takes account of both of these effects. In terms of the general welfare change measure (12), the effect of a change in the output of a commodity z_k by a project can be expressed as:

$$\frac{\partial w}{\partial z_k} = \underbrace{p_k - t_k + t_k \frac{\partial x_k}{\partial z_k}}_{\text{direct effect}} + \underbrace{\sum_{j \neq k} t_j \frac{\partial x_j}{\partial z_k}}_{\text{indirect effect}} \tag{13}$$

Here we consider what is labeled the direct effect, which we can think of as the shadow price of commodity k, denoted s_k. We can obtain a simplified expression for this shadow price by adopting a partial equilibrium approach and supposing that market supplies and demands depend only on own-prices. Then, the market clearing condition can be written $x_k(p_k) = y_k(p_k - t_k) + z_k$. Differentiating this, we obtain:

$$\frac{\partial p_k}{\partial z_k} = \left(\frac{\partial x_k}{\partial p_k} - \frac{\partial y_k}{\partial p_k} \right)^{-1} = \left(x_k' - y_k' \right)^{-1}$$

Then, since $\partial x_k / \partial z_k = x_k' \partial p_k / \partial z_k$, the shadow price of z_k (corresponding to the direct effect above) is:

$$s_k = \frac{(p_k - t_k) y_k' - p_k x_k'}{y_k' - x_k'} \tag{14}$$

This is known as Harberger's shadow pricing rule, following Harberger (1969). It stipulates that the shadow price is a weighted average of the supply and demand prices, where the weights are the proportions in which an increase in z_k comes from increased supply and reduced demand. For discrete changes in project demand, it can be written in the following approximate form:[5]

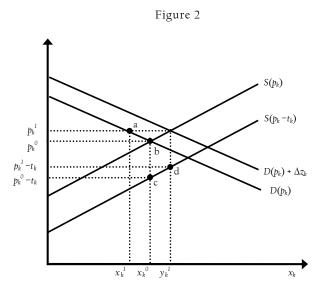$$s_k = (p_k - t_k) \frac{\Delta y_k}{\Delta z_k} - p_k \frac{\Delta x_k}{\Delta z_k} \tag{15}$$

To illustrate the meaning of this shadow price, refer to Figure 2. Suppose before the project is introduced, market output is $x_k^0 = y_k^0$, with the associated consumer price p_k^0 and producer price $p_k^0 - t_k$. The project demand is Δz_k, which causes supply to rise to y_k^1 and demand to fall to x_k^1. The opportunity cost of the project demand consists of loss in value to consumers of the area beneath the demand curve $x_k^1 a b x_k^0 \cong -p_k^0 \Delta x_k$, while the opportunity cost to producers is the relevant area under the supply curve $y_k^0 c d y_k^1 \cong (p_k^0 - t_k) \Delta y_k$. Thus, the opportunity cost per unit of z_k, that is, its shadow price, is given by (15). Notice that if the supply curve is horizontal, the shadow price is the supply price $p_k - t_k$, while if the demand curve is horizontal, the shadow price is the demand price p_k. The former case might be thought to be relevant where the commodity is tradable and the supply price is dictated by the world price. However, in this case, an increment in project supply results in an increment of foreign exchange earnings. To the extent that there are distortions in trade (tariffs, ex post subsides, quotas, etc.), the conversion

---

[5] If distribution is an issue, the weights in this shadow pricing rule could be augmented by distributive weights in a manner discussed earlier.

of foreign exchange earnings to domestic consumption values will itself involve some shadow pricing. We return to that issue below.
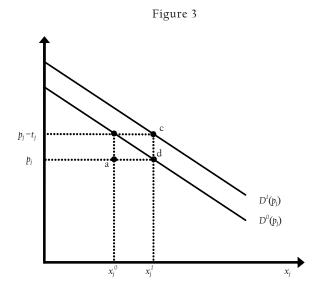
Figure 2

III.2. INDIRECT WELFARE EFFECTS ON RELATED MARKETS

To the extent that the change in project demand for commodity k induces changes in the market output of commodities in other distorted markets, an indirect welfare effect must be taken into account. This indirect effect is captured in the last term in (13), which can be approximated for discrete changes to:

$$r_k = \sum_{j \neq k} t_j \frac{\Delta x_j}{\Delta z_k} \tag{16}$$

The 'augmented shadow price' of z_k would then be the sum of the direct and indirect effects, $s_k + r_k$, where s_k is given by (15).

This indirect effect can be depicted geometrically for the case of one of the related markets in Figure 3. Here, the market for the related commodity x_j has a negative distortion, so $t_j < 0$. This could be a subsidy or a negative externality causing the marginal value to consumers to be less than the marginal cost on the supply side (the marginal social cost in the case of an externality). In the initial equilibrium without the project in place, output is x_j^0 and demand and supply prices are p_j^0 and $p_j^0 - t_j > p_j^0$, respectively. The marginal distortion–the difference between the opportunity cost to suppliers and the value to demanders–is given by t_j. Since the marginal distortion is constant, the welfare gain from any change in x_j is given by $t_j \Delta x_j$, as (16) implies. In Figure 3, the supply price is assumed to be constant for simplicity, and the change in project demand Δz_k is assumed to cause the demand curve for x_j to shift right. Thus,

if $\Delta z_k > 0$, so $\Delta x_j < 0$, it would be the case that $\Delta x_j > 0$ if the two commodities were substitutes in demand. The equilibrium output of good j rises to x_j^1. The welfare change in the market for x_j consists of two components: the increase in the costs of providing the extra x_j given by area $x_j^0 b c x_j^1$ and the increase in the value of the additions output to consumers, $x_j^0 a d x_j^1$. The difference between these is the area *abcd*, which corresponds with the indirect effect $t_j \Delta x_j$ that appears in (16), in this case an additional welfare cost. The augmented shadow price would have to include the analogous indirect effects from all distorted markets that are affected.

Figure 3

Indirect effects of this sort have been important in the evaluation of transportation projects, such as a subway system, an airport, or a major bridge or tunnel. One of the main effects of these kinds of projects, in addition to creating new demand, is to divert traffic from other modes of transport, and often these other forms have distorted prices. For example, public transport projects may divert traffic from road travel, where price (the cost of a trip on the road) will be less than social cost, which includes the congestion imposed on other travelers. In this case, the indirect effect would be a benefit associated with the reduction in traffic on congested roadways. These indirect benefits can be among the most important benefits of transportation projects. Moreover, the existence of indirect benefits arising from diverted traffic can lead to arguments for subsidizing public transit usage as a way of increasing the amount of traffic diverted. This is a classic second−best pricing argument.

III.3.   *SHADOW PRICING OF PARTICULAR INPUTS*

The above discussion of the shadow pricing of inputs and outputs applies in general to any

commodities. There are a few special examples that tend to be of particular importance. Here, we consider two special cases: labor and foreign exchange. When we turn to intertemporal issues below, the weighted average methodology will reappear.

III.3.1.   *The Shadow Wage Rate*

Virtually all projects employ labor as a major input, and labor markets are typically distorted. Taxes of various sorts–both direct and indirect–will drive a wedge between the price firms pay workers and the after–tax wage that workers receive. There may also be unemployment in labor markets, and there may be rigidities that preclude costless movement of labor from one location or occupation to another. Shadow pricing of labor should take all these factors into account.

If the only distortion were taxation, the weighted–average principle would apply directly. The shadow wage rate would be a weighted average of the before–tax wage rate and the after–tax wage rate, where the weights are the proportions in which the project labor comes from workers previously employed elsewhere (forgone demand) and workers induced to enter the labor force (increased supply).

