APPLIED WELFARE ECONOMICS IN PRACTICE

(Notes on a Framework of "Principles and Standards" for Cost-Benefit Analysis)

Arnold C. Harberger

University of California, Los Angeles

September 2010

This paper had its origin in a keynote address that I presented at the Vancouver meetings of the Western Economic Association supplemented by my final remarks at the meeting of the Society for Benefit-Cost Analysis at which my term as president of that organization was coming to a close.

The paper consists of a series of points that should be borne in mind by anyone who attempts to set up a framework of "principles and standards" for cost-benefit analysis. Readers should quickly appreciate the need for such a framework, in order to avoid falling into the trap of grasping for one ad hoc solution after another, as one deals with different problems. This almost inevitably leads to inconsistencies and contradictions as one moves from one problem with one ad hoc solution to another whose ad hoc answer is inconsistent with the first one.

In a sense, this paper can be thought of as the eighth in the series that started with five papers that formed the basis of a short course for AID professionals, given in December, 2008, and that was continued with two recent papers dealing with the cost-benefit analysis of electricity projects. The present paper follows in that sequence, but it is certainly not pitched for neophytes in the field of cost-benefit analysis. It has more meaning for those who are helping to set up a

system of implementing cost-benefit analysis in a developing country or in one of its ministries or agencies.

The paper consists of a series of distinct points, each of which highlights some element or elements that definitely should form part of a framework of principles and standards, or some other element or approach that definitely should be avoided as one tries to set up such a framework.

1. <u>The Classical Dichotomy</u>: The traditional division of the study of economics into a "real side" and a "monetary-macro" side is an extremely powerful tool. Life is complicated enough when we are considering equilibria determining real quantities and relative prices. It gets much more complicated when we insert interactions between those and the absolute level of prices and/or the level of unemployment.

We should not neglect these latter interactions, but they deserve a completely separate chapter in our "book". Most of the chapters should rigorously stick to the world of real quantities and relative prices. This has been the standard practice in cost-benefit analysis; it means that the resources used in our project came from somewhere -- i.e., have an alternative use, an opportunity cost. In the world of investment projects even the investment itself will typically only take place quite some time after the analysis is made, and the project's life may extend for 25 to 50 years after that. It makes no sense to mingle cyclical considerations into our picture of how the economy will evolve, "with" and "without" the project, over this long future period.

2. <u>Choice of a Numeraire</u>: In a world of relative prices, we can rightly ask, relative to what? You can imagine what a mess it would be if we chose the price of oil as our numeraire. Remember that oil was only $10 a barrel as late as 1998. GDP measured in barrels of oil would

have been at a peak then, and at a deep, deep trough when oil got to $100 a barrel. All relative prices would likewise undergo huge gyrations if the price of oil were the numeraire. Everybody who has thought seriously about this problem has come to the same conclusion. For practical work on real problems, the only two numeraires that make sense are the GDP deflator and the consumer price index (CPI). With the first, we measure everything in producer baskets; with the second, in consumer baskets. The need for a numeraire and for expressing benefits and costs in real terms should be highlighted in any attempt to set up a framework for C-B analysis. We can guess by how much real wages will grow over the next 10 or 20 years, but our science does not give us anything like an equally good basis for predicting nominal wages.

3. The Numeraire and the Discount Rate: Remember, if we are measuring our costs and benefits in consumer or producer baskets, these are the units into which every benefit and every cost is converted. The net benefit flows that make up our project's profile are all expressed in numeraire units. The discount rate is how we transit from, say, consumer baskets of period t to consumer baskets of period 0, or vice versa.

When we have an investment project whose investment costs are, say, 1000 consumer baskets, these must have come either from forgone investments, or from forgone consumption (increased saving), or from new foreign borrowing. Each of these has an intertemporal opportunity cost. Say 700 came from forgone investment which would have yielded 12%, 100 from increased saving whose net-of-tax supply price was 4%, and 200 from net new foreign borrowing at a real marginal cost of 8%. (All these are real rates.) These costs combine to $(.7 \times 12\%) + (.1 \times 4\%) + (.2 \times 6\%) = 10\%$.

The point here is that the principles of C-B analysis themselves tell us that this project is not worthwhile unless it yields returns that cover the 10% that the economy has given up in the

act of raising the funds. The key point here is that the resources used by a project, whether in the investment phase or as ongoing operating costs, have to come from somewhere, I have used the capital market as the "sponge" which releases resources when they are needed and absorbs surpluses as they appear. When it comes to government projects, there is no question that day by day, week by week, month by month and year by year, the capital market is indeed the marginal source and use of funds.

