THE ECONOMIC EVALUATION OF PROJECTS

by

Robin Boadway

Queen's University, Kingston, Canada

I. PRINCIPLES OF VALUATION

This paper summarizes the principles that are used to evaluate projects from an economic point of view. The term 'project' should be thought of in the broadest of senses. It can refer to individual investment projects, like the building of a bridge or a dam. More broadly, it can include general expenditure programs, like education, health care or nutrition spending. Or, it can refer to government policies like reforms of the tax-transfer system, or the regulation of economic activities. Our use of the term 'project' in what follows should be taken as including all of these. Economic evaluation of a project implies a measure of its net benefits in monetary terms, as opposed to, say, an evaluation as to its political feasibility.

Attaching a monetary measure to the benefits and costs of a project raises the fundamental question of whose benefits and costs are relevant. We follow the convention used in much of the economics literature referred to as the principle of *welfarism*, which holds that what ultimately counts is the welfare of the individuals in the society.[1] Thus, the objective of project evaluation is to measure the benefits and costs accruing to those individuals who are affected.

Measuring the benefits or costs of a project to any given individual involves asking how much the individual would be willing to pay to have the benefit or to avoid the cost. What they would be willing to pay will typically differ from what they actually have to pay. They may not have to pay anything if the project involves public goods or

[1] The term 'welfarism' is due to Sen (1970), who used it to describe the property of a social welfare function which orders alternative resource allocations according to the levels of utility achieved by members of the society. Sen has been critical of the principle of welfarism, arguing that other characteristics of social well-being, such as freedom, justice, non-discrimination, equality of opportunity, and so on should also count. Project evaluators sidestep this issue by arguing either that the projects under consideration have no particular effect on these virtues, or that if they do it is impossible to measure them so they ought to be weighted according to the values of those ultimately responsible for decision-making.

services provided free-of-charge. But, even if they do pay something, it will likely be less than the amount they would be willing to pay, their *willingness-to-pay*. In other words, they will obtain some *surplus* from it, which implies that market prices will not suffice.

Technically, the willingness-to-pay for a project's benefits (or the minimum amount those affected would be willing to accept to bear the project's costs) is measured by some generalized notion of the *compensating variation* (CV, for short). The standard definition of the CV is obtained implicitly from the following equation:

$$V(p_0, m_0) = V(p_1, m_1 - CV)$$

where V(p,m), the so-called *indirect utility function* is the individual's utility as a function of the vector of prices of commodities p and income m, with the subscripts 0 and 1 referring to the pre-project and after-tax project prices and income. Alternatively, the CV can be measured directly by making use of the *expenditure function*, E(p,U), which indicates the amount of income required to achieve utility level U when prices are p. The compensating variation can then be written:

$$CV = E(p_1, U_1) - E(p_1, U_0)$$

A number of features of the CV are worth noting:

- In the special case in which only one price changes, the CV is equivalent to the conventional consumer surplus, the area beneath the demand curve and the horizontal price line.[2]

- More generally, the CV as defined measures the change in utility at a set of reference prices p_1, those of the final situation. In the two-good case, this is the

[2] Technically speaking, it is the area beneath the compensated demand curve associated with the pre-change utility level. This area will differ from the area beneath the uncompensated demand curve because of income effects. For typical project evaluations, the difference will not be important, given the limitations of data.

distance between the initial and final indifference curves measured with line whose slope reflects the relative prices of the final situation.

- The measure of welfare change is not unique: any set of reference prices could have been used. For example, use of the initial prices yields the so-called *equivalent variation*, EV. More generally, different references prices give rise to different *money metric* measures of utility change. The actual monetary measure of welfare change will differ depending on the money metric used, but they will all be roughly the same.[3] Given the errors of measurement and uncertainties usually involved in actual project evaluation, it will make little practical difference which measure is used. We shall follow the convention of referring to our welfare changes measures as CV's, but it should be recalled that we could be using any of the money metric measures discussed here.

- The prices in the expression for welfare change can refer to either goods purchased by the individual or factors supplied.

- The CV measure refers especially to the case in which all inputs and outputs have prices associated with them, and consumers can vary quantities at will. If neither of these properties holds, the CV formulation must be suitably amended. But the same principles are involved: CV measures the consumer's willingness-to-pay.

- If the benefits and costs occur over different periods of time, they must be converted using a discount rate to a monetary equivalent at a given point in time. Typically, this will involve discounting to the present at the individual's discount rate to give a present value (PV).

---

[3] For example, in the case of a single price change, each welfare change measure will correspond with a consumer surplus area beneath a demand curve, but the actual demand curve that is appropriate will vary according to the measure used. For the CV, the compensated demand curve corresponding to the utility level of the final situation will be appropriate, while for the EV, that corresponding to the initial utility level is used. The two will differ only because of income effects. See Boadway and Bruce (1984) for a more complete discussion.

These CV's are the building blocks of project evaluation. But, any given project will give rise to a spectrum of CV's, one for each individual affected by the project. How do we go about aggregating individual CV's together? The problem is that, while the CV indicates in monetary terms an individual's change in welfare, there is no objective way of comparing CV's across persons. Two alternative procedures may be followed.

The first is to follow the precept advocated by Harberger (1971a) and to treat a rupee as worth a rupee no matter whose hands it is in. In this case, individual CV's can be summed to an aggregate CV intended to measure the net benefits to all members of the society. Though this procedure is most frequently used in practice, the theoretical case for it is highly disputed in the literature.[4] Proponents will usually cite one or more of the following arguments, each one intended to support the view that the aggregate CV measures in some sense the efficiency benefits of the project: i) the government has tax-transfer policy instruments available for redistribution, and should be presumed to be using them to undo any differences in the value of a rupee to various households; ii) there are many projects being undertaken, and their redistributive effects should be roughly offsetting; iii) if the aggregate CV is positive, that is *prime facie* evidence that those who gain could hypothetically compensate those who lose and still be better off. These arguments have failed to convince the critics of the procedure, and the debate stands unresolved. Perhaps the strongest argument for aggregating CV's is a purely practical one: given that aggregate data are all the evaluator has available, it is impossible to do anything but measure aggregate CV's, perhaps supplementing that where possible with evidence about the gainers and losers so that the policy-maker can make an informed judgement.

The second procedure is to incorporate *distributive weights* directly into the project evaluation according to some preconceived notion of deservedness. This is not a straightforward exercise. For one thing, the distributive weights ought, in principle, to be applied to individual CV's. But, that will not be possible for the reason cited above that individual CV's cannot be measured. The best that can be done is to apply distributive weights to different types of benefits and costs according to some evidence about how important they are to individuals of various circumstances (income, needs, etc.). The other main problem with this is choosing the weights. This inevitably involves a value judgment, presumably one that will not command consensus. A common procedure is to parameterize the social welfare function that aggregates household utilities using to a single parameter, a common one being the degree of aversion to inequality. For example, suppose individual i's level of utility measured in some money metric is Y_i. Then, let the social welfare function used to aggregate the monetary measures of various households be:

$$W(Y) = \sum (Y_i)^{1-\rho} / (1-\rho)$$

The parameter ρ is the coefficient of aversion to inequality, or more formally the elasticity of the marginal social utility of Y_i, or $\rho = -W''(Y_i)/[Y_i W'(Y_i)]$. It captures the extent to which one wants to put higher values on monetary gains accruing to various households. Given that there is likely no agreement over its exact value, the evaluator can provide estimates of the net benefits of the project based on different values for *ρ*, leaving it to the policy-maker to decide among them.

What we are left with then are two ways of addressing the distributive effects of a project. The first simply uses aggregate CV's measures to estimate costs and benefits of alternative projects, and reports on whatever patterns of distribution of benefits among individuals of different incomes can be estimated. The other is to incorporate distributional weights into the CV measures themselves, using a range of such weights. In either case, the policy-maker is left to choose among options.

The above procedure relies on individual CV's to capture the full benefits and costs of the project. Two caveats are in order before turning to the details of project evaluation. The first is that sole reliance on CV's reflects fully a welfaristic social objective function, and that may not be universally accepted. Policy-makers may be interested in some non-welfaristic objectives as well. These will often not be measurable in monetary terms, in which case the evaluator may simply report the consequences of

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

the project for these other objectives. Thus, for example, the effect of the project on society's minorities, or on anti-discrimination objectives can be reported alongside the monetary net benefits. The second caveat is that there may be external benefits or costs from the project that should form part of the project evaluation. In principle, these should be measured in willingness-to-pay terms, though that will often be challenging given that external effects are typically difficult to quantify.

II. THE DECISION RULE

The purpose of project evaluation is to calculate the net benefits of a project in such a way as to form a basis for informing policy-makers as to whether the project should be undertaken. The private sector is engaged in these sorts of calculations on a continuing basis, and it is natural to begin by asking why economic project evaluation should be any different from calculations of financial profitability that are used to guide the investment decisions of firms. In fact, there are several factors which make economic evaluation distinct from private profitability:

- Market prices generally deviate from marginal social values if there are distortions in the economy. The sources of the distortions might be government policies (e.g., taxes, tariffs, regulations) or they might be inherent in the market economy (e.g., monopoly).

- There may be externalities, either beneficial or detrimental, which will not be reflected in market prices. Environmental pollution is an obvious example of a negative externality, while the generation of useful information that cannot be appropriated is a positive externality. Because these effects are difficult to quantify, let alone to value, the task of economic valuation is typically more difficult than private profitability calculations.

- Some inputs or outputs may have no explicit market price attached to them, such as the value of time saved on a public transportation facility, or the value of improvements in health and longevity. Though these can be quantified, they are nonetheless difficult to put a money value on.

- Economic values must include indirect benefits resulting from induced changes elsewhere in the economy. This will be relevant when outputs change on markets in which there is a distortion, since a distortion results in the benefit to users of changes in the quantity purchased differing from the cost to suppliers of making the changed quantities available.

- Projects may not be self-sufficient, but may require financing from the public purse. Since it is costly for the government to raise revenues, the excess burden of public financing must be taken into account in valuing projects. Private projects will use private sources of financing whose cost is taken into account directly in the rate-of-return calculation.

- The discount rate used for public projects, the *social discount rate*, will differ from the private discount rate because of capital market distortions

- And, as mentioned, public projects may take into account redistributive equity or other social considerations

All of these points imply that project evaluators must take into account a number of considerations not found in financial profitability studies. How this is done will occupy the remaining sections of this paper. In the rest of this section, we address the rule to be used as a basis for deciding on the economic desirability of a project.

THE PRESENT VALUE CRITERION

A project will be worth doing if the sum of its benefits is at least as great as the sum of its costs, measured in monetary terms. Given that benefits and costs will occur across several time periods, and that rupees today are worth more than the promise of rupees next year, both streams must be converted to a common time period, conventionally taken to be the present period. Thus, the net benefit of the project will be its net present value (NPV), defined as the present value of the benefits (PVB) less the present value of the costs (PVC), or:

$$NPV = PVB - PVC = \sum (B_t - C_t)/(1+i)^t$$

where B_t and C_t are the benefits and costs in period t, i is the one-period social discount rate (assumed constant), and t goes from 1 until the termination date of the project. If this PV is positive, the project should be undertaken. Or, if the policy-maker is restricted to considering mutually-exclusive alternatives, the one with the largest PV should be undertaken. Such alternatives might include identical projects with alternative starting times, projects differing only in scale, projects of differing durabilities, or alternative groups of projects.

There are a number of issues in implementing the PV criterion that ought to be mentioned. They are as follows.

Alternative PV Formulations I: The Benefit-Cost Ratio

Policy-makers often like to have a simple summary statistic indicating how beneficial a project is. The benefit-cost ratio, defined as the ratio of the present value of benefits to the present value of costs or PVB/PVC, provides an intuitively appealing measure of the extent to which benefits outweigh costs. As long as it exceeds one, it can be relied on to indicate whether a project has a positive NPV. But, it can be misleading in the case of ranking projects which are mutually exclusive alternatives, so cannot be used to choose the project which maximizes net social benefits. The reason is that it does not account for the scale of the project.[5]

The fact is that the benefit-cost ratio uses precisely the same information as the NPV, but presents it is a slightly different form. If policy-makers find it to be useful, it is not difficult to supplement it with the NPV to ensure that the NPV criterion is being satisfied.