More generally, where there are other sources of distortions, these would need to be taken into account. Begin with the case of involuntary unemployment. If this exists, workers hired by a project can come from three sources: those who would be employed elsewhere, those who are attracted into the workforce from voluntary unemployment, and those who are involuntarily unemployed. The shadow wage will be a weighted average of the opportunity cost of each of these. In the case of workers employed elsewhere, the opportunity cost is the before–tax wage rate, and for those voluntarily unemployed, it is the after–tax wage rate, as above. For the involuntarily unemployed, they would have been willing to work for the going wage rate but are unable to find a job. Their opportunity cost is less that the after–tax wage rate (since they would be willing to work at the going wage), but greater than the value of leisure (discussed further below). Since the true opportunity costs differs for different workers, and since it is not observed on the market, some arbitrary choice must be made, perhaps midway between the before–tax wage and an estimated value of leisure time, which is presumably above zero.

Matters are more complicated once one takes account of the fact that involuntary unemployment might be an equilibrium phenomenon. Consider the following simple example, taken from the literature on cost–benefit analysis in a developing country context. Suppose there are two sectors in the economy, a rural one and an urban one, and a wage differential exists between them. In particular, the urban wage rate w_u exceeds the rural wage, w_r. There might, for example, by severe underemployment in the rural sector because of, say, technological improvement making labor redundant, but nonetheless families continue to employ family members at subsistence wage rates. A naïve application of shadow pricing might take the shadow wage rate to be a weighted average of w_u and w_r, with the weights being the

proportion of workers for a project drawn from the urban and rural sectors. Given underemployment in the rural sector, the latter may be close to zero, as observed by Dasgupta, Marglin and Sen (1972) and Little and Mirrlees (1974).

However, the persistence of the differential between w_u and w_r might reflect an equilibrium of sorts in the labor market, in which case the simple shadow pricing rule must be amended. One form the equilibrium might take has been proposed by Harris and Todaro (1970), and used by Harberger (1971b) to derive the shadow price. Suppose that w_u is above the market clearing level for some institutional reasons (for example, efficiency wages, turnover costs, etc.). This leads to an unemployment rate denoted by π. If jobs are allocated randomly to the unemployed, if workers are risk-neutral, and if there are no costs of migration, workers will allocate themselves between the urban and rural sectors until the expected wage is equalized, or $w_r = \pi w_u$. Consider now a project that creates jobs in the urban sector. Each job that is filled will cause $1/\pi$ workers to migrate to the urban sector to maintain equilibrium. The opportunity cost of these workers is w_r/π, or equivalently, w_u. So the shadow wage rate is just the wage actually paid to hire them: no weighted average shadow price is required. The same argument can be seen to apply if the project is in the rural sector. The shadow wage would then be w_r.

Of course, this procedure would have to be suitably amended if there were other distortions, like taxes and subsidies, or if distributional concerns were included in the shadow wage. However, note that migration costs would not affect the argument. If migration costs were m per worker, labor market equilibrium would be $w_r + m = \pi w_u$ (assuming migration is from the rural to the urban sector). Now the opportunity cost of each worker attracted to the urban sector would be $w_r + m$. Again, each urban job created would attract $1/\pi$ workers, whose opportunity costs is $(w_r + m)/\pi = w_u$. Thus, the shadow wage rate again equals the wage rate actually paid.

III.3.2.   *The Shadow Exchange Rate*

Suppose now that a project input is a tradable commodity whose world price is taken as given from this country's point of view. Any project demand will result in an increase in the demand for foreign exchange. Suppose that trade is distorted by trade taxes (import tariffs or export taxes). To begin with, and to illustrate the point, suppose that all imports face a common tariff rate τ. Given that the world prices are fixed, we can think of aggregating all traded commodities into a composite commodity, which we can simply refer to as foreign exchange, whose price is e, the exchange rate. The quantity of exports consists of the amount that can be sold for one unit of foreign currency. The supply curve for exports is then the supply curve of foreign exchange as a function of the exchange rate, denoted $S(e)$. It will be an upward-sloping curve: the greater the price of foreign currency, the greater the value in the domestic currency per unit of exports, and therefore the greater the supply of exports by domestic producers. The

domestic price of composite imports will be $(1+\tau)e$, since all imports bear a common tariff. Then, the domestic demand curve for important will be $D((1+\tau)e)$, which will be downward sloping.

The analog of (15) will apply. The shadow price of foreign exchange is given by:

$$s_e = e\frac{\Delta S}{\Delta z_e} - (1+\tau)e\frac{\Delta D}{\Delta z_e}$$

where Δz_e is the net demand for foreign exchange by the project. A geometric interpretation analogous to that in Figure 2 applies directly. Notice that the shadow price of foreign exchange typically exceeds the market exchange rate, $s_e > e$. Thus, a cost–benefit analysis will discourage the use of traded commodities.

The analysis can readily be extended to the case where there are different tariff rates on different commodities. Suppose the tariff rate τ_i applies to import i, whose demand is D_i. The shadow price of foreign exchange can then be written:

$$s = \frac{e\Delta S}{\Delta z_e} - \frac{\sum_i (1+\tau_i)e\Delta D_i}{\Delta z_e}$$

The case of export taxes could also be included. Moreover, the shadow price of foreign exchange could in principle incorporate other distortions in trade, such as quotas or exchange rates that were out of equilibrium due to being less than completely flexible.

IV.   THE EXCESS BURDEN OF PROJECT FINANCING

A typical feature of public projects is that they are not self–financing. Their costs will be financed either by taxes or borrowing (future taxes). In either case, there are welfare costs associated with them that must be taken into account in evaluating projects. Since somewhat separate issues arise with respect to taxation and borrowing, we consider them in turn. To the extent that there is an excess burden arising from project financing, that excess burden must be treated as a cost over and above the shadow pricing of costs that we have already discussed. Note, however, that only the excess burden must be included, not the full cost of the financing. Nonetheless, the methodology outlined below discusses the full costs of financing, both the resources transferred to the project and the excess burden. Care must be taken not to double count the costs.

IV.1.   THE COSTS OF TAXATION: THE MARGINAL COST OF PUBLIC FUNDS

If a project is financed by taxation, the cost of finance is the opportunity costs of raising additional revenues in an already distorted economy. Consider the following simple partial equilibrium model to illustrate the point. The utility function of the representative household is $u(c,\ell,z) = c + h(1-\ell) + b(z)$, where c is composite consumption, ℓ is labor supplied and z is the project output. The function $h(1-\ell)$ is the utility of leisure (assuming households have one unit

of time to allocate to leisure and work), and is increasing and strictly concave. The wage rate is w, and government imposes a tax at the rate t on labor, assumed for simplicity to be a per unit tax. Thus, the household, taking t and z as given, solves the following problem:

$$\text{Max}_{\ell} \quad (w-t)\ell + h(1-\ell) + b(z)$$

The first-order condition is $h'(1-\ell) = w-t$, which yields the labor supply function $\ell(w-t)$. The value function for this problem is the indirect utility function $v(w-t,z)$, which has the following properties by the envelope theorem:

$$v_t \equiv \frac{\partial v}{\partial t} = -\ell(w-t), \quad v_z \equiv \frac{\partial v}{\partial z} = b'(z) \tag{17}$$

Using this, we can obtain an expression for welfare change by differentiating the indirect utility function to give:
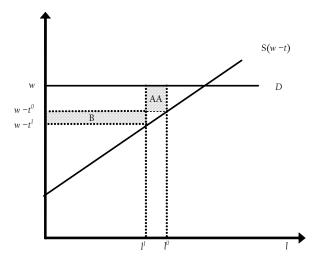
$$dv = v_t dt + v_z dz = -\ell dt + b'(z)dz$$

Feasible changes in t and z must satisfy the government budget constraint: $t\ell(w-t) = z$. Differentiating this we obtain $dt = dz/(\ell - t\ell')$, so the expression for welfare change may be written:

$$\frac{dv}{dz} = \left[b'(z) - \frac{\ell}{\ell - t\ell'} \right] = \left[b'(z) - \frac{1}{1 - t\ell'/\ell} \right] = \left[b'(z) - \frac{1}{1 - \tau\varepsilon} \right] \tag{18}$$

Where $\varepsilon = (w-t)\ell'/\ell$ is the elasticity of labor supply and $\tau = t/(w-t)$ is the ad valorem labor tax rate. The term $(1 - \tau\varepsilon)^{-1}$ is the marginal cost of public funds (MCPF). The cost–benefit analysis rule (18) says that the net benefit of the project consists of its benefits less its costs, where the latter is the MCPF.

Figure 4

The MCPF can be given a geometric interpretation. Figure 4 depicts the labor market with the pre-tax wage assumed to be fixed and the labor supply curve upward sloping. In the initial situation, the after-tax wage is $w - t^0$ and labor supply is ℓ^0. When the tax is increased to t^1, labor supply falls to ℓ^1. The MCPF is the total cost per unit of revenue raised. The increment of revenue raised from the tax increase is $B - A$. The cost is the value of resources transferred, $B - A$, plus the increase in deadweight loss, which is approximately A, for a total of B. Thus, the MCPF is:
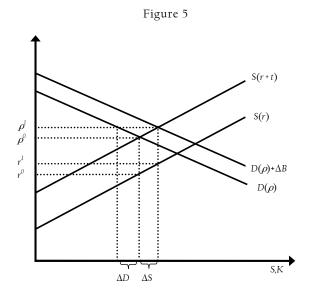
$$\text{MCPF} = \frac{B}{B - A} = \frac{1}{1 - A/B} = \frac{1}{1 - t\Delta\ell/(\ell\Delta t)} \cong \frac{1}{1 - t\ell'/\ell}$$

Of course, if there are many different taxes in place, the MCPF expression becomes correspondingly more complicated.

IV.2.   THE OPPORTUNITY COST OF BORROWED FUNDS

Now consider the case in which some project financing comes from borrowing. If there are distortions in capital markets, that will affect the opportunity cost of the borrowing. Suppose the rate of return on investment, denoted ρ, exceeds the after-tax return on saving, denoted r, because of taxes levied on capital income. These could include both personal taxes and taxes on firms, and we aggregate them into a single tax rate t. (At this point we assume that the economy is closed, a point to which we return below.)

Figure 5 can be used to illustrate the effects of borrowing an amount ΔB. The demand for funds is labeled D: it depends on the before-tax rate of return ρ. The supply of funds, which

Figure 5

depends on r, is the curve S. The project finance shifts the demand curve to the right, displacing some private demand, ΔD, and inducing some additional supply, ΔS. The opportunity cost of these changes constitutes the opportunity cost of borrowing.

Following Marglin (1963) and Feldstein (1972a), the increment of saving corresponds with forgone consumption, so its opportunity cost is simply the amount saved ΔS. The opportunity cost of the forgone investment is the stream of output the investment would have yielded. If the future output would have been consumed, the evaluation of the opportunity cost is straightforward. With a rate of return on investment given by ρ, the return is equivalent to a perpetual stream of output ρ, which has a present value of $\rho/r > 1$. If all returns to investment are consumed, this is the present value of the stream of forgone consumption from the displaced investment. Therefore, we can say that the marginal cost of borrowing, denoted MCB is given by:

$$\text{MCB} = \frac{\Delta S}{\Delta B} - \frac{\rho}{r}\frac{\Delta D}{\Delta B} \tag{19}$$

Thus, for every yen of borrowing, $\text{MCB} - 1$ would have to be added as the excess burden arising from debt finance. Conversely, when the debt is paid down, the reduction in excess burden remaining would have to be deducted (although this may well be offset by an increase in the excess burden of taxation that is used to repay the debt).

The simplicity of (19) is due to the assumption that all investment returns are consumed. If some of them are reinvested, the benefits of that in terms of the stream of consumption generated would have to be taken into account. For example, suppose that a fixed proportion σ of project returns are reinvested at the rate of return ρ. Then, the capital value of one yen of initial investment will grow at the rate $\sigma\rho$ per period. Assuming continuous time for simplicity, the asset value will be $e^{\sigma\rho t}$ at time t, which will generate $(1-\sigma)e^{\sigma\rho t}$ in consumption. In this case, the present value of consumption generated by an initial yen of investment will be $(1-\sigma)\rho/(r-\sigma\rho)$, so the MCB can be expressed as:

$$\text{MCB} = \frac{\Delta S}{\Delta B} - \frac{(1-\sigma)\rho}{r-\sigma\rho}\frac{\Delta D}{\Delta B}$$

Other illustrative examples could be considered, but the principles are clear.

The above discussion of the size of the MCB was for the (unrealistic) case of a closed economy. In reality, capital markets are open, and many countries might face highly elastic supplies of international savings, so the cost of borrowing is effectively fixed. This complicates the determination of the MCB slightly since now there are three sources of project borrowing: forgone domestic investment, increased domestic saving and foreign lending. Different distortions might apply to different sources. A tax on investment will imply that the opportunity cost of forgone investment will exceed unity for the same reasons as discussed above. If there is a tax on savings, there will be a further distortion on domestic investment as well as a distortion on capital inflows. The opportunity cost in terms of domestic consumption will exceed one on this account. Moreover, there will also be a requirement to shadow price any

foreign exchange that is used to finance project borrowing.

It should be noted that in the literature, other ways of taking account of capital market distortions have been proposed. A well-known one is that proposed by Harberger (1969). He suggests using a weighted-average discount rate to take account of capital market distortions, rather than to treat the excess burden of those distortions as a cost of the project itself. The discount rate used in the PV formula (8) would be a weighted average of ρ and r, where the weights are the proportions of project financing coming from forgone investment and increased savings. This method will not generally give the same project rankings as the MCB method described above. For further discussion of this, see Feldstein (1972a), Boadway (1978), and Boadway and Bruce (1984).

V. NON-MARKETED INPUTS AND OUTPUTS

Some project outputs may be of the nature of services or benefits that accrue free of cost rather than being subject to market transactions. Transportation projects may save travel time or reduce environmental degradation; research expenditures produce new knowledge that is freely available to individuals and firms; health expenditures may reduce the incidence of disease and loss of life; environmental expenditures may improve amenities in forests and parks; and education and training programs may improve skills of participants. Since no prices are available to guide the evaluator, some other method must be used to attribute values to these non-marketed or intangible benefits (or costs). The principle is straightforward: benefits should be valued at the willingness-to-pay, that is the amount of consumption or real income that households would be willing to forgo in order to receive the benefits.

In the absence of market prices, other means must be used to infer willingness-to-pay. The procedure used will vary according to the nature of the non-marketed benefit. In some cases, indirect evidence of pricing elsewhere may be used. For location-specific benefits, the effect of the project on property values may be used. In some cases, hedonic pricing methods can be used, where households' behavior in other contexts may be used to infer the value of an item. Finally, survey techniques can be used to ask households anonymously what value they place on an item, taking care to ensure that the survey not affect their incentive to misreport. We proceed by considering various examples that are commonly found in practice.

V.1. *The Value of Lives Saved*

Many public projects have as one of their effects a reduction in lives lost, or more generally, improvements in safety of various sorts. This is sometimes misleadingly referred to as 'the value of life'. From the point of view of the users, it is more accurate to view it as the value of a reduction in the risk of death. Since there are many examples elsewhere of households having to

choose among options that involve different risks of death, it is natural to use hedonic pricing techniques to infer what they might be willing to pay to reduce the risk of death (or injury).

The monetary value to be attached to a reduction in the risk of death or injury should in principle be the willingness-to-pay for such a reduction by the households potentially involved. Note that this approach involves estimating an ex ante willingness-to-pay for reducing the risk of death, rather than evaluating the willingness-to-pay for eliminating the certainty of dying. Presumably the latter would be indefinitely high.[6] For public projects, this might be regarded as inappropriate. It would be known with some degree of certainty how many lives will be saved by, say, road safety improvements even if it is not known precisely whose lives will be saved. Thus, some might argue that using the willingness-to-pay as an ex ante measure of the value of a reduced risk of death may not be appropriate. In any case, some account must be taken of the reduction of lives lost, and we adopt the common approach of evaluating that as an ex ante willingness-to-pay.

If one takes this approach, it may not be necessary to value the willingness-to-pay explicitly: it may be implicitly taken account of to the extent that households can voluntarily choose to use the project. For example, if a new highway or an improvement to an existing one reduces the risk of accident, the value of that reduction is implicit in the demand curve for the use of the highway. It measures the marginal values placed on highway use net of the costs, which include the risk of injury as well as other intangible costs like the saving of time (to be discussed next).

Where an explicit estimate must be made of the reduction in the risk of death, the procedure is to infer it from other situations in which households must choose among alternatives that involve different risks of death. One such situation is the choice of workers among jobs with different probabilities of dying on the job (e.g., mining versus bookkeeping). With the appropriate data, econometric techniques can in principle be used to estimate the amount of money persons need to be compensated for to accept an increase in the risk of death. Of course, it may be difficult empirically to disaggregate wage differentials into various factors including the risk of death on the job. As well, persons may differ according the their 'risk aversion' for death, so the differential attributed to high-risk jobs might underestimate the average cost of the risk of death. Less risk-averse persons will gravitate towards higher risk jobs, and the differential needed to compensate them will be less than average.

V.2.   Value of Time Saved

The construction of transportation facilities including subways, expressways, airports,

---

[6] Actually, income effects are important here. The amount one would be willing to pay to avoid losing one's life would be less than the amount they would be accept to accept losing their life. The former is limited by the wealth they have. For a general discussion of evaluating lives saved, see Jones-Lee (1976).

tunnels and bridges are done to facilitate the movement of people and goods from one place to another in a timely fashion. Some of the traffic they attract is newly generated, while some is diverted from other means, at least partly to reduce travel time and increase convenience and comfort. In evaluating these projects, the benefits of time saved as well as comfort and convenience must be attributed. We focus on time saved, but similar techniques can be used for comfort and convenience.[7]

As with shadow pricing procedures outline above, the value of time saved in traveling depends upon the alternative uses to which the saved time will be put. It may simply allow more leisure time to travelers, or it may be used for productive work. In the latter case, valuation is in principle straightforward. To the extent that working time and travel time are perfect substitutes, and wages correctly reflect productivity, the value of time saved can be estimated using wage rates.

If time saved is devoted to leisure, matters are more complicated. One might imagine that a household attributes different values to working time, MBW, leisure time, MBL, and commuting time, MBC. While commuting time might be fixed to the household, the division of the remainder of time between working and leisure would be done so as the marginal values of each activity are the same: $MBL = w + MBW$, where w is the wage rate. This implies that $MBL < MBL$. If the project increases leisure time, the value of time saved is then $V = MBL - MBC$. Since V cannot be observed directly, it must be inferred. The hedonic approach would be to do so by observing how much people are willing to pay to save time in other contexts where such choices are available. For example, there may be circumstances in which people have a choice between using different modes of transport to travel to and from given destinations. Data of these sorts can be used to estimate the value of time saved traveling. This can be used to value time saved traveling by traffic that is diverted from other modes of transport. For newly generated traffic, the value of time is implicit in the estimate of demand curves for new traffic, to the extent that those can be estimated. (This is analogous to the values of reductions in the risk of death by voluntary project users discussed above.)

Similar techniques can be used to estimate improvements in comfort and convenience associated with new or improved transportation facilities. They can also be used to estimate the benefits of other investment, such as recreational facilities, where access to the facilities takes time. The value of using these facilities can be estimated by using, among other things, the amount of time households are willing to take to gain access to them.

V.3.   *Costs of Pollution*

Projects that are specific to a given location might impose pollution costs on neighboring

---

[7] A good outline of the issues may be found in Harrison and Quarmby (1972).

residents. An airport will increase noise levels in the vicinity. An industrial project may cause air pollution. Estimates of the cost of such pollution might be obtained indirectly from property values. The costs of pollution should be capitalized into residential property values. Empirically, hedonic pricing techniques can in principle be used to obtain monetary measures of the costs of pollution. As before, data would have to be sufficient and care would have to be taken to ensure that one could estimate the effect of the pollution on property values, separate from the other things that determine them. But the principles are clear.

V.4.   Contingent Valuation

A major problem with estimating the effects of intangibles using hedonic pricing techniques, whether it be the value of reductions in the risk of injury and death, the value of time saved or comfort and convenience or the cost of pollution, is that sufficient data may not be available to obtain reliable estimates. In these circumstances, survey methods might be used to generate suitable data. Instead of relying on observed choices to estimate the value of non-marketed benefits and costs, households could be asked to reveal them directly through a survey. The survey could ask how much households are willing to pay for the intangible in question or, alternatively, how much they would be willing to accept to give up some intangible benefit or to accept some cost. To the extent that the survey was complete and households responded truthfully, this would yield estimates of their willingness-to-pay.

The survey would have to be designed to ensure that respondents fully understand the nature of the intangible being evaluated. Those who are surveyed must include a representative sample of those who affected by the project. As well, those who voluntarily accept to fill in a survey must not represent a biased group, such as those who feel most strongly about it. And, as mentioned, responses must be truthful. Those who feel strongly about an issue will have an incentive to exaggerate their willingness-to-pay. As with hedonic pricing techniques, contingent valuation methods must be used and interpreted with due care.

We now turn in the final sections to issues that arise in an intertemporal context.

VI.   THE SOCIAL DISCOUNT RATE

The flow of benefits and costs will be measured in terms of the numeraire in each period. Presuming the numeraire is the value of consumption, the appropriate rate for discounting future net benefits to the present is a consumption rate of interest. Given our assumption that it is the willingness-to-pay of households that determines project benefits and costs, the appropriate rate of discount for aggregating benefits and costs over time is the rate at which households discount present versus future consumption. If we set aside equity considerations for the time being, so that no redistributive weight is attached to future consumption, one might

suppose that the discount rate for public projects–the *social discount rate*–is the interest rate at which households could borrow and lend. This would be the case if capital markets were perfect, if all households faced the same interest rate, and if there were no externalities arising from household saving. The common interest rate would be the rate at which households themselves discount future versus present consumption. If there are taxes on capital income, the appropriate consumption discount rate would be the after–tax return that households obtain on their savings (or pay on their borrowing).

If one or more of these assumptions are not met, the determination of the social discount rate will be more complicated. Some of the more important ones are considered in what follows.

VI.1. *Heterogenous Household Discount Rates*

There are many reasons why the rate used by different households to discount future consumption may differ. Even if there were no constraints on borrowing or lending, after–tax interest rates will differ if the tax system is progressive or if different types of assets face different tax rates, both of which are common in practice. As well, because of costs of intermediation, borrowing and lending rates may differ. Finally, if capital markets are imperfect, households may be restricted in the amounts they can borrow, and the terms on which they can borrow may differ according to their incomes or wealth. If there are liquidity constraints, household borrowing will be rationed, implying that their discount rate will be higher than the market interest rate. Finally, capital markets may not be perfectly competitive, and they may be plagued by problems of asymmetric information.

In all these circumstances, there is no unique consumption discount rate. Since it is impractical to disaggregate project benefits and costs by household type, in practice some compromise discount rate must be used, such as an average of after–tax returns to saving.

VI.2. *Saving Externalities*

Saving may be done for various reasons. One is to smooth consumption out over one's lifetime. Another might be to pass on some of one's wealth to one's heirs. Yet another may be for precautionary purposes, such as to self–insure against uncertainty in the length of life or one's future health. In all these cases, there may be external benefits generated for third parties from whom no compensation is received.

One source of externalities is associated with the investment that savings are used to finance. To the extent that investment provides benefits to society at large, such as through the generation of new knowledge or the accumulation of skills through experience, the future consumption generated from one's savings will be greater than one obtains oneself. In these circumstances, saving generated in a market economy would tend to be too small, leading to an

argument for policies that encourage savings. It is not at all obvious that this has implications for the social discount rate, however. Unless the project in question actually influences the incentive that households have to save, the appropriate rate for discounting future consumption should be the rate at which consumers actually do use, which is the after-tax interest rate. If the project stimulates saving, the external benefits of that saving could be taken into account as an extra benefit of the project.

Similar arguments apply to other sources of externalities associated with saving. In the case of saving for bequests, whether intentionally or unintentionally, external benefits may accrue to households to the extent that they have an altruistic regard for the well-being of future generations. Savings for bequests will take the form of a public good that generates benefits simultaneously for many members of the current generation. Because of the free-rider problem associated with altruism, private saving will tend to me too small on efficiency grounds. As in the case of investment externalities, there will be a case for government intervention to encourage saving for future generations. However, it is not clear that this should affect the social discount rate used for cost-benefit analysis. Unless the project actually encourages (or discourages) saving, no external effect will be created, and the appropriate discount rate is the one that households actually use.

Households may under-save for other reasons. They may simply be myopic. Or, as emphasized in recent literature, their preferences may be time-inconsistent (Laibson, 1997). This kind of behavior poses real problems for cost-benefit analysis since it implies that households are either irrational or otherwise end up regretting their saving choices. It is not at all obvious how a cost-benefit analysis methodology that is built on the foundations of welfarism and non-paternalism can be revised to take account of behavior that seems to contradict the household's own self-interest. This remains an open research question.

The above argument in favor of using the after-tax interest rate as the social discount rate is based on the assumption that the numeraire being used for cost-benefit analysis is current consumption. If another numeraire is used, a different discount rate will generally be required. For example, the numeraire proposed by Little and Mirrlees (1974) and often used by the World Bank is foreign exchange in the hand of the government. The discount rate should then be the relative value of future versus present foreign exchange held by the government. Little and Mirrlees (1974) argue that in a developing country context, marginal foreign exchange in the hand of the government will be used for investment, so the discount rate should be an investment discount rate rather than a consumption one. Thus, capital market distortions get reflected in the discount rate. Procedures for estimating the relevant social discount rate using the Little-Mirrlees approach may be found in Ray (1984) and Squire and van der Tak (1975).

A final argument for departing from the after-tax interest rate as the social discount rate arises when equity concerns are taken into account. In this case, the concerns will be intergenerational equity ones. We return below to the considerations that intergenerational

equity raises for cost-benefit analysis.

VII.   RISK AND UNCERTAINTY

Until now we have assumed implicitly that future benefits and costs were known. This will almost never be the case. At best, project evaluators and households may know the probabilities of various project outcomes occurring. In these circumstances, there will be risks associated with the fact that outcomes may turn out to be good or bad. The literature on decision-making under risk tells us how to evaluate risky outcomes of any sort. We begin with a summary of the costs of risk-taking in general, and then consider how that might be taken account of in the cost-benefit analysis of public projects.

The cost of risk-taking can be illustrated for the simplest case of a representative household facing uncertainty in a single dimension, say, their income. Let y_i be the household's income in state of nature i, where state i occurs with probability π_i. The n possible states of nature are exhaustive and mutually exclusive, so $\sum_{i=1}^{n} \pi_i = 1$. The household's preferences over alternative state-contingent outcomes are assumed to be ordered according to expected utility:

$$E[u(y)] = \sum_i \pi_i u(y_i) \tag{20}$$

where $u(y)$ is increasing and strictly concave, so the household is risk-averse ($u'' < 0$). Then, for any set of state-contingent outcomes y_i, the cost of risk can be defined as the amount of money the household would be willing to pay to avoid the risk. Formally, let k be the cost of risk. Then, it will satisfy:

$$u(\bar{y} - k) = \sum_i \pi_i u(y_i) \tag{21}$$

where $\bar{y} \equiv \sum_i \pi_i y_i$ is expected income. Thus, k is the amount of income the household would be willing to forgo in order to avoid the uncertainty associated with the expected income stream.

The cost of risk k, which is implicitly defined in (21), can be expressed explicitly by applying a Taylor expansion around $\bar{y}$ to obtain an expression for $u(y_i)$:

$$u(y_i) = u(\bar{y}) + u'(\bar{y})(y_i - \bar{y}) + 0.5 u''(\bar{y})(y_i - \bar{y})^2 + R \tag{22}$$

where R includes the higher-order terms. In what follows, we shall ignore these higher-order terms, and effectively use a second-order approximation for $u(y_i)$ in the right-hand side of (20). For risks that are relatively small compared with $\bar{y}$, we can approximate the left-hand side of (21) to the first order by $u(\bar{y} - k) \cong u(\bar{y}) - ku'(\bar{y})$. Combining this with (22) and using the definition of mean income, $\bar{y} \equiv \sum_i \pi_i y_i$, we obtain:

$$k \cong -\frac{u''(\bar{y})}{2u'(\bar{y})} \text{var}(y_i) = \frac{A}{2} \text{var}(y_i) \tag{23}$$

where $\text{var}(y_i) = \sum_i \pi_i (y_i - \bar{y})^2$ is the variance of income, and $A = -u''(y)/u'(y)$ is the coefficient of

absolute risk aversion. Thus, the cost of risk is higher the greater is the spread of outcomes and the more risk-averse is the household.

The cost of risk could in principle enter into cost-benefit analysis in two ways, directly and indirectly. The direct way involves uncertainty about the stream of net benefits of the project. If project returns are uncertain, the project evaluator could proceed in one of two ways. First, the projects benefits and costs could be measured in expected utility terms, using the analog of equation (20) above. That is, the expected utility of the project could be found by estimating the alternative streams of benefits and costs in real income terms and converting them to utility according to some assumption about the form of the utility function. This is likely to be very cumbersome. An alternative is to evaluate the benefits and costs in expected value terms and estimate a cost of risk associated with the uncertainty of the net benefit stream. This too is likely to be difficult.

Fortunately, it may be possibly to avoid taking account risks of a project even if the stream of net benefits is uncertain. That is because the risks may be diluted by risk-pooling or risk-sharing with other projects in the government sector. *Risk-pooling* refers to the case where the variance of a portfolio of projects is reduced if the portfolio is diversified with assets (projects) whose returns are to some extent independently distributed. Suppose a portfolio consists of n assets denoted $i = 1, \cdots, n$, where the return on asset i is r_i, with mean $\bar{r}_i$, and the share of the portfolio held as this asset is a_i. Let the variance of asset i and its covariance with asset j be:

$$\text{var}(r_i) = E\left[(r_i - \bar{r}_i)^2\right] \quad \text{cov}(r_i r_j) = E\left[(r_i - \bar{r}_i)(r_j - \bar{r}_j)\right]$$

The variance of the portfolio can be written:

$$\text{var}(r) = E\left[\left(\sum a_i r_i - \sum a_i \bar{r}_i\right)^2\right] = \sum a_i^2 \text{var}(r_i) + 2\sum\sum a_i a_j \text{cov}(r_i r_j)$$

Thus, when assets are independently distributed, $\text{cov}(r_i r_j) = 0$, so

$$\text{var}(r) = \sum a_i^2 \text{var}(r_i) < \sum a_i \text{var}(r_i)$$

That is, the variance of the portfolio is less than the sum of the variances of the individual assets. In the simple case where the assets are all identical and independently distributed, we have $a_i = 1/n$, so

$$\text{var}(r) = \sum \text{var}(r_i)/n^2 = \text{var}(r_i)/n$$

As n increases, $\text{var}(r)$ approaches zero. To the extent that risk-pooling exists among the projects in the public sector, this would suggest that the cost of risk associated with any one of them can be ignored when evaluating its costs and benefits.

While risk-pooling involves diversification over independently distributed projects, *risk-spreading* involves spreading the risks over a large number of households. Consider a

representative project whose (stochastic) return is *r*. And, suppose there are *m* identical persons sharing in the return of the project. In the case of public projects, *m* could be interpreted as the number of taxpayers. Each person then obtains a return of r/m, whose variance is $\text{var}(r/m) = E[(r/m - \bar{r}/m)] = \text{var}(r)/m^2$. Therefore, using (23) above, the cost of risk to a given individual is:

$$k = \frac{A}{2} \frac{\text{var}(r)}{m^2}$$

Therefore, the total project risk is:

$$mk = \frac{A}{2} \frac{\text{var}(r)}{m}$$

As the populations increases, the total project risk goes to zero, a result due to Arrow and Lind (1970). The implication is that if the project return is shared among a large number of taxpayers, the cost of risk can be safely ignored.

As mentioned, risk effects can also arise indirectly in cost–benefit analysis. To the extent that the financing of a project causes private sector investment to be crowded out, some risk that would otherwise have been incurred in the economy is now avoided. In the case considered earlier in equation (19), the opportunity cost of a yen's worth of forgone investment when there was no reinvestment of asset income was simply ρ/r. Suppose now that the rate of return on private investment included a risk premium of, say, β per yen of return. Then, the opportunity cost of forgone private investment would be $(\rho - \beta)/r$, so the marginal cost of borrowing becomes

$$\text{MCB} = \frac{\Delta S}{\Delta B} - \frac{\rho - \beta}{r} \frac{\Delta D}{\Delta B}$$

In other words, risk–free rates of return should be used to evaluate the MCB.

The upshot of this section is that taking account of the cost of risk is a difficult task since these risks are typically not observable. Ignoring the direct costs of project risk may be justified if one can appeal to risk–pooling and/or risk–spreading arguments. As well, one can account for the indirect costs of risk by using a risk–free notion of MCB. Even so, there may be lingering doubts. In practice, these doubts are often addressed by sensitivity analysis. The project evaluator presents a range of estimates corresponding with optimistic through pessimistic scenarios with respect to net benefits, perhaps with some guess as to the chances of each outcome occurring. It is then up to the decision–maker to attach appropriate weights to the various outcomes.

VIII.   INTERGENERATIONAL EQUITY ISSUES

Many public projects last for a long time, spanning several generations, and as such their net benefits accrue to both current and future cohorts. This is true not only of investment projects,

but of fiscal policies more generally. Changes in government debt, changes in social insurance programs and even tax reforms have as one of their main effects redistribution among generations, as Auerbach and Kotlikoff (1987) have emphasized. Government deficits are effectively postponed taxes, and provided the resulting debt is held long enough, this necessarily entails increased intergenerational transfers from the young to the old. Changes in unfunded public pensions have the same effect, as do other sorts of social programs funded out of current revenues, such as health care. And, major tax reforms have the same effect: substituting a consumption tax for a wage tax is effectively equivalent to an intergenerational transfer from the old to the young.[8]

Evaluating policies that have as one of their prime effects an intergenerational transfer involves taking some view about the relative weights to be placed on future versus present generations. Given the prevalence of intergenerational transfers in government budgets and the importance of the issue for policy, in this section we outline the principles that economists might use to evaluate the intergenerational effects of government policy. Then, some methodologies that might be used to make these principles operational are discussed. Naturally, since we are dealing with transfers among households, value judgments are ultimately involved.

VIII.1.   PRINCIPLES FOR EVALUATING INTERGENERATIONAL TRANSFERS

Four potential principles are discussed, some of which are related and some of which are in conflict.

VIII.1.1.   The Intergenerational Benefit Principle

The benefit principle is has a long-honored history in public finance. Although it has limited appeal as a distributive principle in static settings, it might be more attractive in an intergenerational setting. It would require each cohort to bear the costs of governments in accordance with the benefits they receive. Arguments for funding public pensions or for leaving the environment intact can be viewed as reflecting the benefit principle.

Properly attributing benefits on a cohort-by-cohort basis is a difficult task in itself, but there are problems of principle as well. The appeal of the benefit principle might seem to be that it is value-free in the sense that it entails no redistribution. However, that is a rather narrow interpretation. The benefit principle implicitly takes existing property rights as given and preferable to other allocations with different property rights. The inviolability of property rights does not reflect societal consensus for intra-generational equity, and there is no good reason

---

[8] Some have argued that intergenerational transfers imposed on the economy by the government will be undone by households through their bequest behavior (Barro 1974). However, there is little evidence that this will be the case, and there are good theoretical arguments why it should not be. See the summary of arguments in the *Journal of Economic Perspectives* **3**, Spring, 1998.,

why it would have more appeal in the intergenerational context. There is a further conceptual problem. Public programs presumably generate economic surpluses–total benefits exceed total costs–and these may be of a significant size. There is no natural way to apply the benefit principle to allocate costs among households without making some interpersonal judgment about how to allocate the surplus.

There is an even more serious problem with applying the benefit principle in an intergenerational context. Some assets, such as natural resources and the natural environment, are endowed to societies by nature. Government policies affect the stock of these natural assets, and therefore the amount that are passed on to future generations. Since these natural assets are owned by society at large, their property rights are not defined by cohort. The intergenerational benefit principle cannot be applied without specifically assigning property rights to natural assets to different generations to determine which generations are entitled to which shares of the fruits of these natural assets. Such a problem does not arise with respect to private assets, including those that are obtained from nature, since ownership is well defined. Moreover, it does not apply to government programs that are financed by taxes.

The fact that there is no natural way to assign property rights for natural assets across generations implies that the benefit principle in its usual form is in effect non–operative. To apply it would involve implicitly assigning such property rights, and one cannot do that without invoking some intergenerational equity judgment. Since aggregate intergenerational transfers in the comprehensive sense must take account of the amount of public assets passed on to future generations, some judgment must be made about intergenerational equity, if only implicitly.

VIII.1.2.   Intergenerational Risk Sharing

A major difference among cohorts is the set of exogenous circumstances that affect their well–being. Given cohorts can be born lucky or unlucky because of the circumstances that apply at their date of birth or over their lifetimes, and over which they have no control. The circumstances may be demographic: cohorts that are relatively large face a disadvantage relative to those that are relatively small. There may be shocks, such as a major war, that affect given cohorts. Natural catastrophes, such as earthquakes, weather shocks or disease occur from time to time. And, economic shocks, such as depressions, can affect cohorts during their working lives. In each of these cases, intergenerational transfers can provide a form of social insurance, where the insurance is social because markets cannot provide it.

There is some evidence that intergenerational risk does motivate governments to provide social insurance. Wars are financed largely by debt, and public pension programs have been instituted to assist those who were unlucky during their working lives. Of course, intergenerational transfers instituted for this purpose are by their nature temporary. Risk-sharing cannot be used as an argument for supporting permanent intergenerational

transfers.

VIII.1.3.   Tax Smoothing

A further argument for temporary intergenerational transfers is as a means of smoothing tax rates over time (Barro, 1974). Since the excess burden of tax distortions is convex in the tax rate, a given amount of revenues obtained over time from tax rates that fluctuate will have a larger deadweight loss than if the tax rates were smooth. This is related to the social insurance argument in the sense that the sources of fluctuating tax rates can be similar to those that lead to cohorts having better or worse fortune. But the argument is posed purely in efficiency terms.

The problem with the tax smoothing argument is that, in an intergenerational context, one cannot separate efficiency and equity arguments. Smoothing tax rates over time involves raising taxes for some cohorts and reducing them for others, which will give rise to the usual equity–efficiency trade–off. In this case, the equity side of the argument must rely on some intergenerational welfare comparison.

VIII.1.4.   Intergenerational Equity

As the above indicates, an evaluation of the net benefits of policies accruing to different cohorts necessarily involves an intergenerational comparison of welfare. This is especially apparent when there are systematic and persistent differences in well–being among different cohorts (and not just fluctuations). Intergenerational welfare comparisons necessarily involve value judgments, and settling on interpersonal equity norms is a major and well–known problem in policy evaluation. In the end, there must be some social consensus, and the political process obviously plays an important part in forging, interpreting and applying that consensus. That does not imply that there is no role for normative analysis as a complement to the positive study of the political process. On the contrary, normative analysis and even normative advocacy is an indispensable role of economists and other policy advisors: the political system does not take the form of a political marketplace that yields a determinant outcome.[9]   In fact, one cannot avoid making interpersonal welfare comparisons in policy advice and evaluation, and presumably the basis for these comparisons that one is using should be made explicit.

There is likely broad consensus about some of the basic principles of redistributive equity, whether applied within or between generations. After all, governments are engaged heavily in redistribution, and the forms and extent of that redistribution are very similar across countries. Moreover, they are quite different than one might predict if one were starting with a public choice model based on purely self–interested voters. Since we have already discussed these

---

[9] The case for normative economic analysis as a necessary complement to public choice of positive political economy analysis is made in length in Boadway (2002).

principles earlier when discussing the properties of the social welfare function, we need do no more than recall them here. They include the following, each of which has substantial meaning:

- Individualism: Individuals are the best judges of their own well-being.
- The Pareto principle: Policies that make some persons better off and no one worse off are preferred policies.
- Aversion to inequality: All other things being equal, society prefers outcomes in which welfare is more equally distributed to those in which it is less equally distributed, where the degree of inequality aversion depends partly on ideology, and partly on the perceived severity of the equity efficiency trade-off.
- Equality of opportunity: Not all sources of inequality might be judged relevant for redistributive correction. The 'principle of compensation' suggests that persons ought to be compensated only for those adverse outcomes that are due to factors outside their control, such as their productive ability, their health, and their date of birth. But, if persons are responsible for adverse outcomes (e.g., low incomes) because of the way they have freely chosen to behave, these differences ought not to be compensated: the 'principle of responsibility'. The principles of compensation and responsibility together lead to the idea of equality of opportunity as one of the guiding principles of redistribution: opportunities that households face, as reflected in say their budget possibilities, ought to be equalized.[10] The principle of equality of opportunity gives rise to important policy instruments like education and health care alongside the tax-transfer system.
- Social insurance: Households might be subject to unexpected shocks over which they have no control and against which they cannot fully insure. Or, they may simply be uninsurable at birth. Private insurance might fail to address unemployment as well. These kinds of arguments constitute the major reason for the substantial programs of social insurance implemented by most OECD countries.

While there is obviously not universal agreement on the details of application of these principles, there seems to be enough of a consensus about their relevance to take us a long way in judging in qualitative terms at least minimal standards of redistribution that should apply. Although we may be less used to applying them with respect to intergenerational redistribution, there seems to be no particular reason why such an extension would not reflect societal consensus. For example, widely held concerns about the environment are based on notions of intergenerational equity. Given these principles, we can imagine applying an intergenerational social welfare function analogous to (1) in an intergenerational context. This would lead to using welfare weights in aggregating the costs and benefits accruing to various cohorts.

Some might argue that the normative principles of intergenerational economic justice are in

---

[10] A complete treatment of equality of opportunity along these lines may be found in Roemer (1998). While he applies it to intra-generational transfers, the same principle could apply between generations.

any case hardly sufficient for evaluating public policies that impact on different cohorts. There may be no clear consensus about the degree of aversion to inequality that should be applied to intergenerational transfers, and the future may be fraught with uncertainty. It might also be argued that governments are overly short-sighted and will discount the benefits of future generations who, after all, are not part of the voting constituency. More important, political outcomes may differ from the interests of citizens. If the political process is inherently inefficient, wasteful and captive of special interests, special attention needs to be devoted to deviations of policy outcomes from what the electorate truly wants. On the other hand, political processes may be efficient with political competition leading to outcomes that approximate the social consensus. In that case, if governments discount the welfare of future generations, it is because that reflects the social consensus among those currently alive. If one does not like the social consensus that has been formed, one can always try to persuade the public that it should adopt a different consensus. That is part of the role of normative analysis.

VIII.2.   ACCOUNTING FOR INTERGENERATIONAL EFFECTS

Incorporating intergenerational equity weights into policy evaluation is a difficult matter. To make a well-informed judgment, one would ideally need to have the following types of information, in addition to the effects of policies on the welfare of different cohorts:

- A comprehensive measure of the extent of public indebtedness, including all types of net tax obligations passed on to members to current and future generations that are implicit in existing tax policies;
- An account of the full benefits of government policies to existing and future generations, given some presumption about the path of government policies far enough into the future;
- An estimate of the allocation of the benefits of the stock of natural capital to existing and future generations, given the government policy stance; and
- A measure of the relative levels of well-being of existing and future generations, given the policies that are in place.

These are very difficult to obtain practically as well as conceptually. However, there are some accounting procedures that can help in informing policy judgments. These fall under the general rubric of *generational accounting*, a procedure developed to capture the amount of intergenerational transfers that are implicit in current policy stances.[11]

Generational accounting assigns to members of currently alive and future generations the net costs of financing existing fiscal policies projected into the future. It is purely an accounting exercise, with no account taken of behavioral responses to the fiscal policies. The building block

---

[11] Generational accounting was developed in the USA by Auerbach, Gokhale and Kotlikoff (1991), and has been applied to a number of countries in Auerbach, Kotlikoff and Leibfritz (1999).

is a cohort's *generational account*, which is the present value of the net taxes that the representative household of a given age cohort is liable to pay either over their full lifetime, or their remaining lifetime, as the case may be. Net taxes include tax liabilities of all forms attributed to the household less transfers received and less whatever types of government expenditures can be attributed to the household, such as health and education. Generational accounts can then be converted to generational *lifetime net tax rates* by dividing them by the present value of lifetime income.

Generational accounts or net tax rates themselves can be used to compare the lifetime net tax rates for various cohorts, giving some indication of intergenerational transfers implicit in a given policy stance. Moreover, they can be used to calculate the burden that would be left for future generations from alternative government policies. The basis for this is the government's *intertemporal budget constraint*, which requires that the present value of the government's future stream of net taxes (taxes net of transfers and expenditures that can be attributed to cohorts) equal its existing debt and the present value of its future stream of expenditures. The future stream of net taxes can be disaggregated into those attributable to currently alive and future cohorts. For each currently alive cohort, a generational account is calculated consisting of the present value of net taxes owing from the current period to the predicted end of life. The present value of net taxes owed by all future generations is then the sum of current government net debt and the present value of future expenditure obligations less the sum of generational accounts for all those currently alive. These obligations left for future generations are then typically assumed to be shared equally among all future cohorts to give an idea of the burden that on average is left for future generations. In effect, the net liabilities of the government as of today, its public indebtedness, are amortized over all future generations: they all share in the paying debt.

The resulting calculations are very suggestive, and provide a useful tool for capturing at least some of the effects of the extent of intergenerational transfers implicit in existing policies. There can obviously be disputes about the various assumptions built into the calculations with respect to future policies, population, the assignment of taxes and transfers to various age cohorts, and so on. For our purposes, a more interesting question is how suitable are properly measured generational accounts for evaluating policies involving intergeneration redistribution. We can identify a number of shortcomings of the current methodology of generational accounting as a complete measure of the intergenerational effects of a government's policy stance.

Forward‒ and Backward‒Looking Generational Accounts. The generational accounting method is forward looking in the sense that the accounts for those currently alive include only the net tax liabilities for the remainder of their lives. Thus, one does not get a full picture of the intergenerational transfers that have applied to them over their full lifetimes. This makes it difficult to evaluate policies that affect both present and future generations.

Public Capital. Generational accounts do not take takes account of the public capital stock and infrastructure that exists at the current time. Public capital yields services that are of benefit to current and future generations, and these should be included as an element of the benefits received. In practice, these would presumably be very difficult to estimate and evaluate.

Natural Capital. Natural resource wealth and environmental capital that are commonly owned are not included in generational accounts, despite the fact that they constitute a form of asset wealth that is shared among generations. Natural capital differs from the public capital stock in the sense that it has not been produced using public resources. Nonetheless, the benefit it provides to different cohorts depends on government policy. In principle, one could attribute to generations the benefit that they obtain from natural capital, but that would be a heroic undertaking.

Intangible Public Capital. The most difficult assets to value are the invisible ones, such as accumulated knowledge, social capital or the society's institutions. Yet the passing on of intangible assets from one generation to the next represents an important form of intergenerational transfer.

Measures of Generational Well-Being. Finally, in order to make a judgment about the effects of intergenerational transfers, it would be useful to have measures of how well off future generations will be relative to current generations. There is a presumption in the generational accounting literature that the intergenerational balance is achieved when there is parity of generational accounts or of lifetime net tax rates among different generations. But that might only be the case if all generations are also equally well off. Otherwise, unlucky cohorts should have lower lifetime net tax rates.

Despite these difficulties and drawbacks, the concept of generational accounting is a suitable first step to evaluating changes in intergenerational transfers, and one that serves as a basis for future development. Even if the evaluator is not prepared to subscribe to a particular value judgment about intergenerational equity, it might still be helpful to inform policy-makers of the intergenerational outcomes that are implicit in their policies. The tool of generational accounting provides a promising approach since it focuses precisely on the relative financial burdens imposed on different cohorts, living and unborn, that will satisfy the government's intertemporal spending requirements.

IX.   CONCLUDING REMARKS

In this paper, we have summarized the main issues in the evaluation of projects. It is obvious

from our discussion that project evaluation is very much an art, although one with scientific underpinnings. Our purpose has been to indicate what those scientific underpinnings are, so that readers can have an economic perspective on what is involved. The technical literature on project evaluation is a well-established one, but one which must evolve with the times. Recent advances in economic theory have probably not yet been incorporated into project evaluation principles to the extent that they could be. For example, the importance of asymmetric information and its implications for market behavior and market failure have been very much in the forefront of economic analysis. Yet, little has been done to incorporate imperfect information into project evaluation rules. This is particularly true insofar as the existence of imperfect information has implications for unemployment.   Similarly, there has been considerable research activity into studying the determinants of growth, and whether or not unfettered markets are conducive to high growth rates. Little of this has found its way into applied welfare economics. Finally, the importance of illegal or underground activity has been increasingly recognized. This too might have implications for project evaluation. As with everything else in economics, project evaluation will presumably evolve.

References

Ahmad, Ehtisham and Nicolas Stern (1991), *The Theory and Practice of Tax Reform in Developing Countries* (Cambridge: Cambridge University Press).

Arrow, Kenneth J. and R.C. Lind, 'Uncertainty and the Evaluation of Public Investment Decisions', *American Economic Review 60*, 364–78.

Auerbach, Alan J. and Laurence J. Kotlikoff (1987), *Dynamic Fiscal Policy* (New York: Cambridge University Press).

Auerbach, Alan J., Jagadeesh Gokhale and Laurence J. Kotlikoff (1991), 'Generational Accounts: A Meaningful Alternative to Deficit Accounting', in David Bradford, ed., *Tax Policy and the Economy*, Vol. 5 (Cambridge, USA: MIT Press), 55–110.

Auerbach, Alan J., Laurence J. Kotlikoff and Willi Leibfritz, eds. (1999), *Generational Accounting Around the World* (Chicago: University of Chicago Press).

Barro, Robert J. (1974), 'Are Government Bonds Net Wealth?', *Journal of Political Economy 82*, 1095–117.

Blackorby, Charles and David Donaldson (1990), 'A Review Article: The Case Against the Use of the Sum of Compensating Variations in Cost–Benefit Analysis', *Canadian Journal of Economics 23*, 471–94.

Boadway, Robin W. (1974), 'The Welfare Foundations of Cost–Benefit Analysis', *Economic Journal 84*, (December 1974), 926–39.

Boadway, Robin W. (1976), 'Integrating Equity and Efficiency in Applied Welfare Economics', *Quarterly Journal of Economics 90*, 541–56.

Boadway, Robin W. (1978), 'Public Investment Decision Rules in a Neo–Classical Growing Economy', *International Economic Review 19*, 265–87.

Boadway, Robin (2002), 'The Role of Public Choice Considerations in Normative Public Economics', in Stanley L. Winer and Hirofumi Shibata, eds., *Political Economy and Public Finance* (Cheltenham, UK: Edward Elgar), 47–68.

Boadway, Robin W. and Neil Bruce (1984), *Welfare Economics* (Oxford: Basil Blackwell).

Browning, Edgar (1976), 'The Marginal Cost of Public Funds', *Journal of Political Economy 84*, 283–98.

Dasgupta, Partha, Stephen Marglin and Amartya Sen (1972), *Guidelines for Project Evaluation* (New York: UNIDO).

Diamond, Peter A. and James A. Mirrlees (1971), 'Optimal Taxation and Public Production', *American Economic Review 61*, 8–27 and 261–78.

Drèze, Jean and Nicholas Stern (1987), 'Theory of Cost–Benefit Analysis' in Alan J. Auerbach and Martin S. Feldstein (eds.), *Handbook of Public Economics*, Vol. 2 (Amsterdam: North–Holland), 909–89.

Feldstein, Martin S. (1972a), 'The Inadequacy of Weighted Discount Rates', in Richard Layard

(ed.) *Cost–Benefit Analysis* (Harmondsworth, UK: Penguin), 311–32.

Feldstein, Martin S. (1972b), 'Distributional Equity and the Optimal Structure of Public Sector Prices', *American Economic Review 62*, 32–6.

Harberger, Arnold C. (1969), 'Professor Arrow on the Social Discount Rate', in G.C. Somers and W.D. Wood (eds.), *Cost–Benefit Analysis of Manpower Policies: Proceedings of a North American Conference* (Kingston, Canada: Queen's University Industrial Relations Centre).

Harberger, Arnold C. (1971a), 'Three Basic Postulates for Applied Welfare Economics: An Interpretive Essay', *Journal of Economic Literature 9*, 785–97.

Harberger, Arnold C. (1971b), 'On Measuring the Social Opportunity Cost of Labour', *International Labour Review 103*, 559–79.

Harberger, Arnold C. (1978), 'On the Use of Distributive Weights in Social Cost–Benefit Analysis', *Journal of Political Economy 86*, S87–S120.

Harris, J.R. and M.P. Todaro (1970), 'Migration, Unemployment and Development: A Two–Sector Analysis', *American Economic Review 60*, 136–42.

Harrison, A.J. and D.S. Quarmby (1972), 'The Value of Time', in Richard Layard (ed.) *Cost–Benefit Analysis* (Harmondsworth, UK: Penguin), 173–208.

Jones–Lee, Michael W. (1976), *The Value of Life: An Economic Analysis* (Chicago: The University of Chicago Press).

Laibson, David I (1997), 'Golden Eggs and Hyperbolic Discounting', *Quarterly Journal of Economics 112*, 433–77.

Little, Ian M.D. and James A. Mirrlees (1974), *Project Appraisal and Planning for Developing Countries* (New York: Basic Books).

Marglin, Stephen A. (1973), 'The Opportunity Costs of Public Investment', *Quarterly Journal Of Economics 77*, 274–89.

Musgrave, Richard A. (1959), *The Theory of Public Finance* (New York: McGraw–Hill).

Ray, Anandarup (1984), *Cost–Benefit Analysis: Issues and Methodologies* (Baltimore: Johns Hopkins Press for the World Bank).

Roemer, John (1998), *Equality of Opportunity* (Cambridge, USA: Harvard University Press).

Sandmo, Agnar (1998), 'Redistribution and the Marginal Cost of Public Funds', *Journal of Public Economics 70*, 365–82.

Sen, Amartya K. (1970), *Collective Choice and Social Welfare* (San Francisco: Holden Day).

Squire, Lyn and Herman G. van der Tak (1975), *Economic Analysis of Projects* (Baltimore: Johns Hopkins Press for the World Bank).

Usher, Dan (1986), 'Tax Evasion and the Marginal Cost of Public Funds', *Economic Inquiry 24*, 563–86.

Wildasin, David E. (1984), 'On Public Good Provision with Distortionary Taxation', *Economic Inquiry 22*, 227–43.