4. **Marginal Utility and the Numeraire**: Everybody has demand and/or supply prices for differing amounts of any good or service. These are connected to each person's marginal utility of such goods and services. We cannot measure utility, but people arrive at competitive demand and supply prices by converting utility into real monetary (numeraire) units each using his or her own marginal utility of "money". This fact has been at the root of consumer/producer surplus analysis, and of Slutzky-Hotelling-Hicks general equilibrium applied welfare economics, from day one to day N. One should neither ignore nor try to supplant this whole long history of our discipline.

5. Distributional Weights, etc.: As I have said and written any number of times, the idea of distributional weights (declining as income or wealth rises) is very attractive at first glance to just about everybody, but it is pretty much a dead end as far as cost-benefit analysis is concerned. I will not argue with weights that vary between 1.1 and 0.9. But I do argue with the exponential weighting functions that appear in many places in the optimal tax literature and in other economic writing. The problem is that if we have some people with weights of 2 and others with weights of 1/2, we must (to be consistent) approve a "project" that takes 100 from those with weights of 1/2 (entailing a weighted cost of 50), and gives only 26 to those with weights of 2 (yielding a weighted benefit of 52) -- in spite of the fact that this operation involves

an economic efficiency cost (or waste) of 74. If somebody finds a way of bringing about the same transfer more cheaply -- say wasting only half the amount taken, then we should press the redistribution further and take 100 from those with a weight of 2/3 and give 50 to those with a weight of 4/3. If we can find still less costly mechanisms for effectuating transfers, say with around 20% waste, we'd end up accepting the idea of taking 100 from those with weights of .9 and giving 80 those with weights of 1.15.

These examples should be interpreted in light of the fact that the literature is replete with elasticities of marginal utility with respect to income (wealth) ranging around -1 (often from -1/2 to -2). At an elasticity of -1, the first case (entailing an efficiency cost equaling up to 3/4 of the donors' loss) would occur with the "donors" having an income only four times that of the recipients. With an elasticity of -2, the first case would come with donors having only twice the income of the recipients. With an elasticity of -1/2, the donor/recipient income multiple would be 8.

I've presented similar examples to all kinds of audiences for more than 30 years, and have yet to encounter any serious supporters of the rigorous acceptance of all the implications of distributional weights with elasticities in the -1/2 to -2 range.

The point is that one cannot apply distributional weights where one finds their implications quite acceptable (as they happen to be in the optimal income tax literature), and then simply ignore a huge range of cases where their implications are unacceptable.

6. <u>Basic Needs Externalities</u>: If one is looking for a systematic way to take what are called distributional considerations into account, I have no doubt that "basic needs externalities" provide a much sounder (and more widely acceptable) footing than distributional weights (in the ranges indicated). Basic needs externalities say that "society" (e.g., the taxpaying public) is

willing to pay something extra to reduce the "education deficit", the "health care deficit", the "nutrition deficit" and possibly the "housing deficit" of those who are particularly disadvantaged in these terms.

Basic needs externalities take their motivation from the utility of the donors rather than that of the recipients. This paternalistic vision is supported by the fact that transfers in kind (free education, free or subsidized medical care, food subsidies and food stamps, subsidized public housing) are far more common around the world than transfers in cash. Yet basic needs externalities do not preclude transfers in cash -- to the extent that higher cash incomes cause recipients to spend more on basic necessities, the externalities associated with this incremental consumption of such necessities count as added benefits for the "project".

7. <u>A Shadow Price of Government Funds</u>: This stems from the need to deal with a "loose end" for cost-benefit analysis based on the capital market as the marginal source and use of funds. What about projects (like non-toll highways) whose costs are cash outlays for the public sector but whose benefits are received by the public "in kind", and are not reflected in cash receipts to the government. We "presume" them to be financed by borrowing, which, however, is never paid off out of project proceeds. The natural conclusion is that sooner or later it will have to be covered by taxes. But taxes have excess burdens, so, by the standard rules of applied welfare economics, a marginal dollar raised by taxes has an economic cost of more than a dollar, say (!+λ). When we assign the premium(or perhaps better put, the penalty) of λ to the present value of all accumulated uncovered costs at the end of a project's life -- such assigning turns out to be mathematically equivalent to applying the factor $(1+\lambda)$ to all cash inflows to the government and to all cash outflows from the government, year by year throughout the project's

life -- including its investment phase as well as its operating phase. This $(1+\lambda)$ is therefore the shadow price of government funds, SPGF.

Once the logic behind SPGF is recognized, it is hard to avoid making it an absolutely standard part of cost-benefit analysis. The problem is that there is no standard way of raising an extra dollar via taxation, and each separate way would almost surely have a different λ. This is just one example of a general problem for CBA -- of handling cases where quantification is called for, but hard to accomplish, or subject to a wide range of uncertainty.

My own answer, in the case of SPGF is to set λ equal to something in the range of .15 to .25. I am confident that these are conservative estimates of the marginal excess burden per marginal dollar of tax revenue.[1]

8. <u>On Compensated and Uncompensated Elasticities</u>: This is either an easy problem or a very complicated one -- but probably not very important. Taking the easy side first, there are only two kinds of price changes that really contain an income effect for the economy as a whole -- 1) a real cost reduction in the production of a good or service, where the benefit is the full saving of resource costs on the "old" quantity plus a triangle of extra consumer and/or producer surplus on the change in quantity, and ii) a reduction in the world price of an

---

[1] Take the case of a simple excise tax T on a good in infinitely elastic supply. The increment of revenue stemming from a rise in T is $QdT + T(\partial Q/\partial T)dT$. The increment of efficiency cost is $-T(\partial Q/\partial T)dT$. The ratio of Δ efficiency cost to Δ revenue is therefore $-\eta\tau/(1+\eta)$, where $\eta(<0)$ is the elasticity of demand and $\tau(=T/P)$ is the ad valorem equivalent tax rate. For $\eta = -1$, this is equal to $\tau/(1-\tau)$, which is .25 for $\tau = .2$. For $\eta = -1/2$ this is $.5\tau/(1-.5\tau)$, which is .177 for $\tau = .3$. Empirical estimates of λ range very widely, varying with the particular tax or taxes that are assumed to change, and of course with the country and the period involved. I hold no special brief for the range .15 - .25; only that it strikes me as a conservative range for λ. What seems to me to be completely wrong is to ignore the problem altogether and thus implicitly place a value of zero on λ. Note that the calculation of λ should ideally be done in the context of a computable general equilibrium model in which the consequences of raising given tax rates are estimated, given a whole structure of existing tax rates.

importable good, where the benefit is the price reduction times the old amount of imports plus a triangle of extra surplus on the increment of imports. Price changes due to government policies or to changes in the structure of demand do not involve an aggregate income effect. (Consumer surplus loss is measured under a compensated demand curve.) In the case of taxes, the rectangle of tax receipts of the government is offset by a first-order loss by consumers as they pay higher prices and/or by producers as they receive lower prices. In the case of shifting demand patterns (say, due to an increase in GDP), a rise in price means extra costs for demanders and extra benefits for suppliers (and vice versa for a fall in price).

Mostly, changes in world prices are there "with" or "without" our project. Real cost reductions are also in this category. So the analysis of nearly all projects or policy changes will not entail an income effect as one compares the flows of benefits and costs "with" versus "without" the project.

Just for the sake of the information, however, only in the rarest of cases is there much difference between compensated and uncompensated elasticities. Letting η^c_{jj} be the compensated elasticity of demand for j and η^u_{jj} be the uncompensated elasticity, a_j be the fraction of the budget spent on j and σ_j be the income elasticity of demand for j, the formula is

$$\eta^u_{jj} = \eta^c_{jj} - a_j \sigma_j$$

$a_j\sigma_j$ in turn can be expressed as simply m_j, the marginal propensity to spend on good j. Thus

$$\eta^u_{jj} = \eta^c_{jj} - m_j.$$

Now, for what goods is the marginal propensity to spend very large? Housing is the only big one, maybe as high as .2 or .3. For nearly every other good or service the fraction of the marginal dollar spent on it is almost certainly less than .1.[2]

As another point of information, the link between compensated and uncompensated cross-elasticities is

$$\eta_{ij}^{u} = \eta_{ij}^{c} - a_j \sigma_i,$$

where η_{ij} measures the reaction of the quantity demand of good i to a change in the price of good j.

All the above is the easy part. When it comes to analyzing the costs and benefits of specific subgroups of the population it is clear that suppliers gain from a rise in price, while the demanders lose. Each perceives an income effect and presumably responds to it. But if their marginal propensities to spend on good j are similar, these effects may cancel out. The complexities here are so severe that it is fruitless to pursue this topic for the general case. Dealing with specific cases can often make the problem more manageable.

But please note -- when we calculate the saving in road user costs due to a road improvement, we get a rectangle of benefit to road users at V_o (traffic volume without the improvement) plus a triangle of benefit due to the project-induced increase in traffic volume

---

[2] The other place where the difference between compensated and uncompensated elasticities is important is in the supply of labor. Using ε as the symbol for labor supply elasticity we have

$$\varepsilon_u = \varepsilon_c - m_{\square}$$

where $m_{\square}$ is the marginal propensity to "spend" on leisure, i.e., the fraction of an autonomous rise in income (say, from an annuity) that is taken out in the form of leisure. Most estimates place ε_c in the range of 0.1, while $m_{\square}$ can be as high as 0.2 or 0.3. Thus we have observed secular declines in working hours as real wages have risen.

$(V_1\text{-}V_0)$, there is only a trivial difference between using a compensated or uncompensated elasticity to calculate V_1. The measure of net gain to road users $(C_0\text{-}C_1)V_o +$

$\frac{1}{2}(C_0 - C_1)(V_1 - V_0)$ would differ at most in the third decimal point of V_1 as between using

η_c or η_u to calculate V_1. [The marginal propensity to "spend" on the use of any typical road improvement has to be way below one percent].

9. <u>On Discount Rates Based on Utility Functions</u>: There is a pattern in the literature of assuming a utility function with a built-in discount rate. This is by no means standard economics. It is more like assuming that all production functions are Cobb-Douglas. Not a stupid assumption for classroom or in textbook expositions, but definitely so if it is thought of as a sensible approximation to real-world production functions in general. I see no reason at all to pursue this line as one sets up guidelines for real-world cost-benefit analysis.

The right line to take is to recognize that people respond to market incentives, here as elsewhere. Consumers have different marginal rates of time preference, depending on the source or destination of increments of saving or dissaving. With a 6% interest rate, persons with marginal tax rates of 50% have a marginal rate of time preference of 3%, while those with a 20% marginal tax rate have a marginal rate of time preference of 4.8%. If the relevant gross-of-tax interest rate falls to 4%, these marginal time preference rates drop to 2% and 3.2%, respectively. In the real world of the contemporary USA, lots of people have marginal rates of time preference in the 15%-20% range, with their credit cards determining the operating margin. I believe this approach to time preference is practically required by our choice of numeraire. (The interest rates in question should, of course, be real interest rates.)

Note that in calculating the economic opportunity cost of capital, marginal rates of time preference only enter the picture as the necessary compensation for the fraction of incremental funds sourced from newly-induced saving (a weight of 0.1 in the numerical example of point 3 above.)

10. Issues of Risk: These issues bring up a mare's nest of problems, which cry for a set of straightforward and hopefully realistic assumptions. The problem is how to deal with the many different rates of return that exist in any real-world economy.

The first point to be recognized is that we are not talking about default risk. Default risk simply recognizes that the expected payout on a loan is something less than the promised amount. Economic risk is concerned with variations (and covariations) around expected values, not around promised values.

So the question arises, when we see risk-free indexed bonds that yield 3%, and equities with an expected real yield of, say, 10%, how should we treat the intervening 7 percentage points? Should we say that those investors receiving 10% are really only getting a net return of 3%; with the remaining 7% being simply a compensation for their perceived risk? My answer would be -- just try to take away their 10% and offer them an indexed bond yielding 3% real, and see what is their response. I believe that they would not be happy with that trade.

Where would I like to go on this problem? My preference is to stick with expected values and consider that the expected return of 10% is all "benefit" from the standpoint of "society".

The simplifying assumption is that "society" is neutral with respect to risk. This assumption carries huge benefits for those engaged in the analysis of project and programs,

because it allows us to focus on expected values -- i.e., to build time profiles of net benefits and net costs, where the entry for each future year is the expected net benefit or net cost for that year.

I contend that we are very far from building our profiles in this way. I have worked on several dam-building projects, and on the last three of them I actually asked whether I was authorized to assess the (low) probability of the dam breaking and the (high) cost that society would have to bear in that event. The answer I got was very clear -- you can do it if you want to, but don't ask us to pay for the time you spend on that, and don't let it delay your submission of your evaluation of the project!! Yet leaving out the probability of a dam's breaking is like ignoring the fat tail of a probability distribution. One is bound to overestimate the net present value of the dam project. In other cases we have intermittent spikes in the trajectory of real prices of commodities like oil, copper, sugar, etc. Ignoring these spikes leads one to project something like a median or modal real price, rather than a proper expected value.

Finally, dealing with probable future variations in real prices and quantities, we should recognize that there are complex interactions among them. Most Monte Carlo exercises ignore such interactions and assume independence among the different probability distributions (of prices, quantities, etc.) that they build into their analysis. Much hard work is needed in order to make our Monte Carlo exercises richer, more realistic, more reflective of genuine economic interactions. Is an expected higher real price of copper due to surging demand or lagging supplies? Are we building higher real wages into our projections of a high-growth alternative, and lower real wages for a low growth alternative, as we obviously should do?

Once we have worked these adjustments into our Monte Carlo exercises, then we can take the resulting frequency distribution of calculated net present values, and calculate its mean. This is the "proper" net present value of the project.

I hope it is evident that there is lots of work to be done simply to implement the assumption of risk neutrality. If our generation can do this, for the cases that call for its being done, we will have made a big contribution to the evolution of serious cost-benefit analysis.

11. <u>On the Costs and Benefits of Cost-Benefit Analysis</u>: I have not seen any serious discussion of this issue. We have a methodology of CBA that is soundly based on a long tradition of applied welfare economics, but, guess what?, a really good cost-benefit analysis can cost quite a bundle.

Obviously, it makes no sense to spend $500,000 to analyze a road improvement that will cost $500,000. On the other hand, it might well be worth $2 billion to get a good analysis of a project costing $50 billion.

But even here, it may not pay to do very refined work. If a pretty rough analysis finds an internal real rate of return of 30%, there may be little to gain from a refined analysis if the proper discount rate is only 10%.

Somehow, we have to face up to the facts that: a) we need quick-and-dirty methods for dealing with small projects and b) we need procedures analogous to sequential sampling, where one stops further sampling when, with the existing sample size, one can answer the question at hand (accept or reject the project) with an acceptable level of confidence.

We have to learn how to husband the resources we have available for serious cost-benefit analysis, and use them in the most effective way.

12. <u>Dealing With Future Trends of Relative Prices</u>: It is utterly unprofessional to assume that "this year's" set of relative prices will continue to prevail for the life of a project being analyzed. We use our numeraire to express prices in real terms, but these real prices clearly change over time, and we have absolutely no business holding them constant through the

whole time span of a project, as we calculate its cost-benefit profile. It is true that future real prices are hard to predict, but we should give it our best shot. It is easiest to recognize prices that are low or high for some transitory reason, like a crop failure or a world recession. In these cases we have to build into our analysis a projection of the path by which the relevant price will move back to its normal trend. It is more difficult to project future secular movements in prices, but the rule is simple. The real price of a good will have an upward trend if real cost reduction (TFP improvement) turns out to be below average (as with haircuts or taxi rides). It will have a downward trend if real cost reduction ends up above average (as in computers, cell phones and television sets). These productivity differences are easy to see (and even measure) with hindsight, but most of the time our job in project evaluation is to assess what price trend seems to be most likely over the project's expected future life, and build it into our analysis. We can somewhat hedge these bets with Monte Carlo exercises and sensitivity tests.

Probably the most important area where a systematic trend of real prices should be built into our analysis is the labor market, where the real wages of most occupations will typically have a significant upward trend, which should be built into our analyses.

Discount rates can also vary with time. It is reasonable, at a recessionary time like the present, to incorporate a lower discount rate for the present year and the immediate future. Modifying the example of section 3, let the expected rate of yield on displaced investments be 8%, the supply price of newly stimulated savings to be 2%, and the marginal cost of foreign funds be 5%. Then, using the same weights of .7, .1, and .2, we have (.7×8%) + (.1×2%) + (.2×5%) = 6.8% rather than 10%, as the economic opportunity cost of capital for the year or period in question.

Discount rates varying by year present no conceptual difficulties. If r_1 is the rate linking year zero with year 1, r_2 the rate linking year 1 with year 2, etc., a flow F_2 in year 2 would be discounted to year zero by dividing it by $(1+r_1)(1-r_2)$. In general F_N would be discounted to year zero by dividing it by the product of annual discount factors 1 through N, i.e., by $(1+r_1)(1+r_2) \ldots (1+r_{N-1})(1+r_N)$.

13. <u>Modesty and Professional Pride Go Together</u>: The methodology of cost-benefit analysis applies quite easily and naturally to commercial-type ventures (whether private-sector, public-sector or joint between the two) whose costs and benefits consist overwhelmingly of cash outlays and cash receipts. But what do we do with benefits and costs that are not in this form? The answer here is a bit complicated. Some non-cash benefits and costs can be quantified by direct application of economic analysis. Thus we have economic studies that estimate the value that commuters place on the time they spend going to and from work, and the costs involved in ships waiting in line to enter a port or canal, and the value that recreational users place on their visits to parks, museums, etc. Then we have other benefits which can be set by the analysts themselves (in the absence of other instructions) or by public sector authorities attempting to put values on particular non-market goals. In this vein we have values denoting "society's" willingness to pay for added fulfillment of the basic needs of the poor, or for added economic activities in a distressed region or industry. Finally, we have a range of areas in which neither of the previous answers can plausibly apply. National defense benefits and those linked to a nation's culture, history and traditions come to mind here. For these the standard answer of professional economists is that we have little or no claim of professional expertise in setting values on such elements. Instead, we try to quantify those items which we are professionally

