

# Pioneers in Development

Second Series

Theodore W. Schultz  
Gottfried Haberler  
Hla Myint  
Arnold C. Harberger  
Celso Furtado

---

Gerald M. Meier, editor

---

PUBLISHED FOR THE WORLD BANK □ OXFORD UNIVERSITY PRESS

## Reflections on Social Project Evaluation

I APPROACH WITH CONSIDERABLE HUMILITY the task of organizing and expressing these reflections, because I do not consider my work in the field of project evaluation to be characterized by great originality. What I have strived for over the years is better expressed by the word "professionalism."

This statement applies at two levels. First, there is the level of that largely unsung host of people serving in budget bureaus, planning authorities, and all types of ministries and agencies (among them the World Bank itself) all over the globe, who strive selflessly to see to it, insofar as they can, that projects not meeting adequate standards are rejected, while those in the social interest are accepted. These honorable people (for I exclude the many who demean project evaluation by using it as a device for "justifying" whatever projects their clients or superiors want) must literally number in the thousands. They, in my view, form a nascent profession. To me, the task of helping them develop this profession lies more in distilling the knowledge, wisdom, and common sense that are already part of our heritage than in extending the frontiers of knowledge in any deep sense. Moreover, I have always thought that in doing so one should respect the goal of simplicity as much as possible—seeking principles and rules that can be readily communicated to and understood by this heroic (and hopefully growing) band of practitioners, most of whom belong to professions other than economics.

This brings me to the second sense in which the word "professionalism" is relevant. The economics profession has been with us for a long time. The insights and wisdom it has accumulated over more than two centuries are probably the most important source from which to derive the principles and rules to guide the new profession of social project evaluation. This is the star that I have tried to follow as I have dealt over the years with project evaluation: to distill from the huge corpus of economic science the concepts that were particularly relevant for the

new profession, and to function as a professional myself (rather than as a scientist) in studying different areas of application.

This same spirit motivates the present paper. In the course of it I will surely have to disagree with other professional colleagues who have dedicated their time and efforts to work in the same vineyard. But even here the spirit of the disagreement is better described by "Who has found the simplest and most convenient distillation of our profession's accumulated tradition?" rather than by "Who has invented what?" or even, on a given point, by "Who is right and who is wrong?"

### Three Basic Postulates of Applied Welfare Economics

The grand tradition of applied welfare economics, going back at least to the days of J. A. Dupuit (1844),<sup>1</sup> can be interpreted as being based on three simple postulates:

1. Competitive demand price measures the benefit of each marginal unit to the demander.
2. Competitive supply price (or marginal cost) measures the opportunity cost of each marginal unit from the standpoint of the suppliers (factors of production).
3. In attempting to measure the benefits and costs to a society (or group) as a whole, take the difference between benefits (+) and costs (-).

On the basis of these postulates, the classical principles of economic policy analysis can be derived. Among these are (1) the measure (a triangle between demand and supply curves) of the so-called welfare or efficiency cost of a tax; (2) the traditional demonstration of why the exercise of monopoly power is not in the social interest; (3) the case for the so-called optimum tariff or export tax, which shows how a nation (though not the world as a whole) can benefit from exploiting its monopoly power (if any) over products it imports and its monopoly power (if any) over products it exports; (4) the generalized expression for the welfare or efficiency cost of taxation in a general-equilibrium setting with many commodities (developed independently by many authors but perhaps most elegantly by Hotelling); (5) various rules for optimal taxation, of which the so-called Ramsey rule (1927) is among the most famous; and (6) the so-called Lerner theorem (1936) showing

1. "De la Mesure de l'Utilité des Travaux Publics," reprinted in *International Economic Papers* 2 (London: Macmillan, 1952), pp. 83-110.

the equivalence of a uniform import tax with a uniform export tax under conditions of balanced trade.<sup>2</sup>

These examples are impressive because of the power of the ideas involved and their significant place in the accepted corpus of economic policy analysis. But truly they are only the tip of the iceberg. The tradition of what is called consumer surplus analysis has had many distinguished representatives, among whom Alfred Marshall and A. C. Pigou stand out as giants.<sup>3</sup>

The three basic postulates have sometimes been criticized as not yielding in every circumstance "true" measures of how the utility of individuals changes when some policy is introduced or some other disturbance occurs. But as a basis for the actual analysis of real-world policies and projects the postulates have not been seriously challenged, let alone surpassed.

Some authors have worried about cases of multiple equilibria, others about examples in which the postulates need not always yield the same measure if a sequence of policies is imposed in a different order (the so-called integrability condition). But I do not know of even a single case where the presence of multiple equilibria has been identified as a factor in an important real-world policy problem. And as for the integrability condition, my favorite analogy is that we all know that the distance between two cities will vary with the route we take. But when asked what the distance is, most road maps, geography books, and airlines present only a single number. The logical answer in the case of intercity distance is sometimes "as the crow flies," sometimes "the shortest available route." The corresponding answer in applied welfare economics (when the sequencing of policy steps is an issue) is to choose the most plausible or likely sequence or, if there is none, to assume that all the policies are imposed together (in technical jargon, to assume "a radial expansion of the vector of distortions" from the relevant starting point to the end result).

From my own standpoint, I have always thought of the three basic postulates as exactly that—simple postulates on which a system of measurement (traditional applied welfare economics) has been constructed. All economists know that national income and gross national product are inaccurate measures of national welfare. Yet they are relied

2. Harold Hotelling, "The General Welfare in Relation to Problems of Taxation and of Railway and Utility Rates," *Econometrica*, vol. 6, no. 3 (July 1938); Frank P. Ramsey, "A Contribution to the Theory of Taxation," *Economic Journal*, vol. 37 (March 1927), pp. 47-61; and Abba P. Lerner, "Symmetry of Export and Import Taxes," *Econometrica*, vol. 3, no. 11 (August 1936), pp. 306-13.

3. Alfred Marshall, *Principles of Economics* (London: Macmillan, 1890), and A. C. Pigou, *The Economics of Welfare* 4th ed. (London: Macmillan, 1927).

on for most studies. In spite of their defects, they have performed reasonably well in most contexts. What can be said of the three basic postulates is that they are considerably more subtle and more refined than the rules on which national income accounting is based—they would not, for example, make the mistake of implying that welfare falls when mothers voluntarily leave the labor force to take care of their homes and families. In every case that I know of where the three postulates lead to results different from those derived by the rules of national income accounting, the postulates (as in the above example) win hands down. For the real-world problems they have been used to solve, they have proven more adequate than anything else we have. They thus are the natural sources from which the principles to guide the budding profession of social project evaluation should be drawn.

#### On Distortions and Externalities

Using the three basic postulates makes it easy to understand the sense in which the grand tradition of economics has always looked upon an undistorted and fully competitive economy as an optimum. If demand price as seen by demanders in each market is equal to supply price as seen by suppliers, and competition prevails, marginal social benefits as measured by postulate 1 will always be equal to marginal social costs as measured by postulate 2. This happy state no longer prevails when distortions or externalities are present.

Distortions can take on many forms, of which the simplest to analyze are taxes and subsidies. When these are present, marginal social benefit as measured by the price paid by demanders differs from marginal social cost. If we represent the excess of demand price over supply price in an activity by  $D_i$  and the change in the level of that activity by  $\Delta X_i$ , the introduction of a new policy or project will produce net benefits or net costs, through its effects on other markets, according to whether the expression  $\sum_i D_i \Delta X_i$  is positive or negative.

I have always thought of  $D_i$  as standing for "distortion," with the term construed very broadly. That is to say, the above expression is valid quite generally—not only for taxes and subsidies but also for the distortions present in more complicated public policies, and for those stemming from monopoly and monopsony situations. Externalities that vary with the level of an activity, such as traffic congestion with the volume of traffic or smoke pollution with the volume of a factory's output, can also be treated as distortions ( $D_i$ ) within the same simple formula.

All this is important because, as will be seen below, these distortions are precisely why (within the framework of the three postulates) we have to build a system of social project evaluation that is different from

the simpler economics of a distortion-free world. In particular, the pervasive presence of important distortions is the main reason it was necessary (or at least exceedingly useful) to construct these concepts.

#### On Social Opportunity Costs in a Market Setting

As pointed out above, the postulates would assert that in the absence of any distortions in the economy, the social opportunity costs of marketed goods and services would equal their market prices. When new demand for a good or service is generated (say, by a new project) there are only two sources from which that demand can be satisfied—increased total supply and the crowding out (displacement) of other demanders. Postulate 2 tells us that the increased supply should be evaluated at the supply price, while postulate 1 tells us that the displaced demand should be evaluated at the demand price. In the presence of a distortion, say, a tax, these two prices are different and the social opportunity cost becomes (at least in simple cases) a weighted average of them. Another way of expressing this is that demand price  $p_d$  is equal to supply price  $p_s$  plus the distortion  $D_i$ . If the social opportunity cost of one unit of the good or service is a weighted average  $f_1 p_d + f_2 p_s$ , this can be expressed as  $(f_1 + f_2)p_s + f_1 D_i$ . But  $f_1$  is simply the measure of the amount of displaced demand, so the above formula boils down to saying that the social opportunity cost of a unit can be expressed in the form  $p_s + D_i \Delta X_i$ .

The above is the most rudimentary example of a general rule. Other cases are more complicated but come down to the same thing. A simple capital market example contains two taxes, a capital income tax  $t_c$  and a personal income tax  $t_p$ . There are thus three rates of return:  $\rho$ , the gross of tax return to investment;  $i (= \rho - t_c)$ , the market rate of interest; and  $r (= i - t_p)$ , the after-tax return received by savers. In this market funds raised by a project come in part from displaced investment, which by postulate 1 has an opportunity cost of  $\rho$ , and in part from newly stimulated saving, which by postulate 2 has an opportunity cost of  $r$ . The social opportunity cost of capital can then be expressed as  $f_1 \rho + f_2 r$ . But this is also equal to  $f_1(i + t_c) + f_2(i - t_p)$ , so, given that  $f_1 + f_2 = 1$ , it is equal to  $i + f_1 t_c - f_2 t_p$ —that is, it is a market rate of interest adjusted for distortions by the  $\sum_i D_i \Delta X_i$  principle.

No matter how many sectors we add, the fact remains that it is substantially equivalent in such cases to look at social opportunity cost as a measure that is a weighted average of the demand prices  $p_d^i$  of those sectors whose demand was displaced by the entry of new demand, and of supply prices  $p_s^i$  of those sectors whose supply was stimulated, and as a measure that is a market price (like  $p$ , in the first case or  $i$  in the second) adjusted by a factor based on the application of  $\sum_i D_i \Delta X_i$ . Here

the  $\Delta X_i$  represent precisely the same displacements of demand and stimuli to supply incorporated in the weights of the weighted average. Perhaps the most elegant case is the representation of the social opportunity cost of foreign exchange in terms of a weighted average of a whole host of individual terms, which reflect the different tariff treatment of many categories of imports and the various taxes and subsidies applying to different categories of exports. For a dollar's worth (the dollar being the main foreign exchange unit for most small countries) of imports of good  $j$  the demand price (postulate 1) is  $E_m + T_j$ , and for a dollar's worth of exports of good  $k$  the supply price would be  $E_m + S_k$ . Here  $E_m$  is the market exchange rate for the dollar, and  $S_k$  the subsidy per dollar of exports of  $k$ . If there are many different categories we have weights of  $f_j$  and  $f_k$  reflecting the fractions of a dollar newly demanded (say, by a project) that are brought about through displacing imports of  $j$  or by stimulating an increment in exports of  $k$ . The end result of this exercise is the following expression for the social opportunity cost of foreign exchange:

$$\sum_j f_j (E_m + T_j) + \sum_k f_k (E_m + S_k).$$

This is a weighted average of the demand prices (postulate 1) of many different classes of imports and of the supply prices (postulate 2) of many different classes of exports. But since  $\sum_j f_j + \sum_k f_k = 1$ , the expression above can equally well be written as

$$E_m + \sum_j f_j T_j + \sum_k f_k S_k.$$

This takes the form of a market rate of exchange adjusted by the  $\sum_j D_j \Delta X_i$  principle.

In all the above cases the weighted average of demand prices for displaced demand and of supply prices for newly stimulated supply (based on a quite intuitive application of postulates 1 and 2) is equivalent to the representation of social opportunity cost as a market price, duly adjusted by the same weighted average of the relevant distortions (based on the  $\sum_j D_j \Delta X_i$  principle, a less direct but more subtle application of postulates 1 and 2).

This leads us to see that the standard weighted-average measures of social opportunity cost can be regarded as attempts to follow an approach that (though subtle and correct) asks us to look at the effects of a given action on *all* the distorted activities of the economy, and to assemble the adjustments for distortions into convenient "packages." For the weighted-average measure, the elements in the package are the supplies and demands in the various component parts of the given broad market. Thus, for the social opportunity cost of foreign exchange, the weighted-average measure includes the supplies of all the various export categories and the demands for all the various import

categories, while for the social opportunity cost of capital it would include the demands for all the various investment categories and the supplies of savings from different sources (presumably with different marginal tax rates, hence different supply prices of their savings). Of course, in all these cases, several categories can be lumped together if they have the same, or closely similar, distortions.

But there may be times when the relevant package includes elements above and beyond those in the weighted average. In deciding on the relevant package one must recall that in determining social opportunity cost we are attempting to trace what happens when new demand enters a market. Weighted-average methods trace out those consequences just on the constituent parts of that market—displacing imports and stimulating exports, for example. These displacing and stimulating effects work through the real exchange rate, and one must recognize that there may be other consequences. I have suggested, for example, that if a rise in the real exchange rate ends up displacing some imports of petroleum, there will be a side effect in the form of reduced receipts from gasoline taxes. I have also suggested that if there existed a tax on bricks, and if the introduction of new demand into the capital market displaces construction, there would similarly be a side effect in the form of reduced receipts of that tax. When these effects are relatively minor, the principle of simplicity suggests ignoring them. But when they are significant, they should be taken into account. Most important of all, perhaps, from the standpoint of the present paper, is the clear conclusion that weighted-average measures of social opportunity cost are not in themselves *the* correct solution. In a market situation they will always be a component of the correct solution, but they may at times need to be supplemented in important respects. All of this, of course, follows directly from the three basic postulates that have guided applied welfare economics from its beginnings.

#### *Some Additional Observations*

Let me set the stage by pointing out an important characteristic of a well-functioning capital market or foreign exchange market. Particularly in the latter case, one can almost guarantee that the reaction will be the same: the real exchange rate will go up by the same amount, and the same set of displacements of imports and stimuli to exports will occur—regardless of who is the buyer of a given amount of foreign exchange. In the case of foreign exchange in well-organized markets, the names of the buyers and the purposes for which they buy are not even known by the market, most purchases being made by financial institutions or other intermediaries. The market simply "feels" the pressure of an added demand, and a set of market-determined reactions ensues.

Some people dispute the above proposition by saying that it does not take into account the use to which the foreign exchange is being put, and that when this use is taken into account the social opportunity cost of foreign exchange will vary from case to case according to the attributes of the use. I believe that this line of reasoning misses the whole point of the concept of the social opportunity cost of, say, foreign exchange. The main reason for assembling certain  $\Sigma D, \Delta X_i$  in packages is that these packages turn up with great frequency. When this is so, one can calculate the package solution just once and save lots of time and effort.

This, I believe, is what is accomplished when we calculate the social opportunity cost of foreign exchange. To my mind, the social opportunity cost of foreign exchange (in a market setting) has no relation to its use. To put it graphically, if an enterprise or individual in country A enters the market, buys dollars, and then suffers a fire in which the dollars burn up, what is the loss to the country? This loss is measured by the forces of import displacement and export stimulus and is the same regardless of who was the buyer or what intentions he might have had about the use of the now lost dollars. Here is a perfect example of how social opportunity cost can be separated from use. Of course, the use of foreign exchange cannot be neglected in social project evaluations. But here, by the very nature of the case, a different pattern of effects ( $\Sigma D, \Delta X_i$ ) in distorted markets will take place in almost every operation. Foreign exchange spent on highly taxed imports ( $D_i$ , positive and large) will produce an important benefit, which should be attributed as a benefit of the project. If spent on subsidized imports ( $D_i$ , negative) it will produce a cost. Application of the three basic postulates requires that these benefits and costs be taken into account in the analysis of each project. But this does not annul the usefulness of the package represented by the social opportunity cost of foreign exchange. Quite the contrary, when a project buys foreign exchange the package is a cost, and the analyst need only examine the  $D_i, \Delta X_i$  that are involved in spending it. Similarly, when foreign exchange that has been earned is sold in the market, the package is a benefit and the analyst need only examine the  $D_i, \Delta X_i$  that occurred as a consequence of the particular activities by which it was earned.

The same attribute is pretty much shared by the social opportunity cost of capital. Again, when money is drawn from a reasonably well-functioning capital market, people rarely know for what purpose. But the withdrawal of funds tightens the market and produces a package of responses that is likely to be very similar, regardless of who was doing the withdrawing. There is a difference between the foreign exchange market and the capital market, however. In the former the price paid (say, for a dollar of foreign exchange) is essentially the same for every buyer, and so also are the market reverberations that produce the

difference between the social opportunity cost of foreign exchange and its essentially uniform (among buyers) market price. In contrast, the rates of interest paid by different classes of borrowers vary quite substantially, but the market reverberations that follow from the withdrawal of funds from the capital market are likely to be virtually the same, regardless of who withdrew them. To put it simply, the lender learns who the borrower is and charges him a supply price; but the market, which then reacts to the greater scarcity of funds, does not know who the borrower is. Its reaction is governed only by the fact that funds are tighter.

Thus the reverberations ( $\Sigma D, \Delta X_i$ ) may cause the social opportunity cost of capital to be 4 percentage points above the market price, but the market price might be 8 percent for the government, 10 percent for a good industrial borrower, and 12 percent for a normal commercial borrower. In this case, the social opportunity cost of funds would be 12 percent for the government, 14 percent for the industrialist, and 16 percent for the commercial borrower.<sup>4</sup> Most actual social project evaluations deal with public sector projects, for which the government rate is appropriate. In some developing countries, however, there have been serious analyses of the social costs and benefits of projects to be undertaken by the private sector, and in such cases the appropriate procedure is to adjust a private (real) cost of borrowing (10 percent in the industrial case above) by a premium that takes into account the reverberations ( $\Sigma D, \Delta X_i$ ) from the borrowing.

Furthermore, the time should come when students of the development process take seriously the idea of making extensive ex post evaluations of both public and private investments. When this happens it will also be appropriate, in principle, to consider the social cost of capital to be different for different broad classes of private borrowers.

As in the case of foreign exchange, there is going to be an entirely different set of reverberations ( $\Sigma D, \Delta X_i$ ) stemming from the way the money is spent. These have nothing to do with the market for capital funds and should appropriately be taken into account in the analysis of the outlays of each project.

### The Social Opportunity Cost of Labor

The application of the three basic postulates to the calculation of the social opportunity cost of labor is vastly more complicated than the

4. These rates should in principle be defined net of the actuarial component of default risk, which is not part of the social cost of borrowed funds. But in fact rates do differ by these amounts, and more in many cases, even after correction for the actuarial component.

applications in the preceding section. There are typically huge variations in the wages of labor according to occupation and skill. Even within these categories there are in many countries substantial regional variations in wages. And one must also deal with the fact that individuals can and do move among occupations, regions, and individual jobs and that they have demonstrably different supply prices in at least some of these cases.

#### The Key Role of Supply Price

To continue with the three basic postulates, the starting point for measuring the social opportunity cost of labor must surely be postulate 2. This has powerful implications. If a worker is willing to work at different jobs, but demands more pay for some of them than for others, he does not have a single opportunity cost of labor; instead, he has as many (private) opportunity costs as he has asking prices. But for any one job only one opportunity cost is relevant—his supply price for that job.

Obviously it is folly to think of conducting real-world social project evaluations by finding out the supply prices of all the different workers hired. The key lies in the presumption that, in the absence of clear evidence to the contrary, market wages are being paid by the project for each job, in each skill, and in each occupation. This in turn leads to the presumption that for most workers the wage represents or approximates their supply price, determined mainly by what they could get for an equivalent job in the same market. Some workers may be willing to take their jobs for less than the market wage and may thus be earning economic rents. But if *all* (or in practical cases nearly all) of the workers on a project are earning economic rents, that project simply cannot be paying market wages—it must be paying above-market wages, a case to be taken up presently.

The clear implication of this is that the social opportunity cost of labor must be treated as an essentially microeconomic phenomenon, at least in social project evaluation although possibly not in cases of major macroeconomic adjustments. The prevailing market wage, although differing among types and categories of labor, must reflect the market-clearing supply price of labor under postulate 2.

Yet this supply price is far from being the social opportunity cost of labor. To get to social opportunity cost we must follow a path directly analogous to the one pursued in the case of foreign exchange and of capital. We begin with a somewhat semantic point. When a worker states his asking price (what he expects to be paid), he typically recognizes that he will have to pay income tax out of the proceeds. In this respect the market wage  $w_m$  is similar to the market interest rate  $i$

presented above in analyzing the social opportunity cost of capital. The true supply price of labor,  $w_s$ , must be defined net of taxes (at least in applied welfare economics) because it is only out of net income that workers presumably gain their welfare and satisfaction. If we take the market wage  $w_m$  as the base, we must recognize that there is a distortion  $t_p$  (personal taxes) lying between  $w_m$  and the true supply price of labor  $w_s$ ,<sup>5</sup> just as a similar distortion lies between  $i$  and  $r$  in the case of the social opportunity cost of capital.

The easiest way to think of the social opportunity cost of labor in a given area, occupation, industry, or other category is to start with the market wage and make a series of adjustments. All the labor of each type earns its relevant market wage. The first adjustment is to recognize that taxes  $t_p$  are being paid on those wages; this represents a positive external effect of the form  $D\Delta X_i$ . But now we must turn to the ultimate sources from which that labor was drawn. Some may have been drawn from other regions, some from other industries within the same region, some from competing firms within the industry. In addition, there will be some who newly joined the labor force, and at least in times of less than full employment there may be some net reduction in the number of unemployed.

For all but one of the above sources there will be some sort of external effect of the form  $D\Delta X_i$ . All three of the first groups will have ceased to pay income taxes in their former place. These lost taxes represent a negative external effect of the project. They may just cancel the positive effect of taxes paid by the workers on the project, but this should not be presumed. Wages may have been higher or lower in the source than in the destination; in countries organized on a federal basis, taxes may also differ from region to region. When it comes to displacing labor from other industries, thought must be given to the taxes that may be levied on the products of those industries. These taxes, too, are part of the marginal product of the factors of production in those industries, for the marginal product of the factors must be valued at the prices paid by the ultimate demanders of those products (postulate 1), but their loss may be partially offset if other factors end up substituting for the labor that was *in net terms* drawn away. Labor drawn from the ranks of the unemployed is likely to produce an external benefit. In cases where unemployment compensation is being paid, any net reduction in the

5. I hope readers will forgive my using the term "supply price" sometimes in a gross-of-tax sense and at other times in a net sense. To use "asking price" as a technical term is, I believe, stretching it too far, but when it does not seem suited or possibly confusing I have used it. In other contexts I have stuck with "supply price." I trust that in such cases it will be evident from the context whether the market wage  $w_m$  or the net-of-tax wage  $w_s$  is being referred to.

number who would probably be unemployed in the absence of the project means a corresponding saving of public funds, presumably to the benefit of the citizenry at large. This saving is an obvious external benefit. It is to be expected (and indeed has been demonstrated many times) that the presence of unemployment compensation causes the asking price of the unemployed to be higher. Where there is no unemployment compensation, the absorption of the unemployed creates a benefit in the form of economic rent (consumer surplus) that reflects the excess of the wage received over the true supply price of the unemployed (which is unlikely to be zero as was so often assumed in the early development literature).

The vicissitudes of measuring the social opportunity cost of labor are vast as, indeed, are those associated with foreign exchange and capital. In practice, the actual measurement is always immensely cruder than the underlying concept. I have tried to show that even though the task is more complex and differentiated for labor than for the other two, the three basic postulates lead to essentially the same general solution to the problem of measuring social opportunity cost. Whether hiring labor, borrowing capital, or buying foreign exchange, we must first determine (estimate, guess, or assume) the pattern of the ultimate sources of supply. That pattern will inevitably entail some sourcing from new supply and other sourcing from displacement of other demand. In simple cases one can merely assign the net supply price as the opportunity cost of new supply and the gross demand price as that of displaced demand. In more complicated cases, which the labor market exemplifies, the more subtle approach of working with a market price corrected by distortions is indicated. In the case of labor we start with the market wage and correct it for all relevant externalities: taxes forgone in the various sources, unemployment compensation saved, and consumer surplus gained by newly hired workers. We need also to remember to introduce a credit for income taxes paid by labor in the destination.<sup>6</sup>

#### *The Social Opportunity Cost of Labor with a Protected Sector*

A final complication to the labor story is the frequent existence of what I call "protected sectors." Others have spoken of duality in the

labor market, of traditional versus modern sectors, and so on. I like the term "protected sector" because it quite directly connotes a wage above what would be a free-market equilibrium level. It also connotes that some element (such as minimum wages, union bargains, or "political insurance" by multinational companies) must be at work to create the differential.<sup>7</sup> In addition, the protected-sector concept, more readily than its alternatives, opens the door to the idea of not just two but quite a number of different protected sectors, ordered hierarchically according to the wages they pay for equivalent work.

Another phenomenon in conjunction with protected sectors is that of quasi-voluntary unemployment. Consider that there is a given protected-sector wage ( $w_p$ ) and a free-market wage ( $w_f$ ). The protected sector cannot hire all who want to work there, for to do so would mean bidding up the market wage all the way to  $w_p$ , in which case the dichotomy between the two sectors ceases to exist. With a protected sector of moderate size (compared with the size of the relevant labor market), the normal labor market structure is one in which there are a significant number of people with supply prices between  $w_f$  and  $w_p$ , who were not lucky enough to be employed at  $w_p$  but who have no interest in working at  $w_f$ . These are the quasi-voluntary unemployed. In such a situation an expansion in the number of jobs in the protected sector will draw workers in part from the free-market sector and in part from the quasi-voluntary unemployed. The average supply price and hence private opportunity cost (postulate 2) of those who fill the protected-sector jobs will accordingly be above the free-market wage. Quasi-voluntary unemployment is thus an institutionally induced phenomenon which curiously entails an opportunity cost of the unemployed that is above rather than below the free-market wage.

If the setting of the protected sector is the urban labor market, and if there is ready migration from rural areas, it is likely that the presence of the protected sector will induce rural-urban migration. This has the effect of driving down the free-market wage and swelling the ranks of the quasi-voluntary unemployed. This is really a case of migration-induced unemployment noted early by W. A. Lewis and analyzed in some depth by Harris and Todaro, by myself, and by others.<sup>8</sup>

7. The protected-nonprotected distinction does not refer, for example, to high-skilled as opposed to low-skilled jobs or to experienced as opposed to inexperienced workers. It refers to workers of identical characteristics and abilities receiving different rewards in essentially identical jobs in two sectors of the same labor market.

8. W. Arthur Lewis, "Development Economics in the 1950s," in *Pioneers in Development*, Gerald M. Meier and Dudley Seers, eds. (New York: Oxford University Press, 1984); J. R. Harris and M. P. Todaro, "Migration, Unemployment and (Note continues on following page.)"



price and is not allowed to adjust, then something else (here unemployment) will develop so as to bring supply price up to demand price.<sup>9</sup>

### Social Opportunity Cost in Nonmarket Situations

Except in the preceding section, I have till now dealt with a world of market prices—prices that may be distorted by taxes, tariffs, or subsidies, but nonetheless prices that clear the market. Here I turn to cases where this is no longer true—cases of licenses, quotas, arbitrary rationing, and the like. To begin with foreign exchange, consider a country with a fixed exchange rate system that uses licensing rather than monetary policy as the principal mechanism of adjustment. That is to say, export proceeds are required to be turned over to the exchange authority, which then doles them out among the many applicants. If there is enough foreign exchange to fill all applications at the going exchange rate, the licensing system is redundant and the problem reverts to the cases discussed earlier. But if—as has actually been the case in many countries for long periods—there are far fewer dollars available than are desired, a problem of allocation appears, which is somehow resolved by the licensing authorities. Characteristically in such cases the demands for some “essential” imports will be fully met, others not so fully, others quite partially, and some (the prohibited list) not at all. It is plausible to assume that the tariff structure probably already recognizes this hierarchy, with perhaps zero duties on “essentials,” moderate duties on the next group, and progressively higher ones on the remaining two.

What do the three postulates tell us here? Only for the first group would demand price and supply price be equal to market price. For the other groups the demand price would be above the supply price even if all the desired dollars were made available. Because dollars are not made available, there must be some unsatisfied demanders willing to pay prices well above the world market price plus the tariff. If such demanders are given incremental dollars, their demand prices should be treated as a measure of the benefit (postulate 1).

The easiest case to deal with is that in which there is an open market for the domestic resale of the licensed items. In this case the appropriate demand price is the one prevailing in that market. The social opportu-

9. In this discussion I follow most treatments of this topic and simplify the problem by abstracting from the types of distortions (income taxes, product taxes, unemployment compensation, and so on) treated earlier.

Migration-fed unemployment is the result of the phenomenon of rent seeking, which in turn is easily rationalized in terms of postulate 1 to 3. Wages offered in the destination are above those required to induce migrants to move. The normal workings of an open labor market would entail a fall in wages as increased supplies of workers made themselves available. That is to say, the wage rate itself would serve the function of equilibrating the labor market and in doing so would stem the flow of migration. When wages are maintained at protected levels and are prevented from adjusting, labor market equilibrium is nonetheless brought about in some other way. In the case of migration-fed unemployment, the unemployment rate itself brings it about. One starts with a supply price of migration (a wage at destination that would just barely induce the migrant to move). Initially, the wage at destination is above this supply price, and the migration occurs. But as unemployment at the destination mounts, the supply price of migration increases. That is, in the presence of greater unemployment, the same destination wage becomes less attractive, and the incentive to migrate is reduced. With enough unemployment the tide of migration is stemmed; but unemployment at the destination remains high as a continuing equilibrium phenomenon as long as the wage paid at the destination remains above the supply price of migration that would prevail in the absence of unemployment.

I like this particular way of viewing the adjustment process via migration-fed unemployment, because it is framed directly in terms of supply prices and demand prices. The demand price is the fixed wage at the destination; the supply price is that wage which would just barely induce (or compensate) migration. As unemployment mounts, this supply price rises, and the adjustment becomes complete when this supply price equals the fixed destination wage.

This demonstrates very simply the usual result that under conditions of migration-fed unemployment the social opportunity cost of labor will end up being equal to the fixed destination wage. Raising the latter would only increase unemployment until the supply price of migration again equaled the wage. Not only do we reach this result simply, but once again we see the power of the three postulates. The answer is a profound expression of postulate 2. If demand price exceeds supply

Development: A Two-Sector Analysis," *American Economic Review*, vol. 60 (March 1970); and Arnold C. Harberger, "On Measuring the Social Opportunity Cost of Labour," *International Labour Review*, vol. 103, no. 6 (June 1971), pp. 559-79.

nity cost of foreign exchange will in such cases be a weighted average of the market prices of the various categories of imports. The problem is that the weights here are not derived from the structure of demand and supply and from the normal elasticity of response of different items within that structure to changes in the degree of ease or pressure in the foreign exchange market. Instead, the weights are derived from the policy decisions of the licensing authorities, which—particularly with respect to items not in the “essential” category—tend to change with great frequency; items “serve sentences” on the prohibited list only until enough political pressure can be mustered to take them off. In a case where imported goods can be freely resold in the local market, it would be easy to estimate the social opportunity cost of foreign exchange if there is a stable, consistent, and predictable licensing policy. I know of no case, however, where these conditions have been met. If policy is governed by sufficient rationality to do all these things well, it seems highly likely that the authorities would then take the next logical step of following monetary and exchange rate policies that would render the licensing system superfluous.

When there is no open market (either legal or parallel) for the licensed items the problem becomes one of estimating their value to the users. The principle of measuring the social opportunity cost of foreign exchange remains the same: when the licensing authority gives foreign exchange to a new project, it must come at the expense of other demanders. The opportunity cost is then the value (demand price) that those other demanders would place on the amounts of which they are being deprived. This is obviously impossible even to try to determine, so the estimation of the social opportunity cost of foreign exchange becomes a very crude process indeed. The only solid base to work from in this case is that we know that any licensed import is worth at least the world price plus the prevailing tariff to all license applicants, because that is what they will have to pay even if they get the license.

On the side of the social opportunity cost of capital, interest rate ceilings of various kinds are probably the most frequent source of deviation from a market solution. The key element in these situations is that banks and other financial institutions are sometimes faced by a demand for credit that far exceeds the available funds—especially since the same demanders get in line at many banks. Characteristically, service charges, minimum compensating balances, and other devices appear. They may be powerful enough to produce a market solution—that is, one in which no arbitrary rationing is involved. In this case the actual total payments by the users of credit would have to be determined in order to estimate their demand price according to postulate 1, but no other serious complications emerge.

Frequently, however, the equivalent of a market solution is not

worked out, and banks and other financial institutions are left with a substantial range of discretion. This opens the door to bribery and corruption, which indeed frequently occur. But there are more subtle ways, generally fully or nearly within the law, of doling out credit in these circumstances. For example, when a borrower hires a banker's relative at a good salary, his subsequent loan applications may more readily be approved. Thousands of ruses of this type exist, and given the controlled interest rate the banker can fairly argue that his choice among alternative solvent borrowers has cost his stockholders nothing. The net result is that credit is rationed in inefficient ways, and the marginal productivity of capital, overall, is lower.

To determine the influence of this situation on the social opportunity cost of capital one must inquire what happens when financial institutions are left with lower lending capacity owing to some savings having been diverted to finance a new project. No clear answer exists for this case, but the most reasonable assumption is that, given the ability of financial institutions to ration funds, their basic behavior patterns will not change significantly. Thus we can say, roughly at least, that in these rationing situations the social opportunity cost of the diverted funds will be greater than the controlled interest rate, and less (because of the inefficiency factor noted above) than the typical gross-of-tax productivity of capital ( $p$  in the simplest case) that would prevail in an open capital market.

The above are a few examples of how one can attempt to cope with what I have called nonmarket situations in the measurement of social opportunity cost, using the three basic postulates as a guide.

### The Social Rate of Discount

The social rate of discount has been at the center of controversy for many decades. Early debates raged over whether to use the marginal productivity of capital, the market interest rate (usually the government bond rate), or the marginal rate of time preference. These concepts correspond broadly to  $p$ ,  $i$ , and  $r$ , as presented earlier. The recent discussion has been less simplistic. The major contending views, at least as I see them, all recognize that both  $p$  and  $r$ , when they are different from  $i$ , reflect distortions, and all recognize that when this fact is properly taken into account it leads to situations in which social project evaluations yield different results from private ones. In my opinion, the major contending views are all basically compatible with the three basic postulates. The squabbles are not between saints and sinners, but rather among factions within the same church. Their resolution, in consequence, turns on issues of efficiency, relevance, convenience, robust-

ness, communicability, and the like, rather than fundamental error or heresy.

Three alternative approaches to the social discount rate will be treated; they bear a curious sort of triangular relationship to each other. For example, two of them are based on the convention that marginal funds come from the capital market, while the third adopts the alternative convention that they come from fiscal sources. In a different pairing, two of them use weighted averages of  $p$  and  $r$  as the discount rate; the third uses  $r$  alone.

Consider first the dichotomy between the view that the funds for public sector projects come from fiscal sources, and the alternative view that they come from the capital market. The proponents of fiscal sourcing (Eckstein, Eckstein-Krutilla, and Haveman,<sup>10</sup> among others) argue that most government funds come from fiscal sources and that most of the increment in government funding over time has been, and almost inevitably is, mainly on the fiscal side. In this they are correct. They then proceed to postulate a set of weights  $f_1$  and  $f_2$  reflecting the fractions of an increment of fiscal revenues that come at the expense, respectively, of investment and consumption. Their final formula, a social discount rate equal to  $f_1 p + f_2 r$ , is identical in form to the one derived earlier in this paper. The only difference is that the weights  $f_1$  and  $f_2$  here derive from a hypothetical fiscal experiment rather than a hypothetical capital-market experiment.

As noted earlier, in a functioning capital market we have reason to expect that the reactions to market pressure or ease will on the whole remain quite similar from one case to the next. In particular, we can be quite sure that just about everywhere the elasticity of investment with respect to market pressure substantially exceeds that of saving. In contrast, experience all over the world suggests that each tax change is very different from the last. The weights  $f_1$  and  $f_2$  may bound all over the map, depending on whether one is imposing value added taxes, tightening income tax loopholes, lowering high or prohibitive tariffs to gain more revenue, and so on. To me, it would be perfectly sensible to have a fiscal weighted-average measure of social opportunity cost if, say, the value added tax was the only or principal tax, and if changes in the fiscal situation of the governments were met by changing its rates. In that world, I, too, would be a proponent of this position.

But the world we observe shows a whole panoply of fiscal adjust-

10. Otto Eckstein, "Investment Criteria for Economic Development and the Theory of Inter-Temporal Welfare Economics," *Quarterly Journal of Economics* (February 1957); Otto Eckstein and John V. Krutilla, *Multiple Purpose River Development* (Baltimore, Md.: Johns Hopkins University Press, 1958); and Robert Haveman, *The Economics of the Public Sector* (New York: Wiley, 1972).

ments with little or no predictability about what the next one will be like. I find it impossible to conjure up even a semirealistic mental experiment in which a specified fraction of fiscal revenues typically comes at the expense of consumption, with its complement coming at the expense of investment. In contrast, it is natural for these fractions to be relatively stable for dollars drawn from the capital market. Because of the relative stability of the weighting structure, this argues in favor of the convention that the marginal source is the capital market.

Of course, a stable weighting structure would not mean much if the idea of the capital market as the marginal source of funds did not make sense. But in fact, on a day-to-day, month-to-month, and even year-to-year basis, the capital market is in most countries the marginal source of funds. Most government budgets, even on the day they are first presented, contemplate the borrowing of some funds. In addition, as actual events produce deficits greater than planned, governments almost always turn to the capital market for the difference. When the future smiles and deficits are smaller than expected, the extra money is in effect returned to the capital market. In sum, the capital market is the marginal source when funds are short and the depository for marginal funds when they are abundant.

A third advantage of adopting the convention of treating the capital market as the source of funds at the margin is that one can readily adapt it to incorporate capital funds from abroad. The weighted average would now be  $f_1 p + f_2 r + f_3 MC_f$ , where  $f_1 + f_2 + f_3 = 1$  and where  $MC_f$  is the estimated marginal cost of foreign funds.<sup>11</sup> One should employ here the marginal cost of foreign funds because of the presumption that the supply curve of such funds is not infinitely elastic. With an upward sloping curve, the marginal cost of funds will exceed the average cost,<sup>12</sup> a fact which obviously should be taken into account in calculating the social opportunity cost of capital.

Thus the arguments for building the calculation of the social opportunity cost of capital on a capital-market sourcing model are that (1) the weights are relatively stable, (2) the capital market is the de facto marginal source and destination of funds in the short and middle run, and (3) the calculation can be readily adapted to incorporate sourcing from the world capital market.

Having elected to hold with a capital-market sourcing convention, let us now consider the relative merits of using the social opportunity cost

11. Arnold C. Harberger, "Vignettes in the World Capital Market," *American Economic Review* (May 1980), pp. 331-38.

12. The textbook formula is  $MC = AC[1 + (1/s)]$ , where  $AC$  is average cost and  $s$  is the price elasticity of the average cost curve—in this case the upward rising supply curve (of foreign funds) facing the country in question.

of capital ( $f_1p + f_2r$ ), on the one hand, or the marginal rate of time preference ( $r$ ), on the other, as the discount rate in social cost-benefit analysis.

The best starting point is to realize that modern defenders of the use of  $r$  as the rate of discount (Feldstein; Dasgupta, Marglin, and Sen; and Squire and Van der Tak) do not neglect the existence of the distortions that proponents of the weighted-average approach (Baumol, Bruce, Dreze, Sandmo, Sjaastad and Wisecarver) take account of in the discount rate.<sup>13</sup> Instead of being reflected in the discount rate, these same distortions show up in the "shadow price of investable funds" in the case where  $r$  is used for discounting. The procedure is as follows. Assume that one dollar is extracted from the capital market. A fraction  $f_1$  of this is from displaced investment, which in turn would have generated future income at the rate of  $p$  a year. The complementary fraction  $f_2$  comes from newly stimulated saving, the supply price of which is  $r$  a year. There is therefore an annual opportunity cost of  $f_1p + f_2r$  for each year in the future, as a consequence of withdrawing one dollar this year from the market. This future flow of opportunity costs, discounted back to the present at the rate  $r$ , has a present value of  $(f_1p + f_2r)/r$ . This is the shadow price of investable funds for those who use  $r$  as the discount rate.

If  $p = 12$  percent and  $r = 4$  percent, with  $f_1 = 0.75$  and  $f_2 = 0.25$ , we would have a weighted-average opportunity cost of capital of  $(0.75)(12 \text{ percent}) + (0.25)(4 \text{ percent}) = 10$  percent. This is what the weighted-average advocates would use for discounting. In doing so, their investment criterion would be that the discounted value of benefits should exceed the discounted value of costs.

In contrast, the advocates of discounting by the marginal rate of time preference ( $r$ ) would use 4 percent as the discount rate, but would

13. Martin Feldstein, "The Inadequacy of Weighted Discount Rates," in R. Layard, ed., *Cost-Benefit Analysis* (Harmondsworth, Eng.: Penguin, 1972); Partha Dasgupta, Stephen A. Marglin, and Amartya Sen, *Guidelines for Project Evaluation* (New York: United Nations Industrial Development Organization, 1972); Lyn Squire, and Herman Van der Tak, *Economic Analysis of Projects* (Baltimore, Md.: Johns Hopkins University Press, 1972); William Baumol, "On the Social Rate of Discount," *American Economic Review* (September 1968), pp. 788-802; Neil Bruce, "On Weighted Average Measures of the Social Opportunity Cost of Capital," UCLA Discussion Paper, University of California, Los Angeles, 1985; Agnar Sandmo, "Optimal Taxation—An Introduction to the Literature," *Journal of Public Economics*, vol. 6 (January 1976), pp. 37-54; Agnar Sandmo and Jacques Dreze, "Discount Rates for Public Investment in Closed and Open Economies," *Economic Journal*, n.s. (November 1971); and Larry A. Sjaastad and Daniel Wisecarver, "The Social Cost of Public Finance," *Journal of Political Economy*, vol. 85 (June 1977), pp. 513-47.

require that the present value of benefits be more than 2.5 times the capital costs of the project. The factor 2.5 is exactly  $(f_1p + f_2r)/r$ ; discounted at 4 percent, this is the present value of what is given up for every dollar withdrawn from the capital market.

If the two sets of criteria (10 percent discount rate requiring that benefits exceed costs, and 4 percent discount rate requiring that benefits exceed 2.5 times costs) were properly implemented, I do not believe there would be many serious contradictions in their implications for longer-term investment projects. However, I have never had any doubt about preferring the first approach. Three grounds for this preference are: communicability of the procedure, implications with respect to current expenditures, and implications for handling situations with different rates of time preference for different groups.

With respect to the first, I have always felt that the most basic function of project evaluation was to shoot down the worst projects. Unfortunately, most of the worst projects have strong supporters, usually within and outside of government. To my mind, the project evaluation team weakens its position if it adopts a criterion that requires benefits to be, say, 2.5 times costs.<sup>14</sup> It is hard to beat down a project with the argument that its benefits are *only* twice its costs, and that this is not enough!

With respect to the second point, I have from the beginning been greatly impressed by a powerful argument raised by Sjaastad and Wisecarver in a paper that merits the most serious attention.<sup>15</sup> Once the capital market is accepted as the marginal source of funds, it should be recognized as the marginal source of funds not only for capital outlays, but also for current spending. Most particularly, any money saved through greater efficiency in the police force or in the schools can be used to pay off debt; used in this way it will (in our numerical example) produce a benefit equal to 2.5 times its nominal amount. The logic of the case presses one ineluctably to the conclusion that for current outlays as well, the critical ratio of benefits to costs should be 2.5. In short, the factor  $(f_1p + f_2r)/r$  represents the social value of one dollar of liquid funds, either taken from or placed in the capital market. Any

14. To be sure, within the context of the methodology, the factor 2.5 is necessary in order adequately to reflect the present value of the future costs entailed in borrowing. Once this factor is used as a shadow price to multiply capital costs, the correct rule is that benefits should simply exceed costs, so adjusted. But non-professionals will still wonder about projects being rejected simply because the present value of their benefits is only 1.5 or 2 times actual capital outlays, and powerful forces behind bad projects would surely use such circumstances to convince the public and the relevant authorities that the projects were in fact quite good.

15. "The Social Cost of Public Finance."

funds sourced in the capital market will have a shadow price of 2.5 per dollar, in our numerical example. And since the capital market is being taken as the marginal source of funds (as in truth it really is, in most cases), the 2.5 factor should apply to all cash costs across the board. Cutting out any cash outlays—either current or capital—would likewise permit achievement of benefits at the ratio of 2.5 to 1. A powerful and persuasive argument, to which I have yet to see a convincing rebuttal!

The third reason for preferring the weighted-average discount rate to the marginal rate of time preference is my own (in the sense of my not having encountered it in other writings or discussions). To my mind, the weighted-average discount rate moves quite naturally from the simple  $f_1 p + f_2 r$  to the more disaggregated  $\sum_j f_j p_j + \sum_k f_k r_k$ . In each case the  $p_j$  are the demand prices of displaced investments and the  $r_k$  are the supply prices of newly stimulated savings. A project undertaken today must be able to repay the cost of its assets and to cover these demand and supply prices; otherwise it has not paid its way. There is nothing in the weighted-average approach that insists there be only one  $r$ , and the presence of any number of different  $r_k$  presents no conceptual difficulty or embarrassment.

Contrast with this the use of  $r$  as the discount rate. Its justification stems from the standard treatment of intertemporal consumption decisions in economic theory—the rate of return to the saver, after all taxes, is the rate that he uses in deciding upon his savings. The weighted-average approach builds on this also, but the time preference advocates turn this into the cornerstone of their discounting procedure, on the ground that it is by using  $r$ , not  $p$  or  $i$ , that consumers make their intertemporal choices. Here is where I encounter difficulties in passing from a case in which there is a single  $r$  (a single uniform tax rate separating  $i$  and  $r$  for all consumers) to one in which there are two or more  $r_k$  for different groups (for example, two or more marginal tax rates). If  $r$  is used as the discount rate because it is the fundamental guidepost for consumers as they make their intertemporal consumption decisions, then when different consumer groups use different rates, it would seem that each group's benefits should be discounted at its own rate.

To my knowledge none of the advocates of using  $r$  as the discount rate has broached this problem—I merely say it is a problem that flows naturally out of the conceptual basis of the time preference approach. The underlying problem is that when there are different rates of discount for different groups, postulate 3 needs to be more tightly specified. Suppose groups A and B have different  $r_k$ 's, and suppose there are two projects (say, parks), each of which gives to one of the groups a certain path of in-kind benefits through time. Let these benefit paths be

identical for the two groups. If we simply subtract one benefit profile from the other, year by year, the difference is zero each year. But if we discount future benefits back to the present, the group with the lower discount rate (presumably the wealthier of the two) will have a higher present value of benefits. And if we accumulate the benefits forward to some future time, measured as of that point in time the group with the higher discount rate will have the higher present value of benefits. So by subtracting the benefits of A from the benefits of B year by year we perceive no difference; by discounting to the present we find group A with greater benefits; by accumulating to the future we find group B with greater benefits.<sup>16</sup> All this runs counter to the standard proposition of capital theory and project evaluation that the ranking of projects should not depend on the point in time to which benefits and costs are discounted or accumulated; it matters only that the projects being compared be evaluated as of the *same* point in time. This proposition is always true if a single rate is used for discounting and accumulating—even when that rate varies from one year to the next. But it obviously is not true with different rates applying to different groups of economic agents.

Actually, those who propose using time preference discount rates always work with a single rate, which avoids the ambiguity just referred to. I concur and suggest that in social project evaluation postulate 3 should apply within the year, canceling costs against benefits as they occur, and that a single discount rate (possibly varying from year to year) should be used to carry net benefits or costs from one time period to the next. To my mind, the weighted-average social opportunity cost of capital ( $\sum_j f_j p_j + \sum_k f_k r_k$ ) is ideally suited to this role. For the reasons already expounded, a time preference rate is not.<sup>17</sup>

### Concepts versus Numbers

Practitioners whose main experience is in the field, together with others who function daily in the world of affairs, may wonder at the

16. There seems to be no good reason to treat costs differently in the two cases, as they are supposed to follow identical time paths.

17. I am assuming that there are different groups with different values of  $r_k$ . The weighted-average social opportunity cost has a precise meaning—the amount needed to compensate all losers when money is drawn from the capital market. Using as the discount rate a weighted average of just the  $r_k$ 's will do the trick in a mechanical sense—that is, it will eliminate the ambiguity of working with separate  $r_k$ 's—but no clear concept lies behind such a weighted average covering only marginal rates of time preference.

level of precision that has characterized most of the discussion so far: demand prices, supply prices, distortions, weights based on elasticities that are impossible to observe directly and virtually impossible to estimate exhaustively (that is, it is impossible to estimate all the items necessary for the construction of a system of elasticity-based weights). We have weights that depend on saving and investment behavior, for which econometric explanation is still a matter of controversy even in advanced countries. We have weights that depend on the responsiveness of individual classes of exports and imports to the tightening and easing of general pressure on the exchange market, items that to my knowledge have not been estimated at all.<sup>18</sup>

What then do people actually do? In point of fact, most of the time they use extremely crude estimates based on readily available data. The simplest measure of the social opportunity cost of foreign exchange is the market exchange rate augmented by the so-called force of tariff—the ratio of actual tariff receipts to the c.i.f. value of imports. This measure (1) neglects the export side altogether, (2) makes no attempt to distinguish between import categories that are especially sensitive and those that are especially insensitive to changes in the real exchange rate, and (3) completely ignores the possibility that important distortions (such as petroleum and gasoline taxes) that do not fall directly on trade may nonetheless have a measurable impact on the social opportunity cost of foreign exchange.

The great thing about having a conceptual framework is not only that it helps us to think through problems in a clear way. It also tells us how to try to improve on estimates and measures that are extremely crude and approximate. Starting with a force-of-tariff first approximation, the steps to be taken are those suggested by points 1–3 above. First, we must ask whether exports belong in the weighted average. If the exchange policy is one of simply doling out by licenses whatever foreign exchange comes in, it may well be that no export adjustment should be

18. In deriving the relevant weights in the foreign exchange case, the origin of the pressure is conceived of as additional purchases of foreign exchange in the market. These purchases raise the real exchange rate, either because the nominal rate is flexible or because of the natural workings of the adjustment process under a fixed exchange rate. If the real exchange rate is denominated by  $e$ , the derivatives that are relevant for constructing the weights  $f_i$  and  $f_x$  are  $\partial M_i/\partial e$  and  $\partial X_i/\partial e$ . For example,  $f_i$  would be equal to  $-(\partial M_i/\partial e)/(-\sum_j(\partial M_j/\partial e) + \sum_k(\partial X_k/\partial e))$ , and  $f_x$  would be  $\partial X_x/\partial e$  divided by the same denominator. ( $\partial M_i/\partial e$  is typically negative when  $e$  is defined, as it is here, as the local-currency price of, say, the dollar.) In short, we are asking how each item in the foreign exchange market responds to the general pressure of a tightening of the real exchange rate. This is not the same thing as how each item responds to a change in its own price or to an increase in pressure in just its own market.

made. But if the real exchange rate is determined to a substantial degree by market forces, then one should try to infer how elastic is the response to exchange rate pressure of different categories of exports, and to determine what distortions (such as taxes and subsidies) are relevant for each major category. At the very least, a broad "average distortion" for all exports taken together (for example, the ratio of export taxes or subsidies to their total f.o.b. value) can be estimated and introduced along with the force of tariff to produce a weighted-average estimate (including exports) of the social opportunity cost of foreign exchange.

Second, to break down imports into categories that have different relative sensitivities to changes in the real exchange rate, the practical procedure is to start with simply a value-weighted average of all tariffs. (If exports are not in the picture this actually yields the force-of-tariff ratio.) Then one classifies the categories into those that are judged to be, say, very sensitive, somewhat sensitive, average, somewhat insensitive, or very insensitive to changes in the real exchange rate. Having done so, one adjusts the weights accordingly—doubling or tripling, say, the weight attached to the very sensitive group and perhaps cutting by half or two-thirds the weight attached to the most insensitive group. Any category subject to a quota that is unlikely to change but likely to be continually effective would receive a zero weight, for these imports, even though they may be quite important, are not displaced when the real exchange rate tightens nor do they expand when it eases. A similar treatment can then be applied to exports. Obviously, throughout the process of adjusting weights, care must be taken not to alter their sum.

The third step of adjusting the social opportunity cost of foreign exchange for internal distortions is relatively easy. If, say, the country produces no petroleum, then the petroleum and gasoline taxes imposed internally function in the same way as tariff surcharges—one simply combines them with the tariff to get a picture of the total government revenue generated by each dollar of petroleum imports. In case there is domestic production of a good (such as alcoholic beverages) that is subject to excise tax, an increase in the real exchange rate will give rise to some substitution toward the domestic product, hence the government's loss in excise taxes will be some fraction of what one would predict on the basis of the reduction in imports.

The basic conceptual framework also guides us as we struggle to come up with empirical estimates of the social opportunity cost of capital. The fact that the concepts are framed in real terms suggests beginning with direct rather than indirect measures of the real return to capital. Rates of total return (including real capital gains) in financial markets (especially stock markets) are notoriously volatile and subject to the vagaries of transitory swings in expectations. Much more steady and reliable are measures of the national or sectoral rate of return,

based on estimates of the real capital stock (usually built up via a perpetual-inventory process) together with data on the real returns accruing each year to the capital factor. Considerations of differential taxation suggest distinguishing where possible between predominantly corporate and predominantly noncorporate sectors, and singling out those with special tax treatment such as housing, mining (at times), and agriculture (quite frequently).

The underlying concept is one in which the government draws its marginal funds from the capital market. In most countries, this probably means that there is no systematic mechanism by which one government project derives part of its funds by displacing other investments throughout the public sector. One project may indeed displace another—even in its entirety—but such displacement is likely to be sui generis for each particular project, not systematic and similar for all classes of public projects.

Consequently, when we try to estimate  $p$ , the marginal productivity of displaced investment, we probably should try to measure what I call the social rate of return to private sector capital. The case where government project A in fact displaces government project B can then be looked upon as a combination of two modules: one in which project A is undertaken and is financed by resort to the capital market, the other in which project B is canceled and its funds returned to the capital market.

Some countries have only rudimentary capital markets, often pretty much limited to the banking system itself. In such cases, government resort to the capital market usually means that a roughly equivalent amount of credit is crowded out of the private sector. Therefore the first approximation of the marginal productivity of displaced investment should be the real social rate of return to capital by those segments of the private sector that are the typical recipients of bank credit. Where monetary policy does not permit government borrowing from the banking system to crowd out enough investment, the typical result is that part or all of the government's borrowing gets reflected in inflation. It is then likely that a significant part of the government's outlays will in the end come at the expense of consumption. Inflation processes are exceedingly complex, however, with institutional arrangements (such as indexing) varying widely among countries. In an inflationary environment, therefore, the weights such as  $\beta$  and  $\gamma/\alpha$  should probably be treated as a problem to be solved separately for each country.

With respect to the labor market—perhaps the most vexing case of all—the underlying concepts are essential guideposts to those who undertake the difficult task of empirical estimation. From the very outset, the concepts tell us to seek the ultimate pattern of sourcing from

which an increment to labor demand is filled—and not to take too seriously the provenance of the particular workers that are hired by a project. They may be lured away from other employers; they may quit a job because the project is near their home or otherwise appealing; or they may be picked up by the project out of a temporary spell of unemployment. This direct sourcing pattern may be next to meaningless: if they left another job, the likelihood is they will be replaced; if they were unemployed, they would probably soon have found some other job.

At the other extreme, the concepts warn us away from a demographic-historical approach that considers the ultimate sources of new labor supply to be natural increase, migration into the region, and changes in the rate of labor force participation. The problem with this approach is that most of such changes would have taken place regardless of the presence or the absence of any particular project. What we are seeking is the chain of causation by which the presence of a particular project draws labor (directly and indirectly) from where it would have been (in equilibrium) in the absence of the project. In short, we must think about sourcing in economic rather than demographic-historical terms.

To think in economic terms means to think about markets—in this case the market for the particular type and class of labor involved. One must determine whether the market is a national or a regional one and, if the latter, to what degree increments of demand are likely to be met via migration. One must try to identify the structure of the market and the characteristics of its supply and demand. For highly specialized and highly skilled jobs, the likelihood is that the bulk of any new demand will be met by bidding workers away from other employers. For relatively unskilled, undifferentiated tasks the existing pool of those who hold such jobs has little significance—new taxi drivers, hotel maids, or grocery clerks can come from anywhere.

By fortunate coincidence, dealing with the social opportunity cost of labor is simplified by the fact that most labor market distortions (such as payroll taxes and income taxes) are relatively modest in size (at least compared with many tariffs and with some taxes on the income from capital) and also quite widespread through the labor force. Perhaps the biggest distortions in developing countries are those between the most highly protected sectors (typically multinationals) and the rest of the economy. But this creates relatively little difficulty, because we can be quite sure that the other sectors do not obtain any net labor from the most protected ones. If a worker leaves a highly protected sector, there is always a long list of candidates waiting to replace him. The only time the highly protected sector's wage enters the calculation of social

opportunity cost is when the new jobs that are being created are actually in that sector. In such a case the protected-sector wage enters as the upper limit to the supply price of the quasi-voluntary unemployed.

Similarly, although the problem of determining empirically the social opportunity cost of the unemployed in a cyclical setting poses challenging conundrums to the analyst, it is of relatively minor importance in real-world project evaluations.<sup>19</sup> The reason is that project evaluations are most of the time forward-looking operations. The project being evaluated may not even be initiated until two or three years later, and its economic life may then stretch for ten or twenty or even fifty years. Whether or not cyclical unemployment is currently observed, the best prediction for future years is that the situation will be "normal" or "average." As far as I can see, the issue of the social opportunity cost of labor in a setting of cyclical unemployment comes up largely with respect to programs (perhaps of temporary job creation) specifically designed to deal with the unemployment situation. The context is one in which there will be a relatively rapid absorption of the unemployed, with a duration roughly commensurate to the recession being experienced. The typical longer-term investment project simply does not fit into this context.

With respect to chronic unemployment, two varieties have been discussed: quasi-voluntary unemployment linked to the presence of protected sectors, and migration-fed unemployment. Neither of these, in my view, presents conceptual or measurement problems of a serious nature.

### Distributional Weights and Basic Needs

At the outset of this section I want to reemphasize that all three basic postulates form the roots of the grand tradition of applied welfare economics; all three are needed to produce the series of classic results recounted at the beginning of this paper; all three have been employed by a long lineage of great economists. Nonetheless, it is the third postulate that has most often been questioned and that has raised

19. The principal problem here is to determine what supply price of labor to use. Labor supply, particularly of adult males, is known to be extremely inelastic in most countries, yet in repeated surveys the unemployed state their supply price to be close to the going wage for their age-education-experience category in their own labor market; that is, the stated supply price takes on the attributes of the expected demand price in the market. The question is, to what extent should we accept such statements of supply price at face value, thus attributing little gain in utility to the unemployed person when he or she finds a job?

doubts in the minds of some thoughtful people—economists and non-economists alike. The purpose of this section is to address such doubts and to discuss some of the issues surrounding them.

Let me begin at the end. The great economists who have employed the three postulates have been neither naive nor crass. From their writings one can readily discern that they do not believe that the postulates and their implications are all that count. Each society has its own values, and each has many important objectives apart from economic efficiency. No one, to my knowledge, has argued that the three postulates should overthrow society's values or should supplant important noneconomic objectives. But the three postulates do provide a disciplined and coherent framework for thinking through the economic aspects of a wide range of problems. The appropriate analogy is with accounting—another example of a complex structure of analysis and consequences that rests ultimately on a few fundamental postulates. Accounting tells how to calculate net profit and net worth (wealth). It does not attempt to do the absurd by saying that profits and wealth should never be compromised or sacrificed for other objectives. Accountants know and recognize this. So do economists when it comes to considering the consequences and implications of the three basic postulates in comparison with other objectives.

One must be careful here, however. Traditional accounting, though far from revealing the meaning of life, is important enough in many different contexts to be taken seriously. And it is part of the role of accountants to call the attention of others to the relevance and usefulness of their discipline. So, it seems to me, it is (or should be) with applied welfare economics. It is significant and relevant for many matters of interest to society. It gives a disciplined and structured response to important questions. It embodies some of society's values, but not all. And in many (even most) cases its results are not incompatible with society's other values. In my view, part of the role of economists is to present the answers given by applied welfare economics to the problems at hand, as the particular contribution of the economics profession to society's decisionmaking process. If economists do not present what economics as a discipline has to contribute, who else will?

The approach just described is modest, in the sense that it makes no excessive claims for the postulates and does not try to graft onto them other objectives that might expand their scope. A significant strand of literature, however, follows a different line, specifically with respect to the distribution of benefits and costs. This literature explores the use, within applied welfare economics, of distributional weights. Use of such weights entails multiplying the net gains or net losses of particular groups by specific factors—higher for groups whose welfare is deemed more meritorious, lower for those whose welfare is less prized (presum-



ably by society as a whole). In principle, welfare weights could be classified according to any criteria—ethnic or national origin (Malays and Chinese in Malaysia), educational or class background (untouchables in India), type of economic activity (independent farmers and farm laborers in many countries), and so on. But most of the literature has focused on income or wealth as the criterion to be discussed and has typically assumed that the welfare weights are a declining function of the criterion variable.

Such declining weights have been rigorously applied by several authors (Atkinson, Atkinson and Stiglitz, Diamond, Dixit and Sandmo, and Mirrlees), particularly to deal with the problem of the optimal choice of an income tax schedule.<sup>20</sup> A few attempts have been made to incorporate distributional weights in actual project evaluations. Most of these have been flawed by the application of the weights to the changes in *income* of the various groups, rather than to the changes in surplus (net benefit or cost). All the logic of distributional weights leads to the use of change in consumer and producer surplus as the base to which the distributional weight should be applied. An increase in labor along an infinitely elastic supply curve, for example, implies no increase in welfare to the suppliers; hence there is no relevant base to which to apply the group's distributional weight.

After giving considerable thought to the problem, over a period approaching a decade, I have come to the conclusion that systems of distributional weights do not adequately represent the way citizens feel about redistribution efforts channeled through the public sector. Instead, I feel that a system of "basic needs externalities" reflects people's values and beliefs much better than distributional weights.<sup>21</sup> Basic needs externalities are based on the idea that citizens and taxpayers look for

20. Anthony B. Atkinson, "On the Measurement of Inequality," *Journal of Economic Theory*, vol. 2 (September 1970), pp. 244-63; Anthony B. Atkinson and Joseph E. Stiglitz, "The Design of Tax Structure: Direct versus Indirect Taxation," *Journal of Public Economics*, vol. 6 (1976), pp. 55-75, and *Lectures on Public Economics* (New York: McGraw-Hill, 1980); Fear A. Diamond and J. A. Mirrlees, "Optimal Taxation and Public Production: Tax Rules," *American Economic Review*, vol. 61 (March and June 1970), pp. 8-27 and 261-78; A. K. Dixit and Agnar Sandmo, "Some Simplified Formulas for Optimal Income Taxation," *Scandinavian Journal of Economics*, vol. 79 (1977), pp. 413-23; and J. A. Mirrlees, "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, vol. 38 (April 1971), pp. 175-208, and "The Theory of Optimal Taxation," in K. J. Arrow and M. D. Intriligator, eds., *Handbook of Mathematical Economics* (Amsterdam: North-Holland, 1979).

21. Arnold C. Harberger, "Basic Needs versus Distributional Weights in Social Cost-Benefit Analysis," *Economic Development and Cultural Change*, vol. 32 (April 1984), pp. 455-74.

specific and concrete results when public funds are channeled into helping others. They want to see the recipients turn out to be better educated, better cared for medically, better fed, and better housed. In short, citizens and taxpayers are not interested in having their money used simply to gratify the recipients; they want to see it used to advance the welfare of the recipients as they (the donors) perceive that welfare. This helps explain why the most universally popular government transfer programs have entailed transfers in kind rather than in cash, as in the case of universal free primary education and free or heavily subsidized medical care for the indigent. It also explains why public support has been weak for programs in which transfers in kind can be converted into cash with relative ease. Subsidized food, or the food stamps with which to acquire it, can be resold; subsidized housing, with space judged to be "adequate," may quickly become overcrowded by an inundation of "cousins." In such cases the recipients are maximizing their own welfare by transforming their in-kind transfer into cash, but the donors are unhappy because their purpose (improving the welfare of recipients as the donors perceive it) has not been served.

This and much other behavior, as well as many other opinions and attitudes on the part of citizens and taxpayers, can be rationalized in terms of basic needs externalities. Citizens, the argument goes, are concerned enough to have their money spent to ensure better education, medical care, nutrition, and housing for the less fortunate members of society. But donors are not interested in seeing their money go to finance more elegant dowries for brides, more elaborate funeral ceremonies for grandparents, or more frequent return trips of workers to their native villages. Even if these are the things the recipients prefer, they give rise to no external benefit as perceived by the donor, hence donors (taxpayers) do not want to see their money used directly or filtered indirectly to pay for such things.

Elsewhere I have suggested ways of getting crude quantifications of the basic needs externality for such things as education, medical care, and nutrition.<sup>22</sup> These quantifications are justified by the existence of a basic needs externality, but they are also influenced by a modified version of the least cost principle. The least cost principle tells us not to accept any project (or other action) if a cheaper way of achieving the same or equivalent benefits can be found. This dictum needs to be modified, at least for basic needs externalities, because in most countries there is an almost infinite number of situations in which equivalent benefits from basic needs externalities might be found. The modified least cost principle, then, states that the costs society is willing to incur

in order to meet a basic need should vary inversely with the intensity or urgency of that need. For example, we would typically be willing to incur a greater cost to bring a given group from 85 to 90 percent of some nutritional or health standard, than to bring a similar group from 95 to 100 percent of that standard. Similarly, poor countries can attempt to subsidize only the earlier stages of education for broad groups of the population; the higher the level, the more selective the criteria must be.

In line with this modification principle, the project authority could set, for each basic need, a cutoff level above which no basic needs externality is deemed to exist. And since the externality in the end turns out to be a justification for society to pay some or all of the bill, it is appropriate that the income or standard of living of the recipient households should enter in an important way in defining the cutoff level. For most developing countries, for example, it does not make sense for society to subsidize the diets of the top 50 or 60 percent of the population. Cutoff levels for the attribution of any basic needs externality would presumably vary with the type of need. They might be set at the 75th percentile of the income distribution for primary education, at the 40th for secondary education, and at the 25th for higher education. For basic medical care, they might be at the 60th percentile, while for housing only at the 20th. For society to attribute a basic needs externality, in the terms in which we are speaking, it must be willing and able to bear part or all of the cost.

In addition to establishing a cutoff point for each externality, the modified least cost principle would define a maximum externality (as a percentage of the normal cost of providing the service in question) that society is willing to attribute, even in quite urgent cases. The reason behind this is similar to that which motivated the original version of the least cost principle. If we attribute a maximum externality of 100 percent of normal costs in cases of very low nutrition levels, we can be quite sure that, as a result, society will pay double the normal cost at least some of the time in order to meet the need in question. The logic of cost-benefit analysis is simple and ineluctable. If the normal cost of meeting a given (say, nutritional) need is 10 per unit, and if we attribute a 100 percent externality to meeting that basic need, there are two plausible extreme cases in which that externality would cause the acceptance of a project that would otherwise be rejected. At one extreme is the case in which there are no benefits of the usual type—the externality is the only benefit. In such a case the project would be accepted if its costs were less than or equal to the usual 10 per unit. At the other extreme is the case in which the usual type of benefit is equal to 10 per unit; here the existence of the externality will render a project acceptable if costs range up to 20 per unit. In both cases society ends up

paying a cost of up to 10 per unit as the "price" for meeting the basic need in question.

Could not cheaper ways of meeting this or an equivalent basic need be found? Almost certainly so, and the surest way to bring about a search is to disallow the attribution of externalities as large as 100 percent in the first place. When we attribute a given externality as a benefit, we say in effect that it is worthwhile to incur a corresponding amount of costs in order to produce that benefit. Without a systematic process of searching for alternative, cheaper ways of obtaining given amounts of benefits, the attribution of an externality is an invitation to incur up to that amount of extra costs so as to achieve the external effect.

The modified least cost principle embodies the idea that society is willing to accept the risk of inefficiency in the most urgent cases, such as famine, starvation, and pestilence. But the willingness to accept this risk declines as the basic need being met becomes less urgent, until—at the cutoff point—society simply refuses to attribute any externality, thereby refusing also to invite the conscious acceptance of extra or higher costs.

The practical implementation of the schema just described can be made very simple. The budget bureau, the finance or planning ministry, the project evaluation office, or even the national cabinet would decide that for, say, the lowest percentile of households a basic needs externality of 30 percent of the normal cost of additional nutrients would be assigned. This percentage would decline to zero at the 40th percentile of households (the cutoff point). Thirty percent, in this case the maximum allowable externality, is also the maximum amount of excess cost that the procedure would permit to be incurred. Society would be accepting some inefficiency (by standard criteria) for meeting basic needs. But this acceptance would be tempered by the placing of explicit and conscious limits on the extra costs to be incurred on this account.

Although the use of distributional weights by definition entails the rejection of postulate 3, the attribution of basic needs externalities is fully compatible with it. The positive externality involved in improving the education, health, nutrition, and housing of the disadvantaged takes its place alongside other externalities such as air and water pollution and traffic congestion. All of these, like basic needs externalities, have the attributes of public goods (or public "bads"), but all can be easily fitted as positive or negative distortions into the framework of  $\Sigma D_i \Delta X_i$ .

Once this is recognized, it helps solve what many people consider a troublesome conundrum. As shown above, when the relevant group has an infinitely elastic supply of labor, no distributional weight benefit

can be attributed, even when the employment of low-income workers expands dramatically at the given supply price. This bothers many people, because they firmly believe there should be a social gain. The answer is clear; the people who are troubled by the no-benefit result do not think in terms of distributional weights. Their intuition, it seems to me, runs more along the lines of basic needs externalities. A rise in employment—even along an infinitely elastic curve, and either through increases in numbers employed or in hours worked per person—will increase the cash income of the affected families. These added funds, in turn, will almost certainly be spent in meeting at least some additional basic needs. When we value these increments of welfare as externalities using  $\Sigma_i D_i \Delta X_i$ , we find that the added employment did indeed bring a social benefit. The basic needs approach thus solves the conundrum.

### Epilogue

I have tried in the preceding pages to outline an approach to social project evaluation that is conceptually sound, simple, and deeply rooted in the grand tradition of economic science. It is presented as what economic science has to say or to offer in the area of cost-benefit analysis as society struggles with the many difficult choices it faces.

In the process I have recognized several times that it is an austere analytical structure that emerges when we build upon the three basic postulates. No more than the accountant's rules does that structure reveal the fundamental purpose and meaning of life or the ultimate values that a particular society ought to treasure or to seek.

These things are the source of vast amounts of misunderstanding and confusion. Neither economics as a science in general nor the three basic postulates in particular tell us that individuals and societies should seek to maximize their incomes. Economics postulates that individuals and families seek to maximize their own welfare (as they perceive it), including the various ways—such as basic needs externalities—in which the welfare of others impinges on their own. This welfare-seeking behavior by each individual economic unit forms the foundation stone of the three basic postulates. These, in turn, give rise to an analytical structure that extends to society as a whole a procedure for weighing social benefits against costs that is compatible with welfare-seeking at the level of the individual unit.

The resulting structure is austere. It does not answer all questions. It does not incorporate all of society's values. It does, however, incorporate the value of economic efficiency in the sense of trying to maximize perceived welfare—not income. One very important question for doubters to ask themselves is, how many of society's other goals and

values are really antagonistic to economic efficiency? I, and the whole tradition of economic science, see relatively little problem of incompatibility. But some incompatibility there may nonetheless be, as when a society chooses to favor certain groups (such as disadvantaged minorities) in ways that may not fit neatly as basic needs externalities under the three postulates.

Even in such cases economic science can help us toward an answer. For at the very core of economics is the notion that when we are striving after two or more good ends, but reality tells us that we cannot have more of all of them, we should be prepared to sacrifice something of one in order to get more of another. This is the most fundamental of economic principles. It tells us that in the presence of multiple social goals we should first strive to seek compatibility between those goals and economic efficiency; but if there remains some incompatibility, we should be prepared to make at least a minimal sacrifice of economic goals to gain advances toward other important ends.

But I still remain a defender and champion of the three postulates and of the traditions of economic science that they represent. So far, I have mainly used the same terms as the philosophical doubters. In their discussions, the other values that vie with economic efficiency are somehow almost always deemed to be good values, with high ethical and moral connotations. Almost always they are values that carry us to a world that is somehow better than the one we reach using the three basic postulates.

But what I see, and what I believe was seen by most representatives of the great tradition of applied welfare economics in the past, is a world in which good and evil exist side by side, and pervasively. The noneconomic goals pursued in the real world are not all based on the high principles of ethics and philosophy. Many, far too many, reflect the darker side of human nature.

One important noneconomic goal in most countries—at the very least it is a goal to which the three postulates have been able to contribute very little—is national defense. This is a perfect example of what I mean, for under the label of national defense we have seen things that run the full gamut from the most honorable to the most vile. Those who struggle nobly against aggression and oppression argue their cause in terms of the value of national defense. Those who perpetrate these very acts of aggression argue their own cause using the same vocabulary.

Almost by its very nature social project evaluation is enmeshed in the structure and processes of government. Most often it is public sector projects that are analyzed—the analysis being done by one public agency, the approval of another being required, and the implementation being done by a third.

When we speak, then, of noneconomic goals let us not forget the many projects that have been carried out just to satisfy the caprice or whim of some powerful figure or clan. Let us not forget the corruption that pervades the decisionmaking and contracting process in many parts of the world. And even when these elements are not present, let us not forget how project choice gets intertwined with the political process almost everywhere: how the granting and withholding of projects is used to reward political supporters and to punish enemies, and how in electoral situations governments tend to distribute projects with the aim of winning over constituencies that may be doubtful or wavering.

All kinds of objectives compete with the three postulates, and I have no doubt that, in the final analysis, the less-than-noble motives enter with greater frequency and importance than the lofty ones. Viewed in this light, the three basic postulates provide a way of insulating the methodology of social project evaluation from the banal, crass, even vile pressures just alluded to. They lead to the sort of professionalism discussed in the introduction to this essay. Just like the principles of accounting, the methodology based on the three postulates enables one group of project evaluators to review or "audit" the world of another. The methodology itself should always lead to the same answers; the only serious problems lie in estimating and quantifying future costs and benefits. This aspect is what makes it possible to think of a "professionalized" discipline of social project evaluation in spite of the many human failings that characterize the environment in which project decisions are made and executed. I believe, too, that the development of a rigorous professionalized discipline provides the best hope for improving actual performance and gradually overcoming the institutional weaknesses and personal temptations that have been important sources of the blemishes on record from the past.