INTRODUCTION TO COST-BENEFIT ANALYSIS

P A R T I I I

ADDRESSING SOCIAL CONCERNS

Arnold C. Harberger

University of California, Los Angeles

December 2008

Background

From its oldest roots in the early 19th century, applied welfare economics has focused on the goal of economic efficiency. To be sure, its approach was anything but crude -- it could deal, for example, with the supply price of a person's labor being different for different jobs in the same place or for the same job in different places (see Part II). It could also handle the inefficiencies of consumption stemming from each household getting a given ration of milk (in spite of their having very different intensities of demand) with the same facility as it handled the inefficiency of awarding each farmer a given number of acres on which he could plant wheat (in spite of the productivity of the land being very different from one farm to another). But beneath it all was the treatment of a dollar of cost as counting the same, regardless of which person or group within the society ended up bearing that cost, and similarly for dollars of benefit accruing to different persons or groups.

This focus on efficiency is embodied in the three basic postulates on which applied welfare economics has been based. These can be summarized as: 1) benefits being measured at each step by demand price (= willingness to pay), 2) cost being measured at each step by supply price (= willingness to supply) and 3) the aggregation of these benefits and costs across individuals and groups, regardless of who within the society in question was enjoying the benefits or bearing the costs (= adding up).

It should come as no surprise that even from the earliest days it was the third postulate (adding up) that became the focus of controversy and discussion. Ask 100 people whether an incremental dollar will "do more good" if put in the hands of a poor person or a rich person, and all or nearly all of them will without hesitation side with the poor person. It is only a small step to move from this answer to the idea of "distributional weights" -- of weighing a dollar increment of benefit or cost differently, depending on the individual a family or group to whom it accrues.

But, as we will see later, applying distributional weights systematically, within the frame of cost-benefit analysis: 1) is extremely difficult to do, and 2) carries many implications for policy that most people are unwilling to accept. The standard way to escape from the problems posed by 1) and the dilemma posed by 2) has been to stick to the three postulates of demand price, supply price, and adding up; but at the same time to emphasize that "all we are doing is measuring economic efficiency". This is a key objective in economic policy analysis, and it is

something that we really do know how to measure. Thus we can with a clear conscience say that a given agricultural policy has an efficiency cost of $800 million, or that a given slum-clearance project has an efficiency cost of $50 million, and then leave it to the "authorities" to judge whether or not the non-efficiency benefits of that policy or that project are sufficient to outweigh the $800 million or the $50 million of efficiency costs that we measure.

This "efficiency-only" position is probably the safest one for a defender of cost-benefit analysis to take. It does not claim that we are able scientifically to measure the non-efficiency benefits or costs entailed in a shift of a benefit from one group to another, or in the shift by the country's military to a new weapons system, or in measures that, at the expense of economic efficiency, accommodate the policy demands of one special-interest group or another. We measure the efficiency costs, and let somebody else worry about the non-efficiency aspects of a policy or a project.

I believe that all cost-benefit professionals have to adopt something like an efficiency-only position at one level or another. We have no professional business in placing a dollar value on improved relations with India (which may be an important byproduct of a given project or program), or on many national defense outlays. But move in a little closer to our own terrain and you enter a sort of no-mans-land, where there are good arguments for applying efficiency standards, yet where in order to do so we have to place dollar values on a whole array of benefits or costs that are often very hard to quantify.

The value of a human life is a case in point. The average citizen's instinctive reaction is "no amount of money is sufficient to compensate for the loss of a human life", yet there is a myriad of policies, programs and projects that in effect embody a tradeoff between dollar cost and human life -- the setting of speed limits, the placing of traffic lights and stop signs, the building of median strips on highways, the straightening of dangerous curves are just a few examples related to roads. If we count the life-taking costs and the life-saving benefits of such decisions as being outside our purview, simply to be weighed by the "authorities" as non-economic benefits or costs, we find ourselves with a serious problem. For we actually <u>can</u> estimate with some accuracy how many lives per year would be saved by imposing a national speed limit of 55 miles per hour and we can also estimate the costs (mainly in travel time) that such a policy would entail. Relating those two, we have an economic cost per human life saved that is implicit in either adopting or rejecting a 55 mph speed limit. Then we can do the same for the placing of traffic lights and stop signs, for the introduction of median strips, and for the straightening of specific curves on specific roads. Doing all this we would find that very different implicit values of human life emerge from these different exercises.

Thus we might find ourselves with situations in which we are paying $10 million to save a human life by straightening a specific curve, but where we could save other lives at a cost of $1 million per life by the judicious introduction of stop signs. Obviously, we cannot juggle hundreds of such specific comparisons in our heads as we analyze a host of different policies, projects and programs. The way out is the introduction into our cost-benefit framework of a "shadow price" of a human life. If this price if $5 million per life, then the stop sign projects would generate a net benefit of $4 million per life, while the curve-straightening project would show a net cost of $5 million. Instead of all the projects being juxtaposed one to each of the others, each single project would be assigned a benefit of $5 million for every human life it was expected to save, and a cost of $5 million for every life it was likely to take. All of a sudden the saving and taking of human lives has entered the world of efficiency calculation!

The moment we feel ready to place a monetary value on a given "noneconomic" objective, that value opens the door to incorporating that objective into the efficiency-oriented calculus of cost-benefit analysis. What we can do with life years we can also do with the value of commuter time, with the valuation of free public services such as those of public parks. On the cost side we can, for example, introduce prices for the various pollutants that a project might introduce into our atmosphere or our waterways.

It is here that we rather quickly reach a crossroads. While we may be prepared to set a price on carbon emissions into the atmosphere or on nitrogen spewed into a river or lake, we may not be quite ready to do the same for a battalion to be added to the army or a submarine to be added to the navy. Our readiness to quantify some "noneconomic" benefits or costs varies with how well we think we can pin down those numbers. If we are confident that the value we seek lies within a range of 10% or 20%, or even 30%, then we can still feel like professional economists when we incorporate that range (or a central value within that range) into our analysis. But if our valuation is so uncertain that it spans a range of 300% or 500% or 1000%, then we are probably better off not trying to introduce such an item directly into our quantitative analysis, and simply passing the buck on to the "authorities". Put another way, where the range is very wide, it can end up useless.[1]

In sum, those who want to introduce greater rationality into the decision process on public expenditures have, in cost-benefit analysis, a very worthy product to sell. In every country there is a vast array of projects and programs to which known, readily available techniques can be applied. And one can be quite certain that with enough effort and ingenuity, we will be able to keep on expanding the scope for reasonable application of cost-benefit analysis. But we should beware of overextending ourselves -- there is plenty that we can do while still claiming our work to be "professional", and there are many interesting possibilities for extending the range of projects over which we can function as professionals, but there also is another range of projects for which our ability to quantify benefits and costs is too limited or too vague to be useful.

On Distributional Weights

One of the important areas in which "noneconomic" considerations are often broached concerns the distribution of income and/or wealth. The idea that a dollar in the hands of a poor person is worth more (from society's point of view) than a dollar in the hands of someone much richer -- that idea has deep roots in most people's thinking. And it also has roots in the field of economics. The notion of people's well-being being measured by their "utility" dates back at least to the early 19th century, and has a long history from that point on. "Utility" appears at three levels in our literature -- "ordinal utility", meaning individuals can state their preference (or indifference) as between any two bundles of goods and services (or any two situations); "individually measurable utility", meaning that people can rate differences between bundles (e.g., saying that the difference between A and B is bigger or smaller (in utility terms) than the difference between B and C); and "measurable and interpersonally comparable utility", which says that one person is enjoying more utility than another or that an incremental dollar is worth

---

[1]Like telling a pregnant mother that you can predict what height her new offspring will reach at age 21 -- and then stating the range to be from three to seven feet!

more (in utility terms) to one person than to another. Much of economic theory can be derived just using the notion of ordinal utility (indifference curves, etc.), but the analysis of risk typically requires one to take the next step, to individually measurable utility. Neither of those provides any basis for a distributional weights framework. For that one needs to take the third step -- to measurable and interpersonally comparable utility.

Early utilitarian thinking was based on this latter assumption, but did not pursue its detailed implications. That part came later, particularly on the subject of optimal income taxation. Here we have an extensive literature stemming from the past several decades. This literature assumes that each of us has the same utility function, translating income (or wealth) into utils (the units in which utility is measured). Higher income translates into more utils, but an extra dollar contributes less and less as income rises.

In the optimal income tax literature, the problem is posed of raising a certain amount of money through an income tax, when the weight given to incremental dollars declines with income. Typically, the examples used in this literature assume the weight is cut by a quarter, a half, or three quarters, every time income is doubled. When such an assumption is applied to income distributions similar to those we observe in the real world, the resulting "optimal" income tax structure tends to have an unexpected shape -- marginal rates of tax tend to fall as income rises. This seems counterintuitive at first, but begins to make sense when we take into account that an income tax structure is composed of a series of income brackets. Raising the marginal rate for any one bracket introduces an efficiency cost by creating a new disincentive for work on the part of people in that bracket. But it produces a distributional benefit by shifting dollars to the government, not only from that bracket but also from all higher brackets. Thus, for bracket 1 out of 5 a rise in the tax rate has one efficiency cost plus 5 distributional benefits, for bracket 2 it has one efficiency cost plus 4 distributional benefits, etc. By the time we get to bracket 5 we have one efficiency and only one distributional benefit. It is the fact that fewer distributional benefits come from raising marginal rates in the higher brackets that produces the counterintuitive result of the "optimum" marginal rate declining as income rises.[2]

What is troublesome about the optimal tax literature is that its results come from a pretty fancy set of calculations using distributional weights, and seem to argue for more moderate (i.e., less progressive) income tax structures than those we actually observe in most countries. This is taken as reassuring by many people; it thus serves to foster the general acceptance of the distributional weights approach.

---

[2]To add to the anomaly, the typical optimal pattern has average tax rates rising (up to the top bracket) at the same time as marginal tax rates fall, as income rises. This is due to the existence of an optimal exemption level. Thus an optimal marginal rate structure might be zero up to $20,000 and 30% from $20,000 to $40,000, taking a tax of 6000 (= 15%) from an income of $40,000. It could then go on to take 25% on incomes from $40,000 to $80,000, the total tax on $80,000 being $16,000 (= 20%), and (with a marginal rate of 22% in incomes over $80,000) collect $20,400 (= 20.4%) on an income of $100,000, and $42,400 (= 21.2%) on an income of $200,000. This attribute of the optimal structure being progressive in the average rate at the same time as it is regressive in the marginal rate, was sort of reassuring to readers of the optimal tax literature -- the result was not totally counterintuitive.

Unfortunately, this last step is not warranted. In order to assess the merits of any systematic approach, one has to test it throughout the relevant range. In particular, if we are to use distributional weights in the field of project evaluation we have to test that approach as it would apply to the approval of specific projects and to the problem of choosing among alternative projects.

Looking at one project alone, consider:

	Unweighted	Avg. Weight	Weighted
Present Value of Benefits	500	1.5	750
Costs	-1000	0.5	-500
Net Present Value	-500		+250

If we take distributional weights seriously, and use the indicated weights, we must recommend acceptance of this project, in spite of its net efficiency cost of 500.

Suppose one says no to this on the ground that the government could make a simple transfer of 500 to the beneficiaries, and thus get the same distributional benefit without the efficiency cost. Then, if such a transfer can be costlessly made, it clearly should be made. But costless transfers are hard to find in the real world. Suppose that extracting money from one group and transferring it to another entailed resource costs of extraction, of delivery, and of administration plus the efficiency cost of the taxes themselves, equal to one third of the amount collected. Then one would reject the project above, but would accept the following one.

	Unweighted	Avg. Weight	Weighted
Present Value of Benefits	600	1.2	720
Costs	-900	0.5	-720
Net Present Value	-300		0

One can quickly see that if the distributional weights of gainers and losers from a project have a ratio of 1.5 (= 1.2/0.8), then that project will be at the margin of acceptability when unweighted costs are 1.5 times unweighted benefits. More broadly, when the weights of gainers and losers have a ratio of $(1+\lambda)$, then projects will be acceptable so long as unweighted costs are less than or equal to $(1+\lambda)$ times unweighted benefits.

Always, there emerges a tradeoff in which at the margin society pays in efficiency costs for what it gains in distributional benefits. One doesn't run into much trouble, then, if the lowest distributional weight is 0.9 and the highest one is 1.1. But such weights would not lead to much in the way of redistributive policies or projects. It is when the weights get to be amply different that the implications of a distributional weights framework lead to policies that would be

unacceptable to most people. Those who like the idea of distributional weights, however, are typically not at all happy with weights of 1.1 at the poverty line and of 0.9, say, for millionaires. But give them a span of weights that they like, and you run into the difficulties (of in effect paying huge efficiency costs for most distributional benefits) that were outlined above.

A second significant attribute of distributional weights is that, even though it is quite common for authors to assign weights which decline as income rises, the benefit that is measured under this concept is consumer or producer surplus, not income. This is exactly as it should be; if one is trying to assess whether persons or groups feel that they are better off in situation B (say, with a project) then in situation A (say, without it). But it also leads to a situation where increments to employment are very often assigned only a very modest (or even zero) benefit. This comes straight out of standard economic theory. A rise in the wage rate, say, from $10 to $11 normally pulls into the labor force a group of people who were not willing to offer their services at $10. Their gain, from the act of going to work at $11, is at most the $1 difference between $10 and $11. Those in the group with a supply price of $10.50 gain only 50 cents per hour. Those who are precisely at the margin, with a supply price of $11, are precisely on the margin of indifference as between being out of the labor force or in it at an $11 wage.

All the above is standard economics, which has been built into cost-benefit analysis, and into applied welfare economics in general for as long as we can remember. One should not think it is a mistake, or should be replaced in a system built on "supply price, demand price, and adding up". But what we have to recognize is that the story we have just told focuses on the individuals who supply the extra labor -- on how their utility if affected when some change occurs. There is nothing in applied welfare economics that says that society's valuation of an increment to employment should be based solely on its valuation by the individuals concerned.

But that is precisely the way distributional weights have worked. They have often directly measured the change in utility of the economic agents concerned, based on an assumed utility function. Otherwise, they have focused on changes in consumer and producer surplus which are simply money measures of the corresponding changes in utility. Those who have qualms about this treatment giving so little weight to increases in employment are probably thinking in different terms -- very likely in terms that fit nicely into the concept of basic needs externalities, which are the subject of the next section.

Basic Needs Externalities

Some years ago (1984) I published an article under the title "Basic Needs versus Distributional Weights in Social Cost-Benefit Analysis". The purpose of that article was to clarify that the two concepts were really quite different, and that it was wrong, as many authors were then doing, to treat them as virtual synonymous. To clarify -- distributional weights come out of an individualistic framework, with a focus on the utility of the individual consumer or worker. The concept of basic needs externalities comes out of a paternalistic framework, focusing on the willingness of <u>other sectors of society</u> to pay for an improvement in the economic situation of some individuals or groups. As I like to put it, basic needs externalities apply not to changes in the utility of the affected groups, but in their welfare -- welfare judged not by the affected individuals but, in some form or other, by the <u>rest of society</u>.

One of the standard demonstrations in elementary economics courses is why it is better for the government to give people subsidies in cash rather than in kind. This is shown to be so by demonstrating (usually using indifference curves) which bundle of goods a consumer would choose to buy, having received a given cash subsidy. Obviously that bundle -- call it A -- is the best that that particular consumer can obtain, given the expenditure limit set by the subsidy (plus other income). So if "society" decides to give that consumer a subsidy in kind -- giving a bundle of goods B instead of money, it is only by chance that society would choose bundle A. If B represented any bundle other than A, then the individual would be worse off with a subsidy in kind than with a subsidy in the form of cash. Society can make the consumer happiest by giving a subsidy of $100 in cash. If it gives an in-kind subsidy costing $100, it could in all but one case (choosing precisely bundle A) make the individual equally happy via a cash subsidy of something less than $100.

That demonstration is perfectly correct. But note that it always focuses on the utility of the recipient. Basic needs externalities focus not on the utility of the recipient, but on the utility of the donors. I tell my classes a story of a student in his third year at Yale, writing to his father of the terrible time he is having with his studies, now even worse than in the two previous years, and how he is even more miserable than he was then. He goes on to note that tuition is around $30,000 a year at Yale, and that transportation and living expenses add up to $20,000 more. Then the student makes his final pitch, "Dad, I just heard of a wonderful island in the South Pacific, with beautiful climate, beautiful beaches protected from sharks by a reef, beautiful native girls, and a free port where liquor costs only $4 a bottle and cigarettes $3 a carton. Dad, you could send me there for less than $20,000 a year, and I would be so, so happy!!"

To conclude, I ask my students to imagine 1000 letters like that, from 1000 miserable students all across America, to each of their fathers. My bet is that, if these 1000 letters were sent, I would be surprised at the end of the year to find even two or three of those students on the island -- but I would make no claim as to how many fathers one would find!!

This tells a lot about paternalism, and I believe it has ample reflection in the behavior of governments all across the world. Of the benefits that governments give, particularly under the label of social programs, the great majority come "in kind" rather than "in cash". You don't see governments just handing out checks and telling families to send their kids to school if they choose, but to feel free to spend the money as they like. Free public education comes in kind, not in cash. Even Milton Friedman's school vouchers represent "in kind" payments because they are only valid to pay for children's schooling. In the same way, public programs of medical care provide medical services, delivered "in kind", public housing programs provide housing, or housing-specific subsidies. Nutrition programs take the form of school lunches, food stamps and the like.

Why is this? The pressure for in-kind delivery arises because people want to feel that their tax money is well-spent -- well-spent in the eyes of the taxpayers, not of the recipients. They want education money to go for education, health-care money to go for health care, housing money to go for housing and nutrition money to go for nutrition.[3]

---

[3]Sometimes these desires are frustrated by the ingenuity of the recipients. Thus, when the Indian government provided free good-quality housing to some of its low-paid employees, this in many cases did not end up with those families living with a density of 1 person per 20 square

Cost-Benefit Analysis With Basic Needs Externalities

To start this section, I want to recall some steps we have already taken. When dealing with a situation in which the generation of a saving of foreign exchange brings special merit to a project or program, we end up with the concept of the economic opportunity cost of foreign exchange. When there is special merit involved in the provision of jobs, we end up working with the economic opportunity cost of labor. When judging projects that have effects on the life span of people, we have the concept of the economic value of an added year of life. These "economic opportunity costs" or "economic values" enable us to compare the effects of many different projects, say, in terms of the foreign exchange proceeds they generate, and at the same time judge those projects in terms of their effects in extending human lives, and also at the same time judge them in terms of their employment effects. And after taking all these things into account, we can measure the economic productivity of the capital invested in those projects.

In introducing basic needs externalities into our cost-benefit calculation, we seek to do something similar to find a way of comparing what project A does for nutrition, what project B does for education, what project C does for medical care. Obviously, we want to incorporate "values" for these benefits, but how to do it?

The basic needs approach focuses, as its name implies, on valuing improvements (or the reverse) in the welfare of people in the lower socioeconomic strata of society. And, as indicated above, the welfare in question is judged according to the standards not of the recipients but of the donors (government, taxpayers, voters, society?) Recall that if society wants to follow the recipients' standards and tastes, it should always just give them money.

So what we are looking for is a metric -- a way of placing a monetary value on specific increments to welfare, in specific dimensions.

Let me start with a very practical example, based on a case I once worked on for the Philippines. The concern was with measuring the external benefit of projects that ended up enhancing the educational achievement of the children of poor families. We had census data on educational profiles, showing, of 1000 children in each income decile, how many would leave school after 4^{th} grade, 5^{th} grade, etc., all the way up to the university level. We also had data on

---

meters as the government intended. Instead the old density of one person per ten square meters was restored by the families taking in tenants (usually "cousins", real or phony). This is what good economics predicts. The very analysis that tells us that recipients prefer subsidies in cash to subsidies in kind -- that analysis also tells us that people have an incentive to turn in-kind subsidies into cash, which they can then spend on what they (the recipients) rather than we (the donors) think is most important.

Thus we see in the U.S. the (illegal) use of food stamps to buy non-food items, and even the outright sale of food stamps on street corners. And Indian parents were found to convert at least part of their children's hearty school lunches into cash,, simply by having smaller breakfasts and dinners hat home, with the rest of the family eating a larger lunch than had previously been their custom.

Such misfirings of subsidies in kind are obviously much harder to achieve in the cases of education and medical care, which, perhaps partly for this reason, probably account for the largest share of in-kind subsidies.

the standard costs of each successive level of education. We could then say that if a family's children started with the educational prospects typical of the first decile of the population, and moved as a result of a project to have prospects typical of the second decile, this move would be assigned an external benefit equal to 50% of the standard costs associated with such a move. For a move from the educational prospects of the second decile to those of the third, an external benefit of 40% would be assigned. Further increments would be assigned declining benefits until only a 10% external benefit would be assigned for the move from the prospects of the 5^{th} decile to those of the 6^{th}. Moving beyond this latter level, no basic needs externality would be assigned.

What we have here is a way of conceptualizing a basic needs benefit, and of assigning to increments of filling that basic need a set of values that start higher, then get progressively lower as more of a given need is already met, and finally reach zero at some point. Who assigns these weights -- the government? the education ministry? the planning office? the budget bureau? Obviously any of these, but most likely it would end up in the hands of the entity in charge of the government's system of cost-benefit analysis, operating with lots of advice (and pressure) from any and all of the above.

One way in which a group of beneficiaries could be assigned an education externality would be for that group to experience a move in their actual income, from, say, the first to the third decile. This could readily happen, say, if a shipyard employing 1000 workers were to be located in Fortaleza (in the middle of Brazil's poverty-ridden Northeast) instead of in Santos (the port city of Sao Paulo, Brazil's center of wealth. The project could pay government wages, which in Santos might be just equal to the alternative earnings of the workers, but could represent a big wage premium if paid to Fortaleza workers. These could then be actually lifted from the first to the third income decile, as a result of the shipyard project.

Another way in which an education externality could be generated is for a project to enhance school attendance. In poor areas one way of achieving this goal is to provide full school lunches for the students. This measure has been found to be quite effective in improving attendance and delaying school-leaving, even in cases where its effects on nutrition were dubious. More recently, positive results have been achieved by giving parents specific cash rewards based on school attendance by their children. This type of subsidy is comparatively easy to evaluate within the suggested basic needs framework. Thus, a subsidy of up to 50% of standard annual schooling costs would be justified in an area with incomes in the first decile, if it shifted that area's school-leaving profile to the one corresponding to second-decile families. It's a very good guess that this shift could be bought with a much lower subsidy -- say 10% of average school costs -- with the result that the attendance subsidy project would have an external basic needs benefit of (50-10) = 40% of average costs.

The above example shows that the implementation of a formal basic needs approach is quite feasible. But one certainly needs some sort of indicator of each need (analogous to the school-leaving pattern for different deciles). This is relatively easy for housing, where at the most basic level one uses square meters per person as an indicator, but where one can easily introduce an index of housing quality which also includes the type of floor and roof and the presence or absence of glass windows, running water, indoor toilets, etc. For nutrition, likewise, one can use a crude measure like calories, or a more sophisticated one measuring how closely the typical diet (in each decile of the population) approximated a theoretical norm. Health care is perhaps the most difficult basic need to measure. On the one hand, an index could be based on

something as simple as life expectancy (an outcome). But it could also be based on the typical number of visits made by families in each decile to a medical technician, a nurse, a doctor, a basic clinic, a full-fledged hospital, etc. (i.e., an index of medical inputs). I would want also to include public health measures, like closed or open sewers, the availability of potable water, and the degree of pest control (mosquitoes, flies, rats, etc.) in a health care index.

The basic needs externalities concept has the virtue, as against a distributional weights approach, of giving a value to increments of income (of poor people) even when their behavior reflects indifference on their part as between the work they get and the leisure they give up. This advantage stems from the fact that out of an increase in income a reasonably predictable portion will be reflected in children staying in school longer, in the quality of the family's housing being raised a notch of two, in the family's being better fed, and in some improvement in the various measures of health care. The source of this advantage should be recognized to be the element of paternalism that forms the core of the basic needs approach.