---

[5] A simple example will illustrate. Consider two projects, A and B. Project A has a present value of benefits and costs of PVB = 2,000,000 rupees and PVC=1,000,000 rupees, giving a benefit-cost ratio of 2 and an NPV of 1,000,000 rupees. Project B has PVB=1,200,000 rupees and PVC=400,000 rupees, for a

Alternative PV Formulations II: The Internal Rate of Return

The internal rate or return (IRR) consists of a net present value calculation of a different sort. The IRR is defined as the discount rate which makes the present value of the stream of benefits less the present value of the stream of costs identically zero. Algebraically, the IRR is defined as the value of λ which satisfies the following equation:

$$\sum (B_t - C_t)/(1+\lambda)^t = 0, \quad t = 1, \ldots, \mathrm{T}$$

As can be seen, this is a fairly complex equation to be solved for λ. In fact, it is a polynomial of a degree T, the length of the time horizon. (This can be seen by multiplying the equation by $(1+\lambda)^T$, which leads to an equation of the form $a\lambda + b\lambda^2 + c\lambda^3 + \cdots + k\lambda^T = 0$.) In general, this equation can have as many as T solutions for λ. The number of solutions will correspond with the number of times B_t - C_t changes sign. Fortunately, this will not typically be a problem since, for most projects, B_t - C_t will change sign only once, being negative initially while capital costs are being incurred, and then positive for the rest of the project life. In these circumstances, the IRR will not only be uniquely defined, but it will also indicate to the policy-maker whether the project is socially profitable — if λ exceeds the social discount rate i, the NPV of the project will be positive.

But, like the benefit-cost ratio, the IRR can be unreliable since it may not rank mutually-exclusive alternatives according to their NPV's. Moreover, this problem is not simply one of scale. The problem arises because projects with different time profiles can have different relative NPV's at different discount rates — projects whose benefits accrue later in life will be particularly penalized at high discount rates. This is something that the IRR cannot possibly take account of. For example, suppose two different projects each cost 1 million rupees, so they have a net benefit of -1 million in year zero. Project A generates no net benefits in period 1 and 1.21 million rupees in period 2, while project

---

benefit-cost ratio of 3 and a NPV of 800,000 rupees. While project B has a higher net present value, it

B generates all its net benefits of 1.15 million rupees in period 1. The IRR of project A is .10, while that of project B is .15, so the latter would be chosen on the IRR criterion. Suppose the discount rate is i=.02: the NPV of project A is .163 million rupees, while that of project B is .127 million rupees. Project A would be chosen. Suppose now the discount rate is .07: the NPV of project A is .057 million rupees, while that of project B is .075 million rupees, making the latter the preferred project. Clearly, when the time profile of projects differs considerably, the ranking can depend upon the discount rate, something which the IRR cannot accommodate.

This implies that if the time profiles of projects differ, the IRR cannot be used to determine which one maximizes NPV, even among those that are of similar scales. If the choice is between a long-lasting project and one with a short time horizon, the IRR is prone to be unreliable.

Capital Budgeting

Suppose that, for whatever reason, the policy-maker has an upper limit on the capital budget that can be used for the projects under consideration. The budget can be used for financing various combinations of projects. In principle, the choice of projects is straightforward: choose the combination of projects within the budget limit which maximizes the total NPV of the projects combined, where the NPV calculation is precisely the same as before. This might, of course, entail not undertaking the one which has the highest individual NPV in order that the aggregate NPV is the highest possible. As before, use of the benefit-cost ratio and the IRR criterion will generally be unreliable.

The discount rate might be thought to be an issue here, since the capital budget is fixed, so there is no opportunity to borrow and lend. But, as we discuss below, given our assumption of welfarism, the discount rate for project evaluation, the so-called social discount rate, is the rate at which the benefits and costs of the project are discounted by those individuals in the economy actually obtain them. In the absence of externalities, and assuming that households are free to borrow and lend on capital markets, the social discount rate is the after-tax interest rate faced by households on capital markets.

Although the principles of project evaluation when there are capital budgeting constraints are clear (and not really any different from project evaluation in the

---

yields a lower NPV.

unconstrained case), there are nonetheless a number of conceptual issues that must be dealt with in practice, including the following.

Unused Capital Funds If the collection of projects do not exhaust the capital budget allotted, the issue of what becomes of the unused funds is relevant. If they revert to general revenues and serve to relax the government's overall budget constraint, this must be taken into account. In effect, the excess burden of whatever public financing is available must be incorporated into the project evaluation in a manner discussed below. Projects which use less funds will naturally incur less excess burden on this account. In other words, the procedure for taking account of the actual amount of funding for various options is the same as for project evaluation in the absence of the capital budget constraint. The latter simply puts an upper bound on the capital available.

Multi-Period Capital Costs The evaluation will need to take account of the extent to which different projects incur capital costs over a period of years, and how the capital budget constraint deals with that. Again, nothing new in principle is involved here. As long as all costs and benefits are appropriately accounted for, including the cost of public funds, the only constraint imposed by the capital budget is a restriction on the amount of funds available over time. The capital requirements for different projects may have different time profiles. As long as they are properly costed in the periods in which they are incurred, there should be no problems. The evaluator must still choose the combination of projects which maximizes the aggregate NPV and does not violate the capital budget allotted.

Future Project Capital Requirements Similar considerations apply with capital funding that may be required for expansion of replacement investment some periods down the road. To the extent that capital constraints apply to these, they will obviously have to be satisfied.

The Treatment of Inflation

If the general level of prices is rising over time, market interest rates will reflect that. For example, if the inflation rate is π and it is fully anticipated, the nominal discount rate i will differ from the real one r in a given period according to:

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

(This, of course, neglects any taxes that might be payable on capital income.) In principle, project evaluation should include only real benefits and costs — purely inflationary changes in values should not be included.

There are two equivalent ways to ensure this. The first is to conduct the project evaluation entirely in nominal terms. All benefits and costs would be evaluated in current rupees, using nominal prices projected for each future period. And the nominal social discount rate should be used. The alternative is to use constant-rupee prices, obtained by deflating current-valued ones by the price index relative to some base period, but to discount the flow of net benefits and costs using the real discount rate r. It is straightforward to show that these two procedures will yield the same NPV.[6]

The prescription is perhaps easier than its application. Future inflation rates are difficult to estimate, especially expected ones. What is important is that consistent procedures be used. The use of nominal discount rates must not be mixed with benefit and cost evaluations which do not include an allowance for inflation. Perhaps the safest procedure is to use constant rupee prices so as to avoid the need to estimate future inflation rates.

---

[6] To see this, note that the relation between real and nominal prices is given by $p_t = (1+\pi)^t p_0$, where p_t is the nominal price level in period t, while p_0 is the real price using a base year of zero. The NPV using nominal prices can be written:

$$\mathrm{NPV} = \sum \mathrm{p_t X_t} \big/ [(1+\mathrm{r})(1+\pi)]^t = \sum \mathrm{p_0 X_t} \big/ (1+\mathrm{r})^t$$

where X_t is the vector of net benefits in year t. Note that $(1+\pi)^t$ is the price index for period t.

Terminal Value

A vexing issue in project evaluation is identifying the end of a project's useful life, that is, the terminal period. The terminal period determines the number of periods T over which the project may be evaluated. Presumably, the project should cease once its future discounted present value falls to zero. The problem is that this may be far into the future, where projections become less reliable.

Setting aside estimation problems, whatever terminal date is chosen, there will undoubtedly be some fixed capital left over. To the extent that the capital has some value, its so-called *scrap value*, that value should be included as a benefit of the project. This will also be difficult to measure. If the capital can be put to another use (e.g., office equipment, vehicles), its value in that use should be included as a benefit. If the capital has no alternative use, the scrap value would consist of the value of the materials that could be salvaged.

It is even conceivable that scrap value could be negative. For example, the site of the project may leave an environmental or health hazard if it is not cleaned up. This clean-up should be treated as part of the cost of shutting down the operation.

CHOICE OF A NUMERAIRE

An issue that distinguishes various approaches to project evaluation concerns the unit of measurement, or *numeraire*. The numeraire serves as the standard against which all other benefits or costs are evaluated, given the distortions that exist in the economy. It is important to recognize that the choice of numeraire is basically arbitrary in the sense that it does not affect the outcome of the evaluation: project evaluation done under any numeraire can be converted into that for any other by using the appropriate conversion factors.

In the literature, two approaches have been predominant. The first, more traditional, approach is to value all benefits and costs in terms of current consumption expenditures by households. Although we shall outline later how to value particular sorts of items, it is worth mentioning some key types of benefits and costs that require special attention.

Present Versus Future Consumption

Naturally, items that occur in later periods must be converted to the current period by a discount factor. In particular, having converted all within-period inputs and outputs into their values in terms of consumption, the discounting must use a consumption discount rate (i.e., the rate at which household do transform present into future consumption on capital markets).

Foreign Exchange

When foreign exchange markets are distorted by tariffs and other trade measures, the market exchange rate no longer reflects the economic cost of converting foreign products into domestic consumption. A shadow price of foreign exchange must be determined which incorporates the effects of the distortions on the true opportunity cost to the economy of purchasing acquiring foreign exchange to purchase imported goods, or conversely to sell foreign exchange acquired from the sale of domestic goods abroad. The foreign price of all traded commodities involved in a project must be converted to domestic consumer prices using the shadow price of foreign exchange.

Public Financing

Public funds cannot be treated as having the same value as funds in the hands of consumers because it is costly to reallocate funds from the households to government. Given that there will be a deadweight loss of using the tax system to do so, a rupee in the hands of the government will be more valuable than in the hands of households. The appropriate conversion factor for converting public funds to private fund is the marginal cost of public funds (MCPF) — the opportunity cost of transferring a marginal rupee from the private sector to the public sector. To the extent that the project entails changes in public sectors revenues, these must be valued at the MCPF.

Investment Relative to Consumption

By the same token, because of capital market distortions, a rupee's worth of investment is worth more than a rupee's worth of consumption: the former would yield a stream of consumption whose present value exceeds one rupee. That suggests that to the extent that the project crowds out investment or enhances it, the effect on investment must be valued at the opportunity cost of investment.

Distributive Equity Considerations

It might be judged that a rupee of consumption is worth more to persons of low income than to persons of high income. If it is desired to incorporate such judgments into the project evaluation, the numeraire should specify consumption in the hands of a particular income level of household. Consumption accruing to other households must therefore be discounted by the appropriate distributive weight.

The use of current-period household consumption (perhaps in the hands of a particular income group) as the numeraire therefore entails using conversion factors for future consumption, for changes in foreign currency use, for net public sector funding, for changes in investment, and possibly for consumption accruing to different income groups. What results is a measure of the NPV of the project measured in present consumption to the benchmark income group.

The second approach is that first advocated by Little and Mirrlees (1968) for the OECD, but since widely used by the World Bank (e.g., Ray (1984)). Their numeraire is foreign exchange in the hand of the government. 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. Also, domestic consumption changes are considered to be less valuable than rupees in the hands of the government for a couple of reasons. First, as above, there is a deadweight loss involved in diverting funds from the private sector to the public sector. But, second, it is reckoned by Little and Mirrlees that funds in the hand of government will be available for investment, which as before is more valuable than consumption because of capital market distortions. So a consumption conversion factor is required to evaluate domestic consumption benefits in terms of government revenues. And, such distributive weights as are deemed necessary are also used to convert consumption of different income groups into that of the benchmark group (taken to be the lowest income group).

The upshot, to repeat, is that the two procedures should give the same result, as should a procedure which uses any other numeraire. In what follows, we shall implicitly follow the traditional approach and use present-period consumption as the numeraire.

COST-EFFECTIVENESS VERSUS BENEFIT-COST ANALYSIS

The most complete and informative type of project evaluation estimates the NPV of the benefits and costs of all alternatives being considered. This can be a mammoth task. In some circumstances, it is either sufficient or only possible to measure project costs. For example, if one is comparing alternative methods of delivering the same services, it may be necessary only to measure the costs of the various methods, that is, to conduct a *cost-effectiveness analysis*. Thus, one might be interested in comparing the costs of administering a tax or tariff collection system. Or, one may be comparing the costs of different ways of patrolling one's borders. Assuming that the same services are accomplished by various alternatives, a comparison of PVC's should suffice to pick out the socially desirable one. Of course, as long as PVB's are not estimated, it is not possible to say whether any of the alternatives has a positive NPV.

If benefits are conceptually impossible to measure, one has no choice but to measure only the costs. This might be true even if the benefits differ among projects. This does not necessarily render cost-effectiveness analysis of no use. The policy-maker can be presented with the present value of the costs of different alternatives, and the policy-maker can then assume responsibility for deciding which alternative, if any, should be undertaken.

In principle, the analog of cost-effectiveness analysis might need to be done from the benefit point of view. If costs cannot be measured, it might still be informative to compare the benefits from various options. More generally, if some, but not all, of the costs or benefits cannot be measured, it may still help to policy-maker to know the magnitudes of those that can. Some information is typically better than none.

SENSITIVITY ANALYSIS

Rarely will all parameters be known with full confidence, especially those which are not reflected in market values, which require value judgments, or which will occur in the future. In these cases, presenting calculations using different parameter values will indicate how sensitive the results are to the reported values. The policy-maker will at least know for which parameter values judgment becomes important. There are no general principles for conducting sensitivity analyses. Apart from experimenting to see which parameters are critical for the results of the evaluation, it is also useful to set out the evaluator's judgment of the most likely set of parameter values, as well as lower-bound and upper-bound calculations.

III. VALUING INPUTS AND OUTPUTS

The core problem of project evaluation involves putting monetary values on the various benefits and costs of the project. These should reflect the willingness-to-pay. Benefits and costs can come in a variety of different forms, including the purchase or sale of products and factors of production on markets, the provision of non-marketed benefits or costs such as externalities and intangibles, and the net benefits arising from indirectly affecting resources allocated on other markets, which themselves are distorted. We consider each type of benefit or cost separately.

MARKET INPUTS AND OUTPUTS

Projects that involve expenditures on goods and service (as opposed to, say, transfers or regulations) will typically involve purchasing some primary or intermediary inputs on markets, and perhaps selling some outputs on markets (e.g., electricity, water, transportation services). At the same time, markets may well be distorted. They may have taxes, tariffs or subsidies imposed on them; they may be monopolized; or, they may simply be functioning imperfectly. The measurement of benefits and costs of marketed items involves taking account of these distortions. This results in a social, or *shadow*, value or cost for a marketed item which is typically different from the market value or its consumer surplus.

To understand the meaning of a shadow value, consider the case of a project which uses an input X purchased on a market in which the demand price exceeds the supply price. Let p be the supply price (marginal cost) and $q = p + t$ be the demand price, where t is the distortion. For concreteness, think of the distortion as a tax imposed on the input. Figure 1 (below) depicts the market for X. The demand curve D(q) shows the quantity demanded at various demand prices, while the curve S(p) shows the quantity supplied at various supply prices. By adding the distortion t vertically to the supply

Figure 1

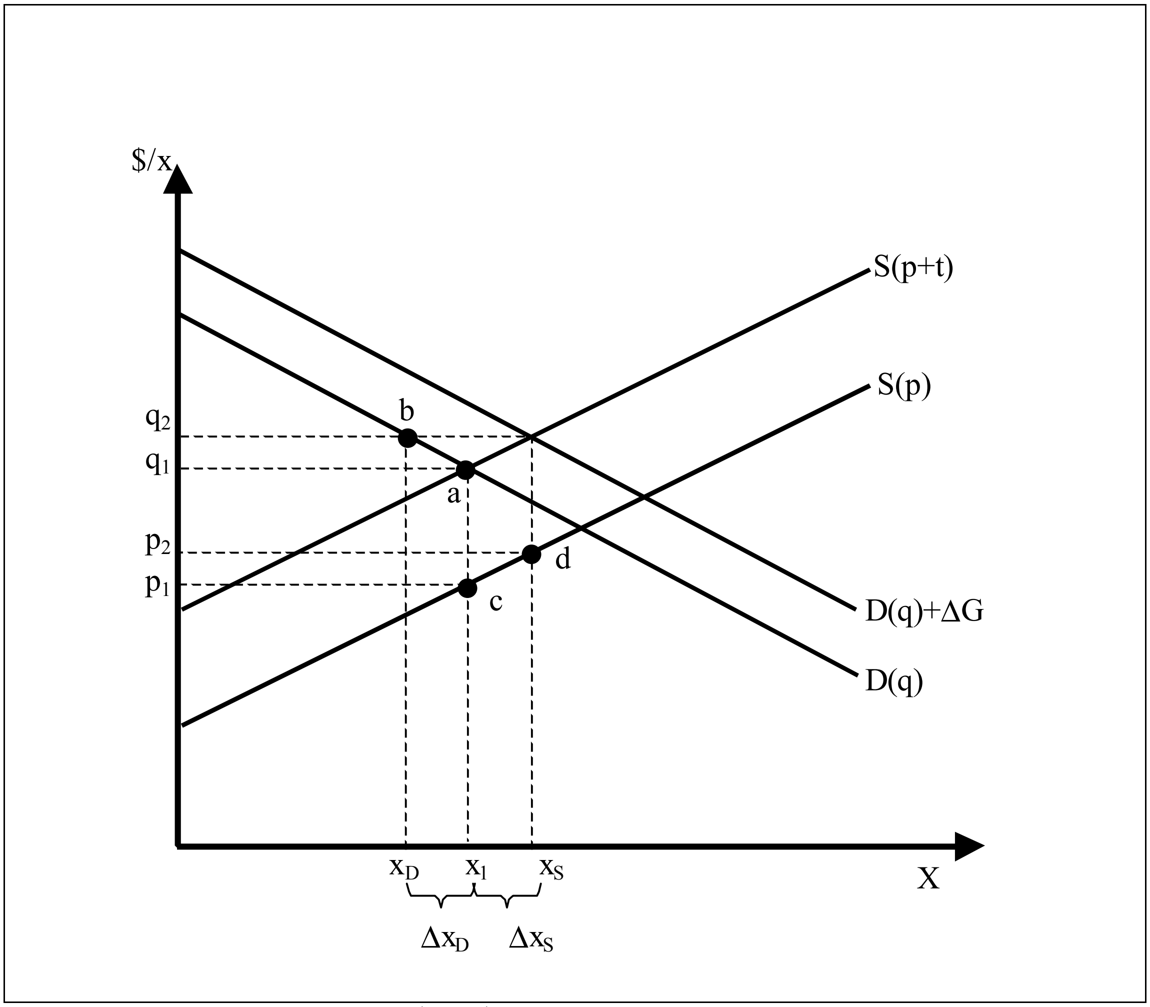

curve, we obtain the curve $S(p+t)$ showing the amount that would be supplied at various demand prices. Market equilibrium occurs at the output X_1 where the supply price is p_1 and the demand price q_1. Suppose now the project purchases an amount ΔG from the market. The demand curve will shift rightward by ΔG, causing the supply and demand prices to rise to p_2 and q_2. As can be seen from the diagram, the project demand ΔG is satisfied partly from an increase in supply, ΔX_S, and partly from a reduction in demand, $-\Delta X_D$. This carries with it an opportunity cost consisting of the cost to suppliers of supplying the additional amount ΔX_S, or X_1cdX_S, and the reduction in benefit to demanders from forgoing purchase of the amount ΔX_D, or X_1abX_D. The sum of these two items gives the shadow value of the input used in the project.

If the project is relatively small, so will be the price changes. Then the shadow value can be written:

$$X_1\text{ab}X_D + X_1\text{cd}X_S = -q\Delta X_D + p\Delta X_S$$

The shadow price per unit of input purchased by the project equals is given by:

$$s = \frac{p\Delta X_S}{\Delta G} - \frac{q\Delta G_D}{\Delta G}$$

This states that the shadow price s is a weighted average of the supply and demand prices, with the weights being the share of the project requirements which are obtained from an increase in supply and a reduction in demand. It is sometimes referred to as Harberger's weighted-average shadow price rule. It can be applied both to the purchase of inputs and the sale of outputs, where in the latter case the project output displaces market supply and induces market demand. In the special case where the supply price (marginal cost) is constant (the elasticity of supply is infinite), the shadow price is simply p; while if the demand price is constant (demand is infinitely elastic) the shadow price is $q = p + t$.[7]

An important class of cases in which prices might be fixed is that of a small open economy which faces fixed world prices. In this case, the shadow price of either a tradable input or a tradable output simply reflects the prevailing world price measured in foreign currency terms. The sale of a traded commodity (even on the domestic market) ultimately gives rise to a supply of foreign exchange according to the world price of the good sold, while a purchase gives rise to a demand for foreign exchange. There are no direct effects on markets for non-traded products. But, if foreign exchange markets are distorted, the conversion of increments of foreign exchange into domestic consumption equivalent values requires a shadow price of foreign exchange.

The Shadow Price of Foreign Exchange

Construction of a shadow price of foreign exchange is analogous to the above procedure. Assuming that trade must balance and the exchange rate is determined as a market clearing price (i.e., is flexible), the demand for foreign exchange reflects the domestic purchase of imports, while its supply comes from the sale of exports. If there were a common tariff at the ad valorem rate τ, and if e is the rupee price of a unit of foreign currency (the market exchange rate), the supply of foreign exchange will depend upon e, while the demand will depend on $e(1+\tau)$. A project which uses one rupee worth of a tradable product will shift the demand for foreign exchange to the right. A similar argument as above then leads to the shadow price of foreign exchange being given by:

$$s = \frac{e\Delta X_S}{\Delta G} - \frac{e(1+\tau)\Delta X_D}{\Delta G}$$

where ΔG is the demand for foreign exchange generated by a project. If there were several different tariff rates τ_i for different products, the shadow price of foreign exchange would become:

$$s = \frac{e\Delta X_{Si}}{\Delta G} - \frac{\sum_i e(1+\tau_i)\Delta X_{Di}}{\Delta G}$$

For project evaluation purposes, this shadow price must be applied to the world price in domestic currency of all traded goods. No domestic taxes need be taken into account.[8] Of course, if the exchange rate is not purely flexible, or if the domestic economy has market power in international markets, those things must be reflected in the shadow price.

---

[7] These weighted-average shadow prices might be augmented by distributive weights if desired. We discuss the use of distributive weights later.

[8] Moreover, if equity is a concern, distributional weights need not be attached to traded inputs and outputs of items of importance to, say, low income groups since they do not directly affect the domestic

The Shadow Wage Rate

One final application of the shadow pricing of marketed items concerns the price of labor. Labor markets not only have significant taxes imposed on them, but they are also prone to imperfections, especially unemployment. This implies that there will be a difference between the demand price for labor paid by employers and the opportunity cost of workers supplying the labor. In principle, a weighted-average shadow price can be devised and used. But, there are some complicating factors. The supply price of labor may be difficult to measure. For example, in the presence of involuntary unemployment, it will be less than the after-tax wage rate. It should, however, exceed zero given that leisure time has a value, but no market price will correspond to it so its measure will be imprecise. Also, wage differentials may exist for the same labor in different locations. To the extent that this reflects costs of moving, no particular problems arise. The supply price of labor is the wage paid in the new location, since that includes compensation for the costs of moving.

Wage differentials may be taken to reflect a segmented labor market where, for institutional reasons, wages are higher in one location (say, urban areas) than in another (say, rural areas), $w_U > w_R$. A worker from the rural area who is hired in the urban area at a wage w_U has only an opportunity cost of w_R, his output in the rural area. In this case, it is sometimes argued that the shadow wage should be a weighted average of w_U and W_R, where the weights correspond with the proportions in which hired workers are drawn from elsewhere in the urban area and from the rural area. Indeed, w_R might even be taken to be zero if there is excess labor in the rural sector, as in the Little-Mirrlees approach and the UNIDO *Guidelines* (Dasgupta, Marglin and Sen (1972)).

Others find this argument unconvincing. Harberger (1971b), for example, using a variant of the well-known Harris and Todaro (1970) model, argues that the wage differential between the urban and the rural sector represents an equilibrium phenomenon, just like the wage differential between two locations on account of the cost of moving. To see the argument in its simplest form, suppose W_U is artificially above the market-clearing level for institutional reasons, but that w_R is free to adjust as workers

---

consumption of those goods. In the Little-Mirrlees approach, which uses foreign exchange as the numeraire, this make the valuing of traded commodities particularly easy: world prices in rupees.

move. In the absence of moving costs and assuming risk neutrality, workers will migrate until their expected urban wage equals their rural wage, $pw_U = w_R$, where p is the rate of unemployment, and urban jobs are assumed to be filled randomly. Suppose a project creates jobs in the urban area, and that they are filled from the pool of unemployed. Each job created will induce a migration from the rural area of 1/p workers, enough to ensure that the equilibrium condition $pw_U = w_R$ is satisfied. The opportunity cost of attracting these workers is w_R each, or w_R/p in total. By the equilibrium condition, this is just w_U, the wage paid to a worker who has been hired. Thus, market wages become the shadow wage rate. Thus, it is important to be sure of how the labor market functions before settling on a shadow wage rate.

If equity is a concern, it will be particularly important to incorporate distributive weights into the shadow wage rate, given that most of the income of lower income workers will be consumed by them. This will be the case whichever shadow wage formulation is used. Again, we shall return to the issue of distributive weights below.

Special Problems with Capital Inputs

The costing of inputs of a capital nature gives rise to some additional problems over and above the need for shadow pricing discussed above. These arise because of the durable nature of capital assets: outlays must be made for them before they generate a stream of benefits. This gives rise to two sorts of problems. First, capital acquisitions must be financed up front, either by government funding or by making use of capital markets. Either source of finance involves distortions, which implies that the opportunity cost of financing exceeds the amount of funds required. Though this problem is endemic to capital acquisition, it is more general than that. Projects may generate insufficient revenues even for ongoing costs. We defer until later the general problem of the opportunity cost, or shadow price, of project financing.

We deal here with the second problem, which is how to measure the costs of capital inputs given that their use is stretched over a number of periods into the future. As in the case of private-sector project evaluation, two methods of capital cost accounting could potentially be used — cash flow or accrual. Cash flow accounting involves simply including all outlays and inflows as they occur. Capital expenditures are costed in full ('expensed') when they are made, with appropriate shadow pricing used if they are purchased from distorted markets as discussed above. Capital expenditures must include all gross investment expenditures — additions to a project's capital stock, replacement investment, as well as any scrap value salvaged at the end of the project's useful life. Costs of financing and ongoing depreciation do not enter directly into the calculation of costs over and above the initial cash flow expenditures: that would be double counting. They may enter indirectly to the extent that financing gives rise to excess burdens, or to the extent that capital that has depreciated has been replaced.

Accrual accounting attempts to attach costs to the use of capital in the future rather than at the time of initial outlay. These costs are of two sorts — depreciation and financing costs. Depreciation is meant to reflect how much capital is 'used up' in each period of use, either due to obsolescence, wear and tear, or due to changes in the relatively price of the asset. In other words, it measures the extent to which the value of the asset falls over the period. The financing costs represent the forgone interest associated with holding real capital rather than putting the same funds into the financial capital market. Again, one must not mix elements of cash and accrual accounting. If the latter is used, no capital expenditures of any kind should be treated as costs when they are incurred; rather they are costed as they are used up.

Cash and accrual accounting for capital costs are alternative ways of presenting the same information. In principle, the present value of the accrual costs of a capital investment should equal its cash flow (or the present value of its cash flow for a sequence of investment expenditures). But, accrual accounting is inherently more difficult to use since it involves attributing a depreciation sequence to the use of capital, something which cannot readily be observed from market prices. Moreover, the principles of shadow pricing are much less transparent when using accrual accounting. For that reason, cash accounting is typically used for project evaluation in the public sector. The private sector prefers to use accrual accounting because of the information it provides to shareholders. It indicates the period-by-period profitability of a firm which is engaged in a multitude of ongoing projects. Presumably if financial accounts were on a project-by-project basis, cash flow accounting would serve at least equally as well.

INTANGIBLES AND NON-MARKETED INPUTS AND OUTPUTS

Public projects by their very nature often produce benefits or generate costs for which market prices are not readily available, or which are intangible and cannot readily be priced on markets. Examples include health and safety improvements, environmental improvements or degradation, time saved travelling, and the acquisition of new knowledge or skills. In some projects, these intangible or non-priced benefits are among the most important outputs. Their valuation should be guided by the same principles as above — willingness-to-pay for benefits, or the analog for costs, willingness-to-accept. The difficulty is of course that no guidance is available from market prices, so the monetary values must be inferred by other means.

Two common means can be used for evaluating intangibles. The first is to use the method of *hedonic pricing*, which is to use households' observed behavior elsewhere in the economy to reveal the value they implicitly place on intangibles. The second is to use survey techniques to ask a sample of households directly what value they place on the intangible under consideration. Consider some examples of each in turn.

Value of Time Saved

Public transportation projects such as roads, airports, bridges and public transit facilities often have as their main objective the saving of time by users of the project as well as by users of alternative means of transportation. The value of time saved travelling depends upon the alternative uses to which the time will be put — whether to productive work or to leisure or non-market activities. In the case of the former, the value of time saved travelling might be the marginal productivity of time spent working, which in a competitive setting can be valued at the tax-inclusive wage rate. This presumes: i) that labour markets are competitive; ii) that workers are indifferent between time spent commuting and time spent working; and iii) that there are no indivisibilities so that time saved can be put to productive use rather than leading to more free time for the worker. If one or more of these is violated, the calculation of time saved must be amended accordingly.

If time saved travelling accrues to households in the form of increased leisure, valuation is more problematic: the household is effectively substituting leisure time for commuting time, neither of which is readily measurable. The wage rate is of relatively little use here, as the following analysis shows. An individual will have different marginal benefits associated with time spent working (MB_W), leisure time (MB_L) and commuting time (MB_C). At the margin, if the individual can freely choose between leisure and working time (for given commuting time), the equilibrium choice will satisfy $MB_L = w + MB_W$, where w is the wage rate. Since MB_W is presumably negative, the value of leisure is less than the wage rate. If the transportation project substitutes leisure time for commuting time, the value of time saved will be $V_C = MB_L - MB_C$, which is even less than the wage rate assuming $MB_C > 0$.

The value of V_C cannot be observed directly, but must somehow be estimated from elsewhere. One method commonly used is to infer V_C using transportation mode choices elsewhere in the economy. If consumers have a choice between two ways of getting from point A to point B which differ in the time cost as well as in resource costs, one can statistically estimate the amount of money consumers are just willing to pay at the margin to take the faster mode of transport. That is an application of the hedonic pricing method.

Once a value of time is obtained, it can be used to generate a monetary measure of the benefits of any project which involves time saved. A transportation project will typically both divert traffic and generate an increase in travel. The benefit of the former will include the monetary value of time saved for all diverted trips plus any changes in real resource costs for diverted trips (fuel, capital equipment, etc.) The benefit of generated traffic requires an estimate of the new demand created by the transportation facility, which will also depend upon the value of time saved. Another example of a project for which the value of time is relevant concerns recreational facilities in remote areas, such as national parks. Users will include both those diverted from other sites and newly generated demand. The willingness-to-pay for the use of the new site will depend upon the value of time.

Value of Reduced Risk of Death

Another example of the hedonic pricing technique involves valuing the saving of lives due to a project. (Similar principles apply to reducing the risk of disease or injury.) Health care programs, safety regulations and transportation projects are all examples where this can be used. Again, 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. In other words, how much would consumers be willing to pay to achieve the given reduction in risk.[9]

In some cases, that valuation is implicit in the measurement of the benefits from using the project. For example, if travellers voluntarily choose to use a transport facility which carries with it a risk of accidental death, and if a demand curve for the facility has been estimated, the latter will include the value that travellers place on using the facility net of any costs associated with risk. But for some projects, the costs or benefits of changes in the risk of loss of life must be attributed separately. An implicit 'value of life' can be obtained by observing other situation in which households implicitly put a value on the risk. For example, different types of jobs have systematically different risks of death, injury, etc., and those risks ought to be reflected in market wage differentials. Statistical techniques may then be used to estimate an implicit or hedonic value associated with different degrees of risk on the job.

Of course, statistical techniques must be used and interpreted with care. There are likely to be very many factors which go to explain wage differentials, and it is important to control for the most important of these. Moreover, it is possible that different households have different degrees of risk aversion: the least risk averse will be willing to accept a lower wage differential to work in riskier jobs. In these circumstances, wage differentials may not measure the cost of risk to the average person.

Costs of Environmental Pollution

Transportation projects or industrial projects may cause various sorts of pollution to neighboring residents. For example, a new airport will increase noise levels in the vicinity. Estimates of the cost of noise might be obtained indirectly from property values. *Ceteris paribus*, property values should be lower in noisier locations. Once again, hedonic pricing techniques can in principle be used to obtain monetary measures of environmental costs. In principle, property values should reflect the monetary value of

---

[9] Evaluating reductions in the risk of death by *ex ante* willingness-to-pay, that is, without knowing precisely who will be saved, is not without controversy. Some would argue that as a society, loss of life should be evaluated from an *ex post* point of view since some persons will be saved for certain. This would give much larger values to each life saved. There will also typically be other benefits and costs associated

attributes associated with various locations. The trick is to control for all the various attributes so that the cost associated with the environmental cost at stake can be obtained. To be useful for the project, these estimates must be for circumstances similar to those of the project. In the case of noise pollution, estimates from other airport locations might be suitable.

Survey Techniques: Contingent Valuation

Data limitations may preclude the use of statistical techniques to obtain hedonic values for intangibles. In this case, other methods must be found to place a value on them. One way to do so is to conduct a survey. Rather than relying on households to reveal their valuations directly or indirectly through their market behavior, they could be asked directly through a survey. Those surveyed are typically asked how much they would be willing to pay for the good or service in question — their willingness-to-pay — or, if appropriate, how much they would be willing to accept to give something up. Thus, households might be surveyed to determine their willingness to pay to create a national park or to protect an endangered species. Or, if a new airport is being contemplated near a residential area, residents might be asked for their willingness to accept increased noise levels at various times of day.

Naturally, surveys must be constructed with some care to ensure that respondents fully understand the nature of the project being evaluated. But even so, there are several potential pitfalls with survey techniques. Two important ones are as follows. First, those who complete a survey may not be a representative sample of those who might be affected by the project. If the sample is relatively small, there may be a biased group in the sample. But a more likely problem is that of self-selection — those who choose to respond may be those who feel most strongly about it. A second general problem is that responses may not be truthful. Since there is no penalty for being dishonest, those who feel strongly about an issue will have an incentive to exaggerate their willingness-to-pay. Thus, contingent valuation methods must be used and interpreted with some care.

---

with project that reduce the risk of death or injury, such as loss of output, and psychic costs to friends and relatives. They are valued in the usual way.

SUBSIDIES

Some projects may provide benefits in the form of subsidies to users. Economists would generally oppose the use of subsidies because they interfere with market efficiency. Nonetheless, there are some circumstances in which policy-makers may be justified in using subsidies:

- there may be externalities associated with an activity, such as innovations or human capital improvements whose benefits partly accrue to others;

- subsidies might be justified to divert resources away from other distorted sectors, such as subsidized public transit to reduce road congestion;

- users might face cash flow problems in the purchase of needed inputs because of inadequate access to capital markets, an example being the use of fertilizer or irrigation in agriculture;
- governments, especially those in developing countries, might have limited instruments at their disposal to achieve redistributive and other social objectives, and must resort to subsidies as second-best policy instruments.

As defensible as these arguments might be if applied with care, there is always a danger that arguments for subsidization can be contrived on market failure grounds, but without quantitative estimates to support them. That is always a danger with second-best analysis in which market prices no longer reflect social values, and potentially anything goes. In any case, project evaluators might simply have to take as given a project as proposed by the policy-maker, and evaluate it as such.

The monetary value of subsidies is obtained from the standard use of CV's, or consumers and producers surpluses. Consider the example of, say, an input like fertilizer or irrigation provided at preferential rates to the agricultural sector. Figure 2 depicts the market for the input being subsidized. In the absence of the subsidy, the price of the input is p_1, and the demand by the agricultural sector is X_1. When the price is subsidized at the per unit rate s, demand rises to X_2 and the price to suppliers rises to p_2. The total benefit to the agricultural sector, or its aggregate willingness-to-pay, is the area beneath the demand curve, X_1acX_2. In the project evaluation, this would enter as a benefit to be set against the costs of the subsidy.

Figure 2

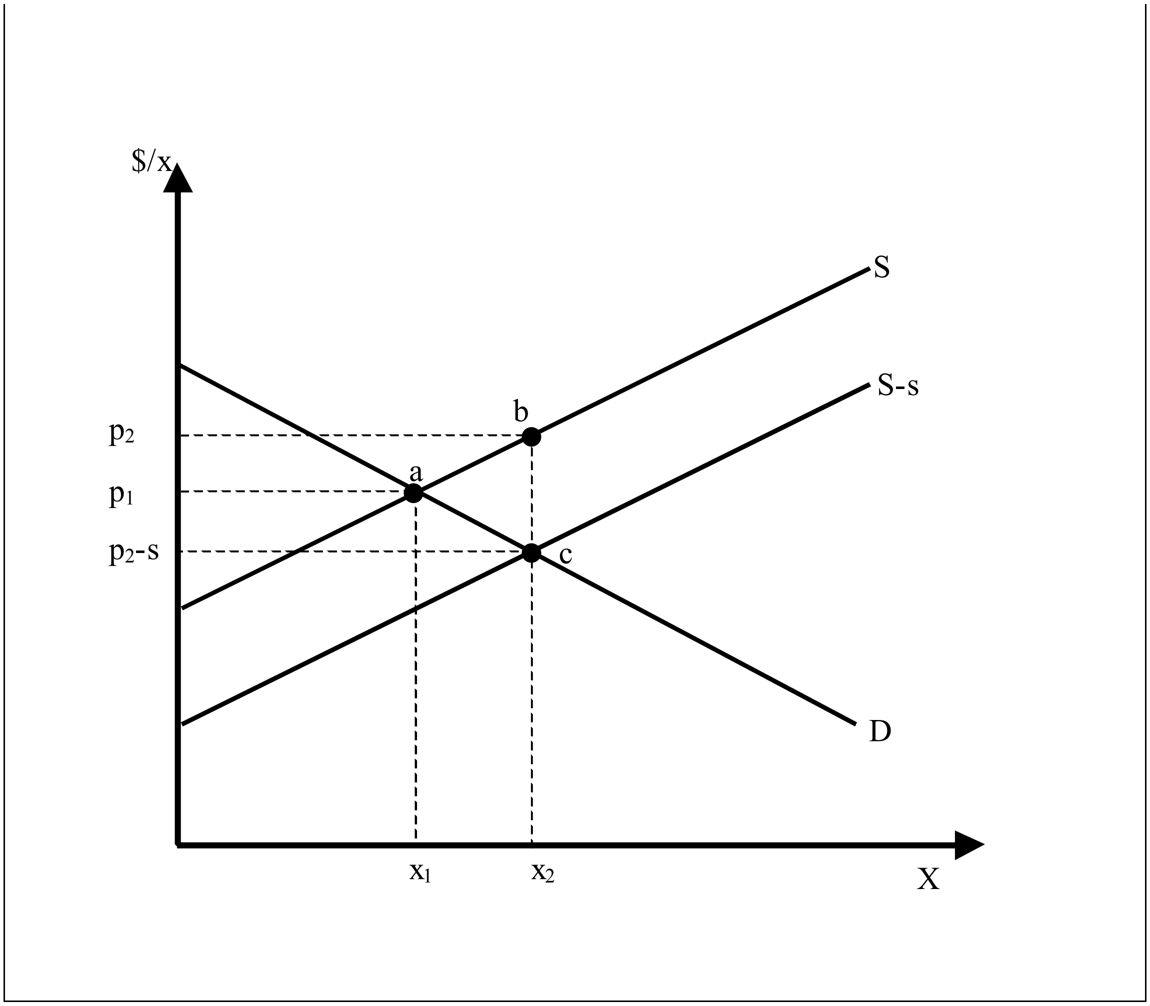

The standard case against subsidization can be readily seen from the figure. If the input comes from a competitive industry, the supply curve represents the schedule of marginal production costs. Then the cost of supplying the increment in demand is given by the area beneath the supply curve, X_1abX_2. This exceeds the benefit by the area abc, the standard excess burden of the subsidy. Thus, the project would be judged not to be socially beneficial unless there were other compensating benefits arising from shadow pricing, distributive weights, etc. For example, if the agricultural product were tradable, but the input being subsidized were non-tradable, it is possible that the premium put on X by using a shadow price of foreign exchange would be enough to make the project profitable in social terms.

An equivalent way to see the same point is to focus on changes in net benefits to the parties involved — consumers, producers and the government. The consumers surplus from the fall in price is the area $p_1ac(p_2\text{-}s)$. The producers surplus from the rise in price that suppliers receive is p_1abp_2. The cost to the government is the amount of the subsidy $p_2bc(p_2\text{-}s)$. Aggregating all these changes leaves a net loss to society of abc as before. The usefulness of this approach is that it identifies benefits and costs by parties involved so enables the evaluator to attach distributive weights if desired.

It should be emphasized that pure transfers of purchasing power from one household or firm to another *per se* should be typically attributed no value. But, they may affect the NPV of a project indirectly. For one thing, as with the subsidy, they may affect the allocation of resources in a distorted market, as in the above example, and as in the more general cases to be considered below. Thus, it is not the amount of the subsidy that is relevant and counts as a social cost, but how it affects resource allocation. Second, if distributive weighting is used, the social value of the transfer may be more to one or other of the donors and recipients.

INDIRECT BENEFITS AND COSTS

In measuring the shadow price of an input purchased on a distorted market, we concentrated on the implications of that distortion alone. But, if there are distortions in other markets elsewhere in the economy, induced changes in outputs on those markets also give rise to net benefit changes which must be accounted for in project evaluation. These are referred to as indirect benefits and costs. The general principle, due to Harberger (1971a), is an application of the theory of second best and is as follows. Consider an economy in which there are d sectors ($j = 1,\cdots,d$) that have a distortion between demand price, q_j, and supply price, or marginal cost, p_j. Let t_j be the distortion per unit of output ($t_j = q_j - p_j$) and let it be fixed for simplicity. Then, if a project causes changes in the output on any of these distorted markets on which it is not directly involved because of general equilibrium interactions, net benefits of the following sort must be included in the evaluation of the project:[10] $\Sigma_j\, t_j \Delta X_j$, where ΔX_j is the change in output on market j.

Figure 3

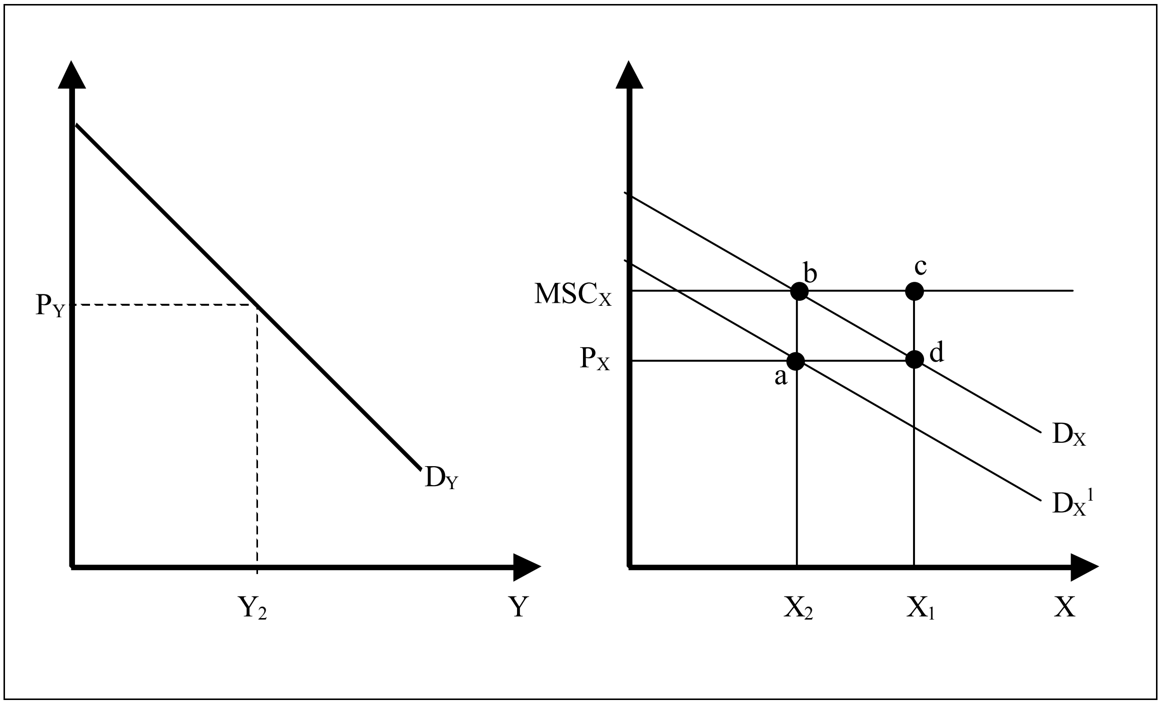

A simple example will illustrate. Consider an urban transit project that partly diverts traffic from congested expressways. Because of the congestion, the marginal social cost of a trip on the road exceeds the cost to the traveler: each traveler increases the cost spent travelling to other road users. As a consequence, there is too much road traffic. By diverting traffic, the urban transit system relieves congestion and generates an indirect benefit. Figure 3 (above) illustrates the indirect benefit. In the absence of the urban transit project, the demand curve for road trips is D_X in the right panel. Given a

---

[10] Formally, let the representative consumers' utility be given by $U(X_1,...,X_n)$. The change in utility from a change in demands is given by $dU = \Sigma(\partial U/\partial X_i)dX_i$. Consumers will set relative prices equal to their marginal rates of substitution so $q_i = (\partial U/\partial X_i)/(\partial U/\partial X_n)$, assuming good n is the numeraire. Then, we can write $dW = \Sigma\, q_i dX_i$, where dW is the change in utility measured in terms of the numeraire. ($dW = dU /(\partial U/\partial X_n)$). Now, suppose public project demands are G_i and market supplies are Y_i; then $dW = \Sigma\, q_i(dY_i + dG_i)$. Since $\Sigma\, p_i(dY_i) = 0$ by the economy's production possibilities frontier, we have $dW = \Sigma\, t_i dX_i + \Sigma\, p_i dG_i$, where t_i ($=q_i - p_i$) is the tax, or other, distortion. Finally, consider the change in a commodity used by or produced by a project, say, dG_k. The welfare change measure becomes $dW = \Sigma_{i\neq k}\, t_i dX_i + t_k dX_k + p_k dG_k$. The latter two terms can be shown to correspond with the value or cost of the public commodity dG_k evaluated at the weighted average shadow price as above. The first term is what we are calling the indirect effect, following Harberger. Further details on this are provided in Boadway and Bruce (1984). See also Drèze and Stern (1987).

cost per trip to travelers of p_X, which includes the time cost, fuel, vehicle operating costs and so on, the demand for road trips is X_1. But, because of congestion, the marginal social cost per trip is MSC_X. For simplicity, we assume that these costs per trip are constant. Next an urban transit project is introduced. As in the left panel, the demand curve for urban transit trips is D_Y. At a price of p_Y, Y_2 trips will be taken. The price includes all costs incurred by the traveller. Since the two types of trips are substitutes, the urban transit system will divert some travellers from the roads, causing the demand curve D_X to shift to the left. The number of road trips falls to X_2. The benefits of the urban transit system now include both the direct benefits, those calculated from the surplus generated in the left panel, and the indirect benefits. The latter are given by the area abcd in the right panel, which is the distortion times the changes in trips.

The same principles apply whatever the source of the distortion. Moreover, the fact the distortion exists suggests that explicit attempts could be made to increase the amount of traffic diverted. For example, the price of urban transit trips, p_Y, could be reduced below marginal cost even though that causes a deadweight loss in the urban transit market. The optimal second-best urban transit pricing policy would be that for which the incremental deadweight loss in urban transit just offsets the incremental indirect benefit created by a price reduction.

IV. THE SOCIAL DISCOUNT RATE

Once benefits and costs in each period are evaluated, they must be discounted to a common value. What discount rate should be used? The principle is clear enough: just like the principle of welfarism dictates that current values should be those reflecting individual households' willingness-to-pay, so it dictates that the discount rate should reflect the rate at which households value future relative to present benefits and costs. If capital markets were perfect, so that all households were able to borrow and lend at the going interest rate, they would organize their personal affairs so that the after-tax market interest rate they face is their intertemporal rate of substitution. Thus, the after-tax market interest rate would be the discount rate for project evaluation, the so-called *social discount rate*. (Of course, either a real or a nominal version could be used depending upon whether benefits and costs were measured in current or constant rupee values, as discussed above.)

In practice, there are a number of complications which must be considered in selecting the actual discount rate. Some of them are as follows.

Heterogenous Household Interest Rates

Households may face differing discount rates. For example, if capital income is taxed at the personal level, different households will be in different tax brackets, or may have access to different types of assets on their marginal borrowing or lending. As well, some households will be creditors and other debtors. If the borrowing rate differs from the lending rate, the social discount rate will be ambiguous. And, different households may face different interest rates because of differences in risk. For example, they may have differing wealth holdings, and therefor different abilities to provide collateral. The social discount rate will therefore have to be a compromise, given that the stream of project benefits and costs cannot typically be disaggregated by type of household.

Capital Market Imperfections

Various kinds of imperfections may exist on capital markets. Households may not be able to borrow freely on capital markets because of liquidity constraints. If they are quantity-constrained, the true discount rate faced by households will not be reflected in any market interest rate: it will typically be higher. Capital markets themselves may not be perfectly competitive, but may contain features of monopoly or of information asymmetries. Again, this makes it difficult to know the true rate at which households who are affected by projects discount future net benefits.

Externalities in Capital Markets

There may be external benefits associated with household saving. To the extent that households save for bequests, that saving may benefit others in the society who attach a value to the well-being of future generations. If so, there will be a free-rider problem associated with household savings, resulting in too little savings. Compounding this, there may also be externalities associated with the investment that is financed by household saving. The ‘new growth theory’ has emphasized that new knowledge comes with investment, knowledge whose benefits cannot be entirely appropriated by those doing the investing. In these circumstances, there will also be too little investment and hence too little saving.

Even if one acknowledges that these sources of externality exist and that as a consequence there is too little saving set aside for investment or for future generations, it is not obvious what implications, if any, this should have for the social discount rate. It has often been argued that the social rate of discount should be lower than the market interest rate on this account, and that would certainly be true from the point of view of supporting measures which increase saving above its market-determined level. But, the application of that argument to project evaluation is more tenuous. If the project involves benefits accruing to future generations, it might be argued that they should be discounted at a favourable rate on externality grounds. But standard projects which are not primarily intended to provide benefits to future generations have no particular claim to a low discount rate on these grounds. The rate at which existing households discount present versus future benefits accruing to themselves is the after-tax market interest rate they face, and that should be the discount rate for project benefits.

The project may well induce changes in the level of savings or investment in the economy. If so, it will create an indirect benefit or cost as a result of capital markets being socially inefficient. It would be appropriate to treat that as a benefit of the project in its own right, but not necessarily as one which calls for a lower social discount rate.

Other Arguments: Risk and Distributive Weights

Two further arguments might arise in discounting present versus future benefits of a project. The future will undoubtedly be uncertain, so that one does not know for sure precisely what benefits and what costs are likely to accrue as a result of the project. There are techniques for dealing with project risk, which may or may not involve the discount rate. We deal with those separately below.

By the same token, the treatment of present versus future benefits may be colored by the fact that those who obtain future net benefits are deemed to be more or less deserving than those who obtain current net benefits. If so, there may be a call for attaching distributive weights to future versus present beneficiaries. Again, we postpone discussion of this until the issue of distributive weights is addressed below.

The Numeraire

We have argued that the after-tax interest rate might be an appropriate project discount rate since it reflects the rate at which household in the economy discount present versus future consumption. In so doing, we have implicitly been assuming that the numeraire for project benefits and costs is current household consumption. As we have mentioned, different numeraires could be used, and the choice of numeraire should have no effect on the outcome of an evaluation. But, the choice of numeraire could affect the discount rate. For example, in the Little-Mirrlees approach, the numeraire is foreign exchange in the hand of the government. The discount rate then reflects the relative value of future versus present foreign exchange in the hands of the government. The determination of this can be quite complicated if it is assumed, as do Little and Mirrlees (1974), that governments can use these funds for investment and there are significant capital and foreign exchange market distortions. For a good discussion, see Ray (1984), and also Squire and van der Tak (1975).

V. OPPORTUNITY COST OF FINANCING

We have mentioned that pure transfers of funds among households, firms and governments should themselves have no effect on project benefits and costs. But if such transfers occur through distortionary means, or if they induce changes in outputs on distorted markets, they will have efficiency consequences. One important example of these induced welfare effects concerns the use of public funds to finance projects. Two issues are involved here. First, the transferring of funds from the private to the public sector through taxes is costly if distortionary taxes must be used to facilitate the transfer. That being so, account must be taken of the additional costs incurred in the financing of projects out of public funds. Second, projects which involve significant capital costs may be (justifiably) financed by borrowing from capital markets. Given that capital markets are distorted, this sets in train welfare costs which must be included as a cost of undertaking the project. We discuss these two in turn. It should be borne in mind that the costs being identified here are over and above the benefits and costs associated with project outputs and inputs already discussed.

Figure 4: Market for Composite Consumption Good

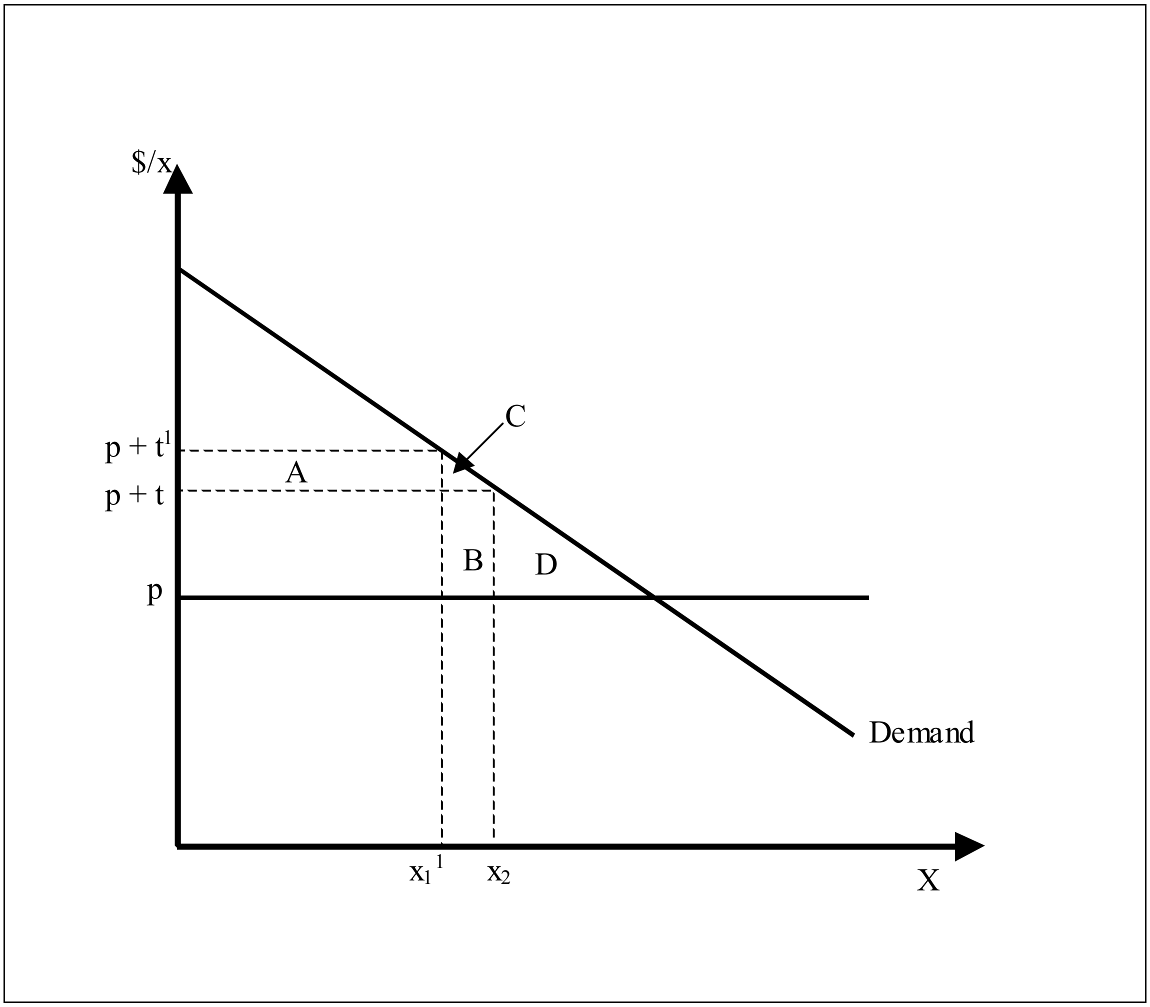

The Marginal Cost of Public Funds

Many projects in the public sector do not cover their costs; they require public funds. The problem is that obtaining a rupees worth of funds through taxation costs more than a rupee. That is because there is a deadweight loss or excess burden associated with raising revenues through distortionary taxation.[11] In evaluating projects, what will be important is the deadweight loss at the margin. Given that the deadweight loss is typically convex in the tax rate — it is approximately proportional to the square of the tax rate — the marginal deadweight loss can be significantly higher than the average deadweight loss per rupee of the tax system as a whole. The marginal cost of public funds (MCPF) is

[11] Even if the funds are not obtained from current taxes, there will be a deadweight loss involved. Funds raised through borrowing will induce a deadweight loss. Since borrowing is simply postponed taxes, the deadweight loss will be incurred later on when the loan is eventually repaid. As well, borrowing through

simply defined as one plus the marginal deadweight loss of raising additional tax revenues. If the MCPF is high, the hurdle rate of return that a project which relies on public funding would have to achieve can be substantially higher than for private projects.

To understand the MCPF, consider Figure 4 (above), which depicts the market for, say composite consumption, denoted X. In the absence of the project being evaluated, there is a per unit tax of t on consumption. It could be interpreted as a general sales tax, or as a general payroll tax. For simplicity, we assume that the supply price of composite consumption is fixed at p, so consumers face a price of p+t. If the project were introduced, the tax rate would have to rise to t′ to finance it. We can identify the various costs and benefits of the additional financing required. The change in tax revenue, ΔR, is simply area A minus area B. The deadweight loss from the tax before the tax increase is the area D. The change in deadweight loss from the tax increase is then given by C+B. The MCPF of transferring an increment of resources from the private to the public sector by increasing the tax rate is given by one plus the incremental deadweight loss per rupee of revenue, or:

$$
\begin{aligned}
MCPF &= 1 + \frac{\Delta D}{\Delta R} \\
&= 1 + \frac{(B + C)}{(A - B)}
\end{aligned}
$$

For small changes, we can neglect the small triangular area C. Then the MCPF may be written:

$$
\begin{aligned}
MCPF &= 1 + \frac{B}{(A - B)} \\
&= 1 - \frac{t\Delta X}{(X\Delta t + t\Delta X)} \\
&= [1 + t\Delta X / X\Delta t]^{-1} \\
&= [1 - \tau\eta]^{-1}
\end{aligned}
$$

printing money will create a deadweight loss as a result of inflation. Of course, if there are unemployed resources, this deadweight loss may be mitigated.

where τ is the tax rate expressed in *ad valorem* terms — $t/(t+p)$ — and η is the uncompensated price elasticity of demand for consumption.

The MCPF will typically exceed unity. In fact it can be considerably above unity. For example, suppose that the elasticity of demand is 2, and the tax rate τ is .25. Then, the MCPF is 2! The methodology can be extended to take account of more than one tax base, and even to take equity considerations into account (Ahmad and Stern (1991)).

The MCPF is relevant for attaching a shadow price to funds used to finance public projects. For each rupee of financing, a charge of (MCPF-1) should be attributed as a project cost: it is the excess burden induced by the requirement to use distortionary taxes. If lump-sum revenue sources were available, project financing would represent a pure transfer, and would therefore have no efficiency implications.

Figure 5: Effects of Borrowing

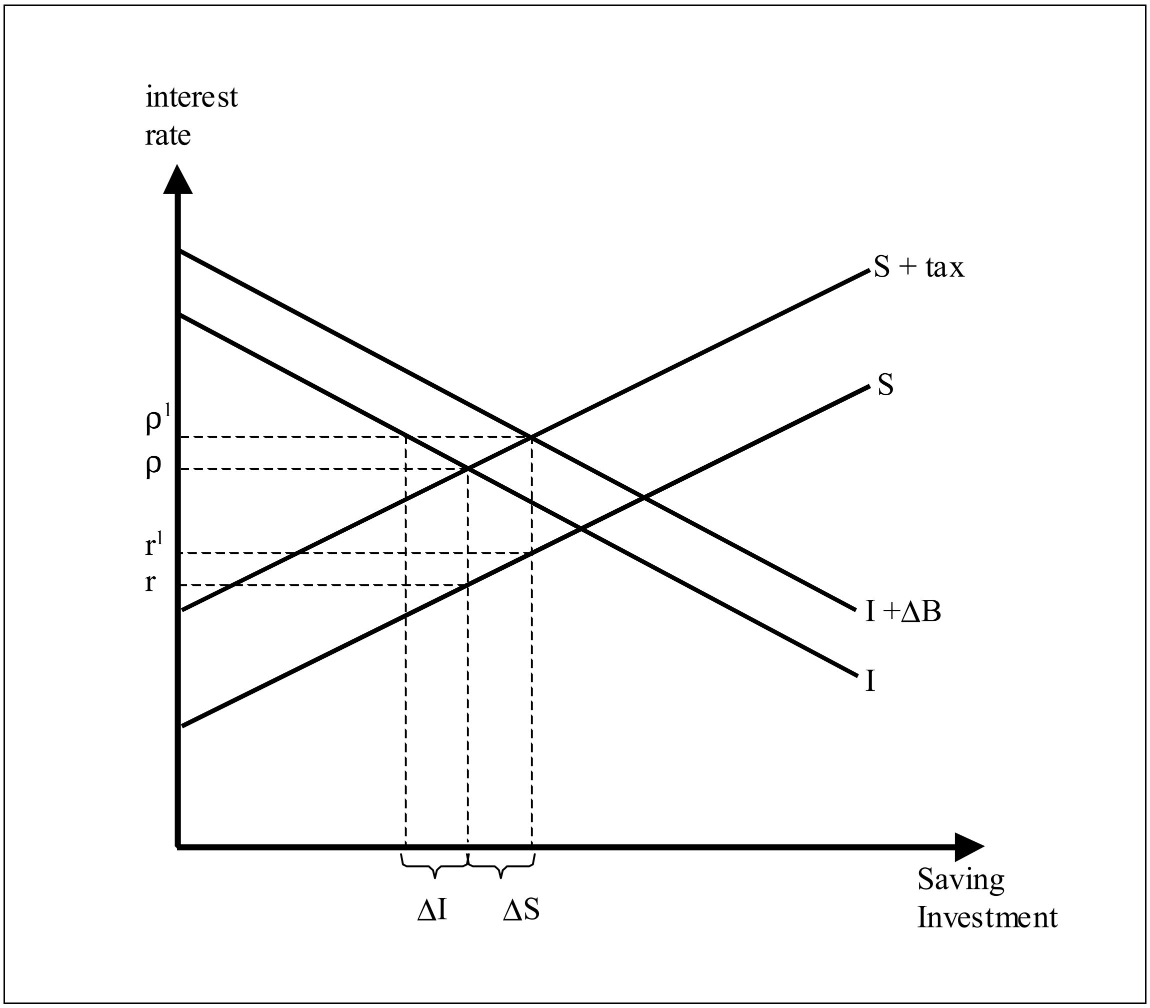

The Opportunity Cost of Borrowed Funds

Additional problems arise in a dynamic setting when capital markets are distorted. Consider the case of a project whose initial capital costs are financed by borrowing.[12] Let ΔB be the amount borrowed and invested. Suppose the capital market distortions arise from a tax on capital income. Then the gross (pre-tax) rate of return on investment, ρ, exceeds the net (after-tax) return on savings, r. Figure 5 (above) illustrates the effect of borrowing. The project borrowing of ΔB will represent a demand for funds on the capital market, and can in general come partly at the expense of reduced investment, -ΔI, and partly from induced savings, ΔS. As in our discussion of the weighted-average shadow pricing problem above, the opportunity cost of the additional financing depends on the opportunity costs of the forgone investment and induced saving.

The opportunity cost of an additional rupee of saving, measured in terms of current consumption which we take to be our numeraire, is simply one rupee. This is the amount of current consumption forgone. To obtain the opportunity cost of a rupee's worth of forgone investment, we need to identify the stream of consumption it would have yielded. The simplest case to consider is that in which the returns to the forgone investment would have been entirely consumed. One rupee's worth of investment with a rate of return ρ will then yield a stream of consumption benefits of ρ in perpetuity. The present value of this perpetual stream is then ρ/r, which is then the opportunity cost of one rupee's worth of forgone investment. Note that because of the capital market distortion, $\rho/r > 1$ — the opportunity cost of a rupee's worth of forgone investment exceeds one dollar.

Taking these two opportunity costs together — unity for savings and ρ/r for investment — we can construct the social opportunity cost per rupee of borrowing (SOC):

$$SOC = \frac{\Delta S}{\Delta B} - \frac{(\Delta I/\Delta B)\rho}{r}$$

[12] This example draws on Feldstein (1972a), which in turn draws on a seminal paper by Marglin (1963).

The SOC is used in a manner similar to that of the MCPF just discussed. In this simple case, the amount that would have to be added as a cost to the project at the time of borrowing would be simply (SOC -1)ΔB. If the debt was paid down later rather than being held in perpetuity, there would be a net welfare gain determined by the difference between SOC and unity.

This simple case illustrates the principles involved in determining the excess burden associated with providing debt financing to a project. (Naturally, if tax finance induces changes in investment and /or savings, a similar calculation could be carried out for that.) More complicated situations can readily be imagined. One simple extension is to suppose that a fraction σ of the returns to investment are re-invested, the remainder being consumed. Then a rupee of investment will accumulate at the rate $\sigma\rho$ per period, so by time t, the amount accumulated will be $e^{\sigma\rho t}$, leading to an amount of consumption at time t of $(1-\sigma)\rho e^{\sigma\rho t}$. Discounting this stream of consumption to the present yields a present value of forgone consumption of $(1-\sigma)\rho/(r-\sigma\rho)$. Now, the SOC becomes

$$SOC = \frac{\Delta S}{\Delta B} + \frac{(\Delta I/\Delta B)(1-\sigma)\rho}{(r-\sigma\rho)}$$

Other assumptions could be made about the stream of consumption that would be obtained in the future from one rupee’s worth of current consumption.

The SOC obviously depends upon the extent to which the borrowed funds crowd out private investment as opposed to inducing increased investment. That depends on the responsiveness of savings and investment to the interest rate, on which evidence is rather mixed. But, generally, the more elastic is the supply of savings relative to the demand for investment, the greater will be the proportion of the debt coming from induced savings, and vice versa.

There is a special consideration which merits some attention, and that is that savings might be invested in foreign assets rather than in domestic investment in an open economy setting (which is, of course, the relevant one). Thus, project borrowing might come from three potential sources now — domestic savings, domestic investment, net foreign capital flows. The relevant SOC then depends upon the type of tax distortions in place. If a tax applies to domestic investment, the opportunity cost of a rupee of forgone investment will exceed one rupee for the reasons discussed above. If there is a tax on capital income earned by savers, that will be a distortion between the rate of return on savings and both the return on domestic investment and that on foreign assets. If project financing displaces the accumulation of foreign assets, the opportunity cost will again exceed unity using analogous arguments to the above. But, if no tax on capital income is applicable at the personal level, which may be approximately so in developing countries, there will be no distortion on the accumulation of foreign assets: the opportunity cost of displacing a rupee's worth of foreign assets will simply be one rupee, the same as for increased savings. Therefore, in a small open economy setting, where all marginal finance involves changes in foreign asset holdings, the SOC of one rupee's worth of project financing will just be one rupee — there will be no need to make an adjustment for the cost of financing.

The procedure described above is based on the net present value criterion, which seeks to identify all sources of benefits and costs measured in terms of a common numeraire and then discounts it at the social discount rate. If domestic consumption is the numeraire, an appropriate discount rate would be the after-tax rate of return on savings, at least assuming that capital markets worked reasonably well. Capital market distortions would be taken account of as a form of cost entering into the numerator of the project's present value calculation. An alternative procedure for taking account of the cost of financing when capital markets are distorted has been championed by Harberger (1969). He advocates using a weighted-average discount rate for discounting project benefits and costs: the pre-tax and after-tax rates of return to capital are weighted by the shares of project financing coming from forgone investment and increased savings, respectively. Though this method has simplicity as its virtue, its results will typically not correspond with the procedure outlined above, so it will not accurately reflect the true NPV of the project. More details of the circumstances under which the two methods differ may be found in Feldstein (1972a) and Boadway and Bruce (1984).

VI. RISK AND UNCERTAINTY

One of the most vexing problems in project evaluation concerns the treatment of uncertainty. The problem is that the stream of future benefits and costs is not known with certainty; at best, the evaluator may know the probability distribution of benefits and costs. Perhaps more to the point, the households who are affected by the project will not know with certainty the stream of costs and benefits. Given our assumption of welfarism, it is from their perspective that the benefits and costs must be evaluated. Project evaluation practice typically proceeds on the assumption that the distribution of future benefits and costs is, in fact, known to the affected households. This undoubtedly an heroic assumption, but one whose consequences we first consider, since it is the standard approach.

Consider the simple case of a household whose real income next period is y_i with probability p_i, where $\Sigma p_i = 1$. The various outcomes represent different exhaustive and mutually exclusive 'states of the world' that are possible. Under certain reasonable assumptions, the household can be supposed to rank alternative combinations of y_i according to expected utility, defined as $Eu(y) = \Sigma p_i u(y_i)$. The fact that the values of y_i differ implies that the household faces some risk. Assuming the household to be risk-averse ($u''<0$), the household would be willing to pay something to avoid this risk. This willingness-to-pay, referred to as the cost of risk and denoted k, is defined implicitly as follows:

$$u(\hat{y}-k)=\sum_i p_i u(y_i)$$

where $\hat{y}=\sum_i p_i y_i$ is expected real income. The value of k, for a given distribution of y_i values, evidently dependent on how risk-averse the household is. To see this, note that to a first-order approximation, the cost of risk can be shown to simplify to:[13]

[13] To see this, we can use a Taylor series approximation to obtain $u(y_i) = u(\hat{y}) + u'(\hat{y})(y_i-\hat{y}) + \frac{1}{2}u''(\hat{y})(y_i-\hat{y})^2 + R$, where R is the sum of the higher order terms. Similarly, given that k is relatively small, a first-order approximation of u(_ -k) can be obtained as $u(\hat{y}-k) \cong u(\hat{y}) - ku(\hat{y}-k)$. Combining these in the above definition of k yields:

$u(\hat{y}) - ku(\hat{y}-k) \cong \Sigma p_i [u(\hat{y}) + u'(\hat{y})(y_i-\hat{y}) + \frac{1}{2}u''(\hat{y})(y_i-\hat{y})^2]$

which can be rearranged to obtain the expression for k in the text.

$$k \cong -\frac{1}{2}\left(\frac{u''(\hat{y})}{u'(\hat{y})}\right)\cdot \operatorname{var}(y_i)$$

This indicates that the cost of risk depends both on the dispersion of incomes, as indicated by the variance, and the degree of risk aversion, measured here as the co-called coefficient of absolute risk aversion, $-u''(\hat{y})/u'(\hat{y})$.

As with other project benefits and costs, there may also be indirect costs of risk associated with a project. For example, recall the social opportunity cost of project financing, which included the opportunity cost associated with crowded-out investment. Suppose that the investment crowded out had a risky return, so that its pre-tax rate of return ρ included a risk premium of , say, β. If the investment is forgone, this risk is no longer borne. Thus, the above SOC formula would have to be amended to reflect that. For example, in the simple case in which all project proceeds are consumed, the SOC would now be given by:

$$SOC = \frac{\Delta S}{\Delta B} + \frac{(\Delta I/\Delta B)(\rho - \beta)}{r}$$

The cost of risk identifies the willingness-to-pay of a household faced with a known distribution of outcomes. It is what, in principle, would have to be included as a cost in a project evaluation which otherwise uses discounted expected benefits and costs to calculate the NPV.[14] The difficulty, even assuming that the distribution of future costs and benefits is known, is that the cost of risk is not readily observable since it does not correspond with any market price. One is left with making arbitrary adjustments to account for the riskiness of benefits and costs.

There may be a way out in certain circumstances. Even though a project's benefits and costs may be risky, those risks may be diluted by risk pooling or risk spreading. Risk pooling occurs if the project's benefits and costs be diversified in the

[14] As in the case of capital market distortions, some persons advocate taking account of risk by incorporating it into the discount rate. There may be special cases in which the use of a risk-adjusted discount rate to discount expected benefits and costs is equivalent to treating the cost of risk as a cost and discounting at a risk-free discount rate which represents how consumers actually discount future versus present consumption, But, in general, the two procedures will not be equivalent.

‘portfolios’ of households. If the project returns in question are independently distributed compared with other sources of income accruing to the household, the variance of the entire portfolio will be less than the sum of the variances of individual elements. The more independent income streams there are, the smaller will be the portfolio variance, with the latter approaching zero as the number of assets increases. The extent to which this may be applicable in the case of a project’s returns depends on the case in question. But to the extent that the riskiness facing the household is made negligible by risk pooling, the cost of risk on a given project can be ignored.

Alternatively, public projects may be subject not to risk pooling by the representative household, but risk-spreading within the public sector itself. Instead of pooling risks among a number of independent projects, a project’s risk might be diminished by sharing the risk among a large number of households, something which may well occur for the typical public project. Returning to our analysis of the cost of risk-taking above, consider an asset with a random return of y, which is shared equally among n households. The variance of the return faced by each household, var (y/n) can be shown simply to be equal to var $(y)/n^2$. Therefore, the cost of risk k can then be written:

$$k \cong -\frac{1}{2}\left(\frac{u''(\hat{y})}{u'(\hat{y})}\right)\frac{\mathrm{var}(y)}{n^2}$$

As n increases, the risk per person obviously decreases rapidly. Moreover, total risk from the project, nk, also diminishes, approaching zero as n rises. This result, due to Arrow and Lind (1970) suggests that when the risk is spread among a large number of taxpayers, the cost of risk can be ignored. Again, whether the conditions are right for this to be the case depends upon the circumstances. The project must be small relative to total national output, and independent of it. And, as the number of persons increases, the return per person must fall, a condition that would not be satisfied for public goods, for example.

As the discussion in this section indicates, costing risk in project evaluation is a difficult proposition. The risks are not readily observable and their costs cannot be inferred from market information. Moreover, the analysis is premised on the notion that households have a good idea of the probability distribution of outcomes, something which may well not apply in practice. Given this, the analyst is often left with taking account of risk in an *ad hoc* manner. One common practice is to report optimistic (upper-bound) estimates, pessimistic (lower-bound) estimates, and best-guess estimates, and let the policy maker weigh the alternatives according to some notion of how likely each scenario might be.

VII. DISTRIBUTIVE WEIGHTS

The principles of project evaluation have been presented with reference to measuring benefits and costs according to some numeraire (e.g., rupees worth of consumption), where values are imputed according to those apply to the actual households in the economy. Little explicit account has been taken of the fact that a unit of numeraire might be 'worth more' to one household than to the next. If so, summing up rupees values of consumption would be akin to summing up numbers of apples and oranges. We have discussed the rationale for this procedure of treating a rupee as a rupee no matter whose hands it is in. But the procedure remains controversial, if only because this procedure itself implicitly assigns welfare weights to households. There are suggested ways of getting around that, such as by invoking hypothetical compensation arguments, but they are not entirely compelling to all persons (see, especially, the summary statement by Blackorby and Donaldson, 1990).

Concern with the normative implications of simply aggregating rupee values of benefits and costs has led some observers to advocate incorporting distributive weights into the calculation to take account of the fact that one rupee is worth more to one household than another from a social point of view (e.g., Drèze and Stern, 1987; Ahmad and Stern, 1991). At the outset a vexing question arises: what distributive weights should be used? There is obviously no single correct answer to this question, since a value judgment is involved. Thus, the practice has been to report results for a range of judgments, and let the policy-maker decide. For example, a common procedure is to adopt a very simply form of social weighting procedure by parametrizing the social welfare function into one involving a single parameter — the coefficient of aversion to inequality. We saw earlier that a social welfare function of the form

$$W(Y) = \sum \frac{(Y_i)^{1-\rho}}{(1-\rho)}$$

has only a single parameter ρ. Varying ρ from zero to infinity spans the range of coefficients of aversion of inequality for all inequality-averse social welfare functions. In principle, benefits and costs accruing to various households could be weighted by their marginal social utility β_i, which is given by

$$\beta_i = \frac{\partial W}{\partial Y_i} = Y_i^{-\rho}$$

(with $\rho = 0$ corresponding to the case where no welfare weights are used). The problem with this procedure is that it is very difficult to attribute benefits and costs to households according to their real income levels. For some categories, it might be possible. Thus, wage payments for workers might be so weighted, as has been the practice in various applications of the Little-Mirrlees methodology. But generally that will not be possible.

An alternative, less demanding procedure, is to attribute welfare weights to commodities according to what might be known about the mix of persons that consume them. To see how this works, consider the change in social welfare from a given project:

$$dW = \sum \beta_i \cdot dY_i$$

where dY_i is the change in real income accruing to household i as a result of the project. Suppose the project involves changes in prices facing households. Then, using the principles of consumer theory, the change in real income for household i can be written $dY_i = -\Sigma_j \, x_j^{\,i} dp_j$. Substituting this into the expression for dW, one obtains:

$$dW = \sum_j R_j \cdot X_j \cdot dp_j$$

where X_j is aggregate consumption of good j, and $R_j = \Sigma_i \beta_i x_j^i / X_j$ is referred to by Feldstein (1972b) as the distribution characteristic of good j. As can be seen, it is a weighted average of the β_i terms, where the weights are the proportions of good i consumed by household j. Costs and benefits can be weighted by their distributive characteristics, and the distributive characteristics can themselves be constructed using various assumptions about the marginal social incomes, β_i, or equivalently the coefficient of aversion to inequality, ρ. Obviously this is much less demanding than disaggregating each benefit and cost to individual households, though still not a trivial exercise.

Figure 6: Market for Commodity X

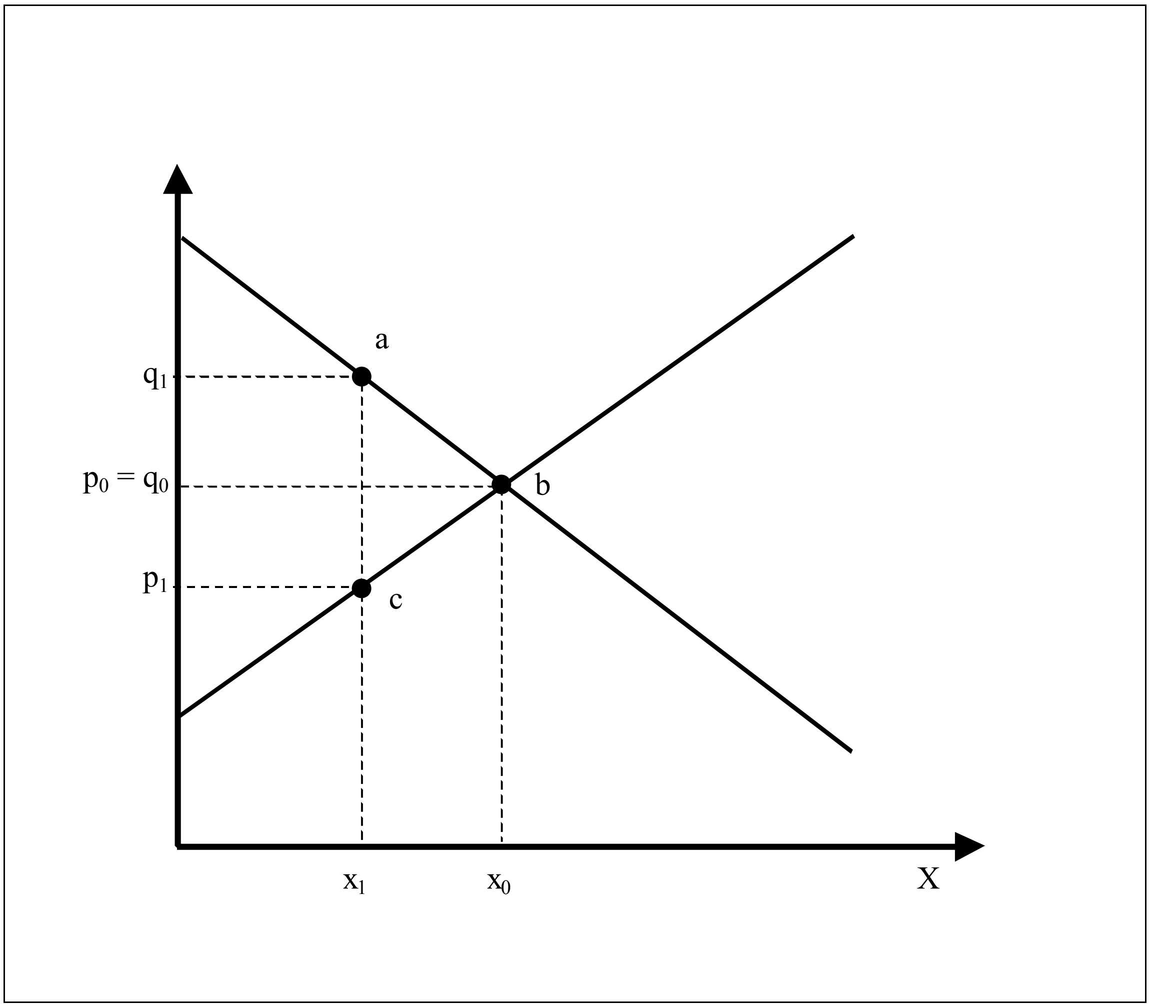

An example of the use of distributive characteristics of this sort was presented by Harberger (1978). It is simply an illustrative case involving the measurement of the welfare cost of an excise tax. Figure 6 (above) depicts the market for commodity X, with demand curve labelled D and supply curve labelled S. An excise tax drives a wedge between supply and demand prices and results in an equilibrium price and output of q_1 and X_1, compared with q_0 and X_0 in the absence of the tax. Producer prices in the two equilibria are p_0 and p_1. Consumers lose q_1abq_0 from the tax, while producers lose q_0bcp_1. The government obtains tax revenues of q_1acp_1. If distributive weights are attached to these gains and losses according to how various income groups share them, the change in welfare will be:

$$\Delta W = -R_D\left(q_1 \mathrm{ab} q_0\right) - R_S\left(q_o \mathrm{bc} p_1\right) + R_G\left(q_1 \mathrm{ac} p_1\right)$$

where R_D, R_S and R_G are the distributive characteristics associated with consumers surplus, producers surplus and government revenue respectively. Of course, if there were no aversion to inequality, these distributive characteristics would all equal unity, so the welfare change would be simply the standard loss in surplus, -abc.

VII. 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.

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

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

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* **IX**, 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.

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.

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

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).