

INTRODUCTION TO COST-BENEFIT ANALYSIS

PART I

PROFILES, FOREIGN EXCHANGE AND CAPITAL

Arnold C. Harberger

University of California, Los Angeles

November 2008

This paper is not intended to give its readers the knowledge and skills needed to carry out economic evaluations of real-world projects. That task would require a major study program, ideally extending over several months or more of hard, full-time work, comprising classroom hours, extensive readings, and many exercises in which participants deal with a whole gamut of issues in a sequence of simulated real-world cases.

Instead, we try here to convey an understanding of what is involved in cost-benefit analysis: what are the analytical foundations on which such analysis is based? What are the main issues that have to be faced? What are the key variables that have to be calculated or estimated? And, finally, what determines the degree of confidence that one can have in the results?

Without a doubt, the best starting point is the ex post analysis of a pure business project. Here we are interested in its financial profitability over its effective lifetime. To estimate this, we build up a profile of the “cash flows” associated with the project during each year (or quarter, month, