ARNOLD C. HARBERGER[*]

1.   THE BACKGROUND

This paper attempts to summarize some of the ways in which the field of economic project evaluation has advanced since the great explosion of literature on the subject that occurred in the late 1960s and early 1970s. That explosion can almost be thought of as the genesis of the field as a subdiscipline of academic economics. It focused on evaluating projects or programs not on the criterion of their financial or commercial profitability but rather on the net benefit that they bring to the "economy as a whole". Quite obviously, this focus led to questions like: a) how to value benefits and costs that are not reflected in financial flows to the project or program in question, b) how to take into account a project's general equilibrium effects on the economy, and c) how to recognize the consequences of economic distortions in creating important gaps between the direct financial benefits and costs of a project and its true economic benefits and costs.

From the outset in the 1960s, the literature of economic project evaluation pursued these questions. Important items of discussion during the "explosion" period were: i) the choice of the appropriate discount rate for deriving present value of benefits and costs, ii) the choice of an "economic opportunity cost of foreign exchange" (EOCFX) that reflected the full economic cost of a project's buying foreign exchange (and the full economic benefit of generating it), and iii) the appropriate way to recognize circumstances where the economic opportunity cost of labor (EOCL) differed significantly from the market wage rate.

Methodological controversies centered on three points. First was the issue of how capital market distortions were to be reflected. One group, to which I belong, advocated incorporating these intertemporal costs in the discount rate. The discount rate would ask that "our" project yield a rate of return sufficient to cover the future productivity that was lost, as other investments were displaced by our project's drawing funds from the capital market. The other view advocated the use of a lower rate of discount, based on the net returns received by savers in the economy. It used this rate to get a present value of the future productivity of displaced investments, and employed this present value as a "shadow price of investible funds" (SPIF). A simple example shows the difference between the two approaches. Suppose that the typical investment yields 12% per year and that, owing to corporate, property and personal taxes the typical saver receives 4%. Suppose, too, that funds drawn from the capital market come 3/4 from displaced investment and 1/4 from newly stimulated savings. These facts would lead the

*   Department of Economics, University of California, Los Angeles.
        Email: grams@worldnet.att.net

first group to use a discount rate of 10% = .75(12%) + .25(4%). The second group would use a 4% rate for discounting project benefits and costs, but would incorporate the present value of future displaced productivity in a SPIF, in this case equal to 2.5. This would be derived by taking the present value, at 4% of the future stream of 10% per year, needed to cover the costs of the investment displaced and the savings stimulated by "our" project. This controversy probably will never be resolved, but real-world practitioners have overwhelmingly opted for the first approach, of reflecting intertemporal capital market distortions in the discount rate.

The second controversy centered on whether to use so-called "border prices" or national market prices in valuing a project's benefits and costs. The use of border prices had the merit of automatically "disallowing" the internal market values that were artificially created by a whole variety of protectionist measures. But this method required that the prices of nontradables be somehow translated to make them comparable to border prices. This was done by using a conversion factor that was in fact the reciprocal of the economic exchange rate (EOCFX) used to reflect international prices under the alternative methodology. The border price approach, too, has fallen into disuse, not because it contained any fundamental error, but because, of two essentially equivalent approaches to the problem, it was the more awkward.

The third controversy swirled around the use of distributional weights in valuing the benefits and costs perceived by different income groups. At first glance it seems both reasonable and appealing that we should give greater weight to a marginal dollar of benefit (or cost) perceived by a poor person than to a similar dollar accruing to the wealthy. Indeed, one can trace applications of this principle in different places in the economics literature, going back at least as far as Jeremy Bentham and the utilitarians. It is interesting that this idea did not occupy center stage — the main stream of applied welfare economics instead being based on a principle of neutrality with respect to the identity of those who perceive benefits and costs ("a dollar is a dollar is a dollar"). Advocates of distributional weights were quite prominent, however, in the project evaluation literature of the 1960s and early 1970s. Once again, this appears to have been a passing fad. At the practical level, real-world project evaluators were typically at a loss as to how to assign weights, and had good grounds for avoiding them simply because of the controversy and opposition that any given choice of weights was bound to generate. Of course this proves nothing about the potential correctness or wisdom of using distributional weights; it only helps explain their falling into disuse. On a more serious level, it can be said that few if any of those who applied distributional weights went much beyond using them for the particular problem they were dealing with — most prominently the problem of so-called optimal taxation in the field of public finance. The failure to pursue their implications more generally was a major flaw, for they have extremely uncomfortable implications when applied to other problems or in other areas. To me, the most convincing case is one in which the government engages in a "take and transfer" scheme, taking from the rich and giving to the poor. Suppose it takes 100 from a rich person with a distributional

weight of 1/2. That person's "weighted" loss is equal to 50. Now suppose that the "take and transfer" scheme is very inefficient and that because of administrative and monitoring costs, etc., the government can only deliver 25 to the transferees. But suppose they have a distribution weight of 2 — in that case the weighted benefit is also equal to 50. This example tells us not only that the operation in question is acceptable — marginal social cost equals marginal social benefit — but that take and transfer schemes with lower than 75% waste-in-transit should definitely be undertaken. If 100,000 of income gave one a weight of 1/2 and 10,000 of income gave one a weight of 2, then a 75% wastage factor would in effect mandate the government to keep on taking and transferring until all incomes were squeezed between 10,000 and 100,000. If the marginal wastage factor was less than 3/4, this band would be correspondingly narrower. Most people, when confronted with this type of implications of distributional weights, quickly withdraw their allegiance to the idea. I believe it can be fairly said that not even voluntary transfers within families — between parents and children and among siblings — come close to reflecting the type of behavior that one would see if these relatives really thought and acted in terms of seriously declining distributional weights.

This exposition up to now simply gives the background for the rest of my story. The sections that follow will highlight new developments in each of the three general areas mentioned, plus a few more.

2.   "CONTEMPORANEOUS" DISTORTIONS LINKED TO THE SOURCING OF PROJECT
       FUNDS

Somehow, even while standard project evaluation procedures recognized distortions with respect to, say, taxes on specific commodities, plus those (tariffs, export taxes, etc.) affecting foreign trade, plus those (corporate income, personal income and property taxes) affecting the intertemporal aspect of capital market sourcing, those same procedures neglected the "contemporaneous" aspect of capital market sourcing.

My colleague and collaborator, Glenn Jenkins and I ran into this problem as we were working on the macroeconomic aspects of a proposed bridge across the Rio de la Plata, between Argentina and Uruguay. One key feature of our problem was the fact that both these countries had value added tax rates that ranged above 20%. Even if we, like the rest of the profession, had previously ignored value added taxes, we certainly couldn't do so in this case. Taking them into account led us down a road that produced important modifications to what previously had been our standard methodology.

The first step, of course, was to recognize that value added tax would likely apply to some project expenditures (outlays) and perhaps to project receipts (sales) as well.

That step was very easy, but once it was taken we had to ask, what about the value added taxes that were displaced in the act of sourcing our funds? These were not the future corporation and property taxes that would now not be paid on

the investments displaced by our sourcing operations. Those reflected the intertemporal effects of our actions and were already dealt with in the standard methodology, as shown above. No, the "new" effect we were concerned with was the *contemporaneous* effect of our entry into the capital market.

It quickly became clear that this contemporaneous effect involved not only value added taxes that were lost as consumption was displaced by our action but also import tariffs that were lost due both to displaced investment and displaced consumption, and excise taxes on specific commodities that were lost, as demand for those commodities was displaced by our capital market sourcing.

One standard way of discussing the external effects of any given action z^* is

$$\int_{z=0}^{z^*} \sum_i D_i(z)(\partial X_i / \partial z)dz$$

Here D_i are distortions, each applying to the distorted activity X_i. Pre-existing D_i's are taken as given. Our action, z^*, is treated as if it came drop by drop, but the integral measures the full effect. In this case the action being taken (z^*) is the raising of funds in the capital market, and the external effects that we are measuring stem from the taxes, subsidies and other possible distortions that exist on activities likely to be affected by our capital market operation. In this case we can use a simpler version of the above expression, namely $\sum D_i \Delta x_i$. Here D_i is the distortion affecting the i^{th} activity and Δx_i represents how the level of the i^{th} activity changes when we raise our project's funds in the capital market.

The end of the road in this exercise could be an overall correction factor of λ_0 per dollar of funds employed by the project. It would represent the combined contemporaneous external effect $\sum_i D_i \Delta X_i$ stemming from the capital market –including lost value added taxes, lost import duties, lost excise taxes on specific products, etc. In some cases it may be convenient to calculate the correction for λ_0 at this point, but in general it is better to wait. The reason will become clear in the next section; basically, it turns out to be easiest to combine the correction for contemporaneous sourcing distortions with another correction reflecting whether the funds are spent on tradables or nontradables.

3.   INTRODUCING SPNTO — TWO SHADOW PRICES INSTEAD OF ONE[1]

Even as our thinking about foreign exchange externality evolved over the years, the basic measure of the foreign exchange externality remained the same. If one thinks of just one import tariff τ_m and one export subsidy z_x we have

---

[1]   SPNTO is the Shadow Price of Nontradables Outlays and EOCFX is the economic
opportunity cost of foreign exchange.

$$\text{EOCFX} = \frac{-(\partial M / \partial E)E(1 + \tau_m) + (\partial X / \partial E)E(1 + z_x)}{-(\partial M / \partial E) + (\partial X / \partial E)}.$$

Here M is imports and X exports, measured in foreign currency units (e.g., dollars), and E represents the country's real exchange rate (the number of real pesos per real dollar). This is obviously a weighted average so one could express EOCFX as

$$\text{EOCFX} = f_1 E(1 + \tau_m) + f_2 E(1 + z_x)$$

We thus thought of EOCFX as a weighted average of : a) the country's real exchange rate adjusted upward for its tariff in imports and b) the same exchange rate adjusted upward for its subsidy to exports (an export tax being a negative subsidy).

When faced with many different tariffs (on different imports) and with different subsidies or taxes on different exports, we could (and did) naturally extend this to

$$\text{EOCFX} = \sum_i f_i E(1 + \tau_i) + \sum_j f_j E(1 + z_j)$$

where f_i represented $-\partial M_i / \partial E$ and f_j represented $+\partial X_j / \partial E$, both divided by $-\sum_i (\partial M_i / \partial E) + \sum_j (\partial X_j / \partial E)$. The EOCFX was here the "market" real exchange rate, modified by how tariff revenues were reduced and export subsidies increased, as a project entered the market to buy the foreign exchange it needed.

We used to add, rather as a footnote than in the main text, that if a given import good M_i was subject not only to a tariff τ_i but also to a commodity tax t_i, then we ought to replace $(1+\tau_i)$ in the above formula with the augmented form $(1+\tau_i)(1+t_i)$ in order to recognize the additional distortion represented by the commodity tax. Needless to say, this adjustment needs much more than a footnote when we recognize that todays value added tax networks cover a wide range of goods, especially tradable goods.

But the needed adjustment goes beyond just incorporating value added and other commodity taxes into the old formula. We have to recognize the inadequacy of the old version of the process which in effect thought in microeconomic terms — i.e., managers go into the foreign exchange market to buy foreign exchange, bid up its price, and thereby displace all sorts of other imports M_i and simultaneously stimulate all exports X_j through a real exchange rate mechanism.

The new way of thinking starts with the project's getting its funds out of the capital market (or some other source of funds if that is more relevant). This sourcing in the capital market displaces demand for both tradables and nontradables. In principle it might happen that the project called for outlays on tradables and nontradables in exactly the proportions in which they were released as the funds were "sourced" in the capital market. In that case there would be no externality

effect to be taken into account. But such a nice matching of spending and sourcing patterns would be a pure accident. What we need is a robust procedure that can handle the full range of possibilities.

To build such a procedure we think of two polar cases –one in which all the project's outlays are in tradables, the other in which all its outlays are on nontradables. In each of these cases we attempt to quantify the whole package of distortions that are involved– up to but not including the taxes or other distortions that apply to the specific goods the project buys (more will be said below concerning project purchases).

Thus, in the case of tradables, we envision raising, say, 100 of funds in the capital market. This may displace 60 of demand for tradables, and 40 of demand for nontradables. But then we spend the full 100 on tradables generating an excess demand of 40 in the tradables area together with an excess supply of 40 in the nontradables. A real exchange rate adjustment is required to close this gap, stimulating exports and displacing imports in the same manner as was envisaged in the traditional treatment of EOCFX. But now imports are displaced only in part by a real exchange rate movement. An important part of the displacement already occurs as the funds are raised in the capital market. And similarly, a significant stimulation of exports occurs as the demand for exportable goods is reduced via the project's funds being sourced in the capital market. The new version of EOCFX incorporates all the distortions involved, both in the sourcing of the funds and in the real exchange rate adjustment needed to bring about market clearing, always under the assumption that the funds in question are being spent on tradables.

Our second major case is that in which all the funds are spent in nontradables. Here, following our earlier example, we have 60 of demand for tradables and 40 of demand for nontradables being displaced as the funds are raised in the capital market. But now the full 100 of funds is spent in nontradables, causing an excess supply of 60 in the tradables market. Again a real exchange rate adjustment is needed in order for markets to be cleared, but it is obviously in the opposite direction from that in the previous case. The peso prices of the dollar has in this case to fall, imports have to be stimulated and exports displaced. In this process the country gains tariff revenue and pays less in export subsidies. This pretty much guarantees that SPNTO, the shadow price of nontradables outlays, will be lower than EOCFX, the economic opportunity cost of foreign exchange, which could equivalently be called the shadow price of tradables outlays.

As I have described them above, the two concepts of EOCFX and SPNTO incorporate the contemporaneous capital-market-sourcing distortions dealt with in Section 2 of this paper. Thus, even though the conceptual framework is considerably more complicated than the traditional one, our revised approach is almost as easy to apply. It is simply a matter of classifying project outlays (and receipts) into two grand categories –tradables and nontradables, and of applying EOCFX to one of these categories and SPNTO to the other.

4.   SEPARATE TREATMENT OF THE DISTORTIONS ON PROJECT EXPENDITURES AND RECEIPTS

The point to be made in this section is in a sense an old one, but it is also one that has not been adequately emphasized in the project evaluation literature. The issue concerns how we should draw the line between the distortions we take into account via standard concepts like EOCFX and SPNTO, and the distortions that we treat individually and specifically in the analysis of each project. On this question the general rule is that concepts like EOCFX and SPNTO are useful only to the degree that they represent operations that are repeated time after time in the world of projects. If a standard set of repercussions is always involved when money is raised in the capital market and spent on tradables, then a concept like EOCFX makes sense. And obviously, one would not want to incorporate in that concept a set of other repercussions that will be different, not only from project to project, but also even from outlay to outlay.

Thus, we know that value added tax is displaced when funds are raised, but we do not know if any particular project will have to pay such a tax on some or all of its outlays (some outlays may be on exempt or zero-rated products; also, some projects may be executed by government agencies that are not required to pay such taxes). Hence both the size of the adjustment and even whether it is made at all cannot be known until the specific details of the project receipts and expenditures are known. Moreover, the adjustment will differ from outlay to outlay even on the same project.

So, when we calculate EOCFX and SPNTO, we naturally are taking into account whether given expenditures fall into the tradables or nontradables category, but we should not build in the value added taxes or the commodity taxes or the import tariffs, etc., that will be paid on those expenditures. This does not mean that such distortions affecting project outlays will be neglected. It only means that because they are so sure to differ from project to project, they must be accounted for separately in the analysis of each project. They cannot be handled as part of a standard adjustment like SPNTO or EOCFX.

5.   YET ANOTHER ADJUSTMENT: THE SHADOW PRICE OF GOVERNMENT FUNDS (SPGF)

At several points in this paper, I have alluded to the conventional assumption that project funds are sourced in the capital market. I have also noted that this is a sensible assumption, for in most cases the capital market is indeed the sponge that absorbs surplus money at the margin, and that can be squeezed to provide for extra outlays as needed. The capital market sourcing convention has been standard in the project evaluation literature for nearly 50 years. A few efforts (by Eckstein and Haveman, among others) were made to promote a convention of sourcing via taxation, but as mentioned earlier, they failed to gain significant support

because there simply is no such thing as a "standard" way in which additional tax money would be raised, and different ways of doing so generate very different patterns of displaced consumption and displaced investment.

Thus the profession has gravitated to the capital market as the canonical source of funds. But that is not the end of the story. In this type of work, we try to make assumptions that help us in our task, but we should not let them carry us into the realm of the unreal or implausible. Let me then take readers down the road that convinced me, a little more than a decade ago, that our methodology had to include a shadow price of government funds (SPGF).

I started by pursuing the capital market sourcing assumption, thinking in terms of an electric power project executed by a public sector power company. The capital market was used, in this example, in a very traditional way — the company borrowed the money to erect a generating station, then sold the output of that station over the life of the project, using the proceeds to pay back the loan, and having enough left over to compensate all the externalities involved. This case fits neatly into the standard way project evaluators have done things.

But my next step was to consider a highway project, one that I assumed to have an identical project profile to the power station of the earlier example, *except* that the benefits in this case were not in the form of cash to the project entity but rather in the form of reduced travel costs (including time costs) to the users of the road. Following the capital-market-sourcing convention, this project borrows to cover its investment costs to begin with and borrows further to cover its operating and maintenance costs over the course of its economic life. At no point in this economic life does the project generate a cash flow to pay back these debts — even though its benefits are sufficient to do so, these benefits are enjoyed by the road's users, not collected by the project authority.

Thus, it is as if, when the project ends, there is a huge accumulated debt that it generated during its lifetime. What then happens to this accumulated debt? As I see it, the implicit assumption of the capital-market-sourcing convention was that this debt just keeps on accumulating forever, with interest being added every year, and with the debt never being paid off, even in our conceptual vision of it.

I found this a very unsatisfactory state of affairs, and was led to add a new assumption or convention — namely, that we try to "close the books" at the end of a project's life by assuming that tax proceeds are used at that point in time (year N), to pay off any net debt that under our capital market assumption would have accumulated during the project's life. Raising this tax money would almost certainly entail a significant "excess burden". We can say this even though the marginal cost of extra money differs for each source, because there is a positive and significant marginal excess burden for just about every tax in a typical governments arsenal. The suggestion then, is to make a reasonable estimate l_g of the likely excess cost per dollar of extra tax revenue, and assign this as the extra cost of paying off the net accumulated debt at the end of a project's economic life. A number of estimates have been made of the marginal economic cost of extra tax money, each under a different set of assumptions. My own inclination was, and is, that the course of wisdom for us is to choose a factor for SPGF that falls within this

plausible range, but that is on the conservative side. A premium λ_g of 20 to 30 percent would fulfill this requirement for the U.S. and many other countries. This probably underestimates the true marginal cost of tax money, but has the clear virtue of not exaggerating that cost.

The final step in this process is to recognize a key equivalence. Letting B_{gt} and C_{gt} equal cash benefits and cash costs to the government at time t, the unpaid debt at time N will be

$$L_N = \sum_{t=0}^{N}(C_{gt} - B_{gt})(1 + i_e)^t$$

Thus if we are going to impose on the project, at period N, an extra charge equal to $\lambda_g L_N$, it is mathematically equivalent for us to introduce an SPGF equal to $(1+\lambda_g)$, and apply this shadow price to each and every cash outlay C_{gt} and cash inflow B_{tg} of funds to the government.

This is the story of how I personally became convinced of the necessity and usefulness of a shadow price of government funds. Many others have since been persuaded by this same line of argument. I hope that readers of this paper will also find the argument persuasive.

Obviously, it is not possible to somehow incorporate SPGF into the major macroeconomic shadow prices SOCFX and SPNTO, for they apply to all cash outlays and receipts of a project, while SPGF only applies to such outlays and receipts when they come from or go to the government. Hence, in this case, further "simplification by merger" is ruled out.

6.   BASIC NEEDS EXTERNALITIES INSTEAD OF DISTRIBUTIONAL WEIGHTS

In Section 1 it was explained why distributional weights had been pretty widely rejected or had otherwise fallen into disuse over the past 25 years or so. This does not mean, however, that so-called "social" considerations have been banished from the field of economic project evaluation. Quite to the contrary, such considerations have probably increased in importance in the most recent quarter century. What has occurred is that these considerations have been taken into account in ways that are different from distributional weights.

Probably the most common way to take such "social" considerations into account is by ad hoc decisions, outside the mechanisms of formal project evaluation. I must emphasize here that there is nothing wrong with doing it that way. For example, a shipyard built in the port of Santos (near Sao Paulo) might have an estimated net present value of U.S.$200 million, while a similar shipyard built in Fortaleza (in Brazil's impoverished northeast state of Ceará) might have an estimated NPV of, say, U.S.$180 million. It seems to me that the Brazilian authorities could be perfectly justified in deciding that it was nonetheless worthwhile to put the shipyard in Fortaleza. This would reflect the implicit judgment that the "social" benefits of generating, say 1000 new jobs for shipyard workers in Fortaleza, over an extended future period, was worth the "price" of $20 million.

This, however, puts the social decision outside the purview of the project evaluation process as such. Is there a sensible way of bringing it back into the project evaluation process? The answer is yes, and the relevant device is the concept of "basic needs externalities". In what follows, I will give a simple illustration of how such externalities can be brought into play.

Basic needs externalities can conveniently be put into four major groups: education, health care, nutrition and housing. Let us consider measuring these attributes in the form of indexes, with 100 being a fully acceptable level. The notion of basic needs externalities can then be made quite concrete, by saying that society is "willing to pay" a premium of $\underline{x}$ to help bring a family's housing level from index 80 to index 81, but only $1/2X$ to bring it from 90 to 91, and nothing at all to bring it from 100 to 101. Similar premia would apply in the cases of education, health care and nutrition.

Now we come to the shipyard in Fortaleza. Its 1000 jobs, we suppose, will carry wages significantly above the typical alternative pay that those workers would earn. By studying patterns of expenditure of Fortaleza families we can estimate the likely improvements in the levels of education, health care, housing and nutrition that would be generated, for those 1000 families as a result of placing the project in Fortaleza. Valuing these increments according to society's schedule of "willingness to pay" as described above, we can reach a total sum to represent the external social benefits generated by placing the project in Fortaleza rather than in Santos (the assumption here is that the Santos workers would be above the poverty level, with or without the project.)

This type of mechanism for dealing with the social benefits of a project is still in its infancy, but it has great promise as a way of bringing some rationality into government decisionmaking in this very central area of public policy.

In conclusion, I cannot resist pointing out that Chile's longstanding policy of measuring different families' level of poverty via the CAS index represents a[2] very interesting start along the lines that are contemplated in the use of basic needs externalities. I can easily imagine a future in which the official CAS index for a family would actually be composed of sub-indices indicating how the family is faring in terms of education, health care, housing and nutrition. At that point the only further step that would be needed would be to set a schedule of "rates" representing society's "willingness to pay" to see a poor family lifted from one index level to the next, in each of the four basic needs categories. Indeed, the thought of using an improved and extended CAS index as an important input into the valuation of basic needs externalities may help stimulate its early development. If it does, practitioners in the field of economic project evaluation, and in applied welfare economics generally, will have good reason to applaud.

---

2   CAS index is a Social Stratification Instrument used for selecting beneficiaries to several social programs. Among them, family money allowances (SUF), Old age pensions for indigents (PASIS) and housing subsides targeted to poor people.
      In the CAS index, each family receives an score that summarized the impact of four factors: housing, householder educational level, householder occupational position and income and durables.