equipped to estimate, and derive measures of costs and benefits for just those items. We then confront our audiences with statements like "In this project, the strictly economic costs exceed its strictly economic benefits by $200 million. We leave it to the ultimate decisionmakers to determine whether its national defense or other benefits are worth this cost."

Modesty enters here in the sense that when we realize that we are not experts at quantifying many kinds of non-market benefits, it is best for us to acknowledge this fact openly. Professionalism enters in the sense that only by "sticking to our professional last" can we fairly claim to have a special voice concerning the analysis of costs and benefits.

14. <u>A Final Caution on Discount Rates</u>: There is a very old literature, associated particularly with the name of Steve Marglin, but joined by a number of others, which focuses on using the marginal rate of time preference as the discount rate. I will illustrate that literature using the numerical example of section 3, where the weighted average opportunity cost of capital was 10%, and the marginal rate of time preference was 4%.

Marglin, et al. would be adamant about using 4% in this case as the discount rate, but they were astute in recognizing that, in our example, displaced investment really did have a productivity of 12% and funds drawn in from abroad really did have a marginal cost to the country of 8%. To take these costs into account, they conceived of a shadow price of investible funds, representing the present value (at the 4% rate) of the flow of future costs adding up to 10% that was generated by the extraction of funds. Using their methodology but our section 3 figures, the shadow price of investible funds would be 10%/4%, or 2.5. Thus an investment of $1 million, would be "costed" in the C-B analysis at $2.5 million. Correspondingly, any newly or autonomously generated investible funds would be credited with producing a benefit that was 2.5 times their dollar amount.

If we use 10% as the discount rate, then the "shadow price of investible funds" is 1.0, so we can forget about it. But if we use any rate less than the opportunity cost of funds, I believe we are logically forced to reintroduce a shadow price of investible funds. This greatly complicates the analysis, and introduces the huge problem of explaining to enthusiasts for a long-term project, who want to use a 2% discount rate, that of course they can have their 2% rate, but only if they pay the price of "costing" their project's outlays at 5 times (= 10%/2%) their actual cash value.

This dilemma can be partly solved by asserting that the excess of the expected yield on displaced investment (12% in our case) over the marginal time preference rate (4%, or 2% just above) was not perceived as a benefit by investors. But I do not believe any convincing case can be made for such a sweeping assertion.

I see no way of even partially solving this problem as far as funds drawn from abroad are concerned. Here the marginal cost of such funds is paid to outsiders, and is all pure cost as far as the paying economy is concerned.

As I see it, there is no true escape from the economic opportunity cost of capital (derived from gross-of-tax yields on investment, net-of-tax supply prices of saving, and effective marginal costs of money drawn from abroad). We can, on the one hand, use this as our discount rate and be done with it. This is surely the simplest solution, and is the one I favor. The alternative is to use some lower rate (presumably a time preference rate). But if we do this the economic opportunity cost enters by the back door, in forming the basis for the calculation of a shadow price of investible funds. This creates a whole plethora of problems, especially of communicating our analysis and our results to the general public (e.g., "This bridge costs $10 million in cash and the present value of its benefits discounted at 4% are $18 million. But it is a

terrible project, because the true cost of the investment is not $10 million but $25 million"). How can one convince even highly-educated non-economists of the logical necessity of applying the shadow-price multiple of 2.5? How much simpler it is to use the actual cost of $10 million and compare it to the present value of benefits, which would here probably be only $8-9 million, when discounted at the opportunity cost rate (10%)!!

15. <u>Dealing With Cyclical Downturns -- Absorbing the Unemployed</u>: This area has typically been neglected in discussions of cost-benefit analysis, and for pretty good reasons. Most operative CBA exercises look to a long-term future (a project's life), and it doesn't seem sensible to build predictions of future downturns into such analyses.

However, an issue arises with respect to <u>ex post</u> analyses, where we can clearly specify recession years, and with respect to recessionary times like the present, when some projects can arise which are going to be implemented right away, or in a quite near-term future when some elements of recession are likely still to prevail.

The big question is, what sort of "extra credit" should be given to projects that tend in one way or another to ameliorate a recession?

The easiest answer concerns the externality involved as a project has the result of reducing the ranks of the unemployed. Here the saving in public outlays for unemployment compensation is a clear externality, easily quantified as what a typical unemployed worker actually gets from the government.

But one quickly runs into a much tougher problem. If one creates 100 jobs in a new project or program, how many of these jobs will be reflected in a reduction in the ranks of the unemployed? For normal cost-benefit analysis the standard answer is zero. Let's say that this answer applies when the general unemployment rate is 5%. At the other extreme some of us still

remember the great depression, when the unemployment rate reached 25% and more. Here one can probably consider that 100 new jobs might actually absorb 100 (or even 120 or 130) unemployed.

What seems totally implausible is that we would pass from zero percent absorption (at a 5% overall unemployment rate) to 100% absorption at a 6 or 7 or 8 or 10 percent overall unemployment rate. Alfred Marshall opened his Principles book with the motto natura non facit saltum (nature is continuous). That motto seems to fit well with the problem at hand -- the absorption fraction has to respond gradually and continuously, as the overall unemployment rate rises. Maybe at a 10% unemployment rate this absorption fraction is 1/3. Maybe at a 15% overall rate, the fraction is 1/2. I find these numbers, to be plausible, given the fact that unemployment is heavier in some occupations, and much lower in others and given the fact that most good jobs are typically filled in the first instance by people moving from other jobs.

Since our present recession started, I have played the role of a modern-day Diogenes, going around with my lantern asking every labor or macro economist I encounter "What does the literature have to say about how the absorption fraction varies with the overall unemployment rate, or with any other key variables." Up to now the answer has been that nobody has focused their research on the absorption fraction. This is certainly a big hole to be filled by future serious researchers. But this is not a task for those trying to set up a framework of principles and standards. But they can certainly warn of the folly of assuming that every 100 jobs created during a recession will end up absorbing 100 unemployed.

Moreover, whatever assumptions anyone may make concerning the absorption fraction at a given moment during a recession, they also have to face the problem of projecting how that fraction will move, presumably toward zero, as the economy is expected to recover. It would

certainly be audacious to say that a permanent job created today will seriously influence the rate of unemployment 3, 4, or 5 years or more from now.

16, Dealing With Recessions -- Issues of Supply Price: Where we are so far is that we have an externality equal to the project-induced reduction in unemployment compensation payments, times the number of new jobs in the project, times the absorption fraction. The benefit here all goes to the government, as a saving in compensation payments. In addition there is likely to be a benefit to the unemployed workers themselves, which presumably is the producer surplus that is generated as they move from unemployed to employed status. What is this producer surplus? Most economists would say, the excess of the wage a worker receives over his/her genuine supply price of labor for the jobs they get.

The problem here is that data on workers' "reservation wage" do not give us the right answer. The 1970 census asked questions about the reservation wage of the unemployed, and got answers that were very close to the prevailing wage of each respondent's occupation. In short, the responding workers were giving the "expected demand price" for their labor, not their true supply price.

And here we come to a genuine clash between what we get empirically and what we really are looking for. When we draw supply and demand curves, or schedules, the question we ask is what would each agent offer (as supply) or want to buy (as demand) if each given wage or price were the equilibrium wage (received by all suppliers) or price (paid by all demanders) in that market? An unemployed worker in this sense might have a genuine supply price equal to half the equilibrium wage (i.e., would be willing to work for half the currently prevailing wage, if that were indeed the going wage for his/her occupation). But that is not the answer that

workers are going to give when asked for their reservation wage under current market conditions.

So we've got another big problem in determining the amount of benefit accruing to a worker who is absorbed out of the ranks of the unemployed. It certainly is not likely to be zero, which is the conclusion to which responses about reservation wage would lead us. But the worker's benefit is also certainly not the entire wage of the new job, which would imply a genuine supply price of labor equal to zero.

I do not think that these questions can be answered as we try to set a framework of principles and standards, but we should certainly warn people of the pitfalls they face, and perhaps trace out some reasonable guidelines.

17. <u>A Final Note on the "With Project/Without Project Comparison</u>: Every serious treatment of cost-benefit analysis points out that the relevant annual net cost or benefit to be attributed to a project or program is the difference between the benefits and costs that are generated, year by year, in a world "with" as compared to a world "without" that project or program. But my impression is that lots of people pay lip service to this principle without really appreciating its meaning and its implications.

I think the easiest way to describe this comparison is to ask people to envision two moving pictures of the relevant aspects of the economy -- one describing how the economy would likely evolve year by year in the presence of the project, the other describing how that evolution would likely take place in its absence. The project profile is derived by taking the year-by-year differences of benefits and costs as derived from these two motion pictures.

One common mistake treats the spending on a project as being an increment to aggregate spending. It is essential that the "sourcing" of the resources used in the project be taken into account -- a point already mentioned.

For the present, I would like to focus on the labor market and point out the difference between good and bad answers to the question of the "sourcing" of the labor used in a project.

If asked to respond to this question, Mr. Gallup or Mr. Roper would probably suggest taking a survey of those employed by the project, asking them what was their status (previous job, out of labor force, involuntarily unemployed) just before they came to work on the project. This is not the right answer because (among other things) those who came from other jobs left vacancies there, of which ordinary supply and demand economics tells us that some will be filled, and others not.

If asked to respond to the same question, economic historians and demographers would probably say that future new jobs will be filled by the natural increase of the population, by within-country and international migration, and by changes in labor force participation. This can be true at an aggregate level but is not useful in answering the question for a given project or program.

The right answer is that each of the two moving pictures that we compare should trace the likely equilibrium of the economy from the beginning to the end of the life of the project that is being analyzed. This is a tough problem, but let me describe how it was dealt with in a basic study undertaken by Banobras, a development bank that houses Mexico's central project evaluation authority. The basic data for this study were labor market surveys in around 40 different labor market areas of the country. In each of these areas the surveys provided data on

wages, average hours, average earnings, etc. for some 100 different occupations, with separate data for men and for women.

The first task was to simplify the problem by amalgamating the 100 different occupations into meaningful and manageable groups. Thus they had common, semi-skilled, and skilled labor, ordinary clerical and skilled clerical workers, technicians, professionals, managers, and perhaps one or two more categories. The rule for putting several occupations into one category was that these occupations had to pay similar wages.

So now one had the average wage rates of, say, 10 labor categories, in 40 labor markets covering basically the whole territory of Mexico. For each category and each location, the tax payments linked to that category in that place were estimated.

Next, they had to make assumptions about the geographical sourcing of incremental labor for each of the 40 centers. In all cases it was assumed that half the labor came at the expense of other jobs in that same labor market area. Different assumptions were made about the sourcing of the other half of project labor. Under the first alternative -- called the donut -- the second half of project labor was assumed to come from the contiguous labor market areas, in proportion to their existing numbers of workers in the category (e.g., semi-skilled) in each such area. Under the second alternative -- called the nationwide -- the second half of project labor was assumed to be drawn from all 39 other labor market areas, again in proportion to their existing numbers of workers in the category.

So, if w_{ir} is the monthly wage of workers of category i in region r, one estimated the amount of T_{ir} -- equal to the monthly taxes paid, linked to the wage w_{ir}. Then the economic opportunity cost of labor of category i in region s would be

$$w_{is} - [T_{is} - \sum_{r} a_{ir} T_{ir}].$$

Here a_{ir} for region r would be the fraction of workers in category i (in this case semi-skilled workers) coming from each sourcing region. In the Mexican study, a_{ir} was always 1/2 for the region in which the project was located. The remaining half would be distributed across the other sourcing regions in proportions of their numbers (N_{ir}) of workers in the category in question. That is, a_{ir} for regions other than s would be equal to $(N_{ir} / \sum_{r \neq s} N_{ir})(1/2)$.

This may look more complicated than it really is. What we are doing is taking as a benefit the taxes to be paid by (or on account of) workers in the project, and as a cost that taxes that would have been paid by (or on account of) them in the employment they likely would have had in the absence of the project.

Where one has sourcing among major cities in a country, one typically does not expect to find significant difference in the taxes paid (per worker) as between source and destination. However when a project is in the main metropolitan area of a country, and is likely to be sourced largely by migrants from the rural hinterland, then there can be large differences in taxes paid, which will tend to lead to an economic opportunity cost of labor somewhat lower than the market wage. This is likely to be the case for low-skilled labor in many developing countries (notably China and India). There the wage paid in the urban destinations of migration tends to be very substantially higher than the wage in the rural sources of migration. This is taken as a genuine economic cost, as it reflects the fact that labor probably would not move to the cities in the absence of a significant premium (e.g., city wage perhaps double the rural wage).

The bottom line is simply that most of the time one takes the wage actually paid (or to be paid) by the project as the first approximation measure of the economic opportunity cost of labor. This is then adjusted by a difference in taxes linked to the wage (payroll taxes, personal income

taxes, and sales and excise taxes) as between employment on the project in question and the alternative employment "at the source" (conceptualized as the difference in employment at each source as between the two moving pictures -- "with" and "without" the project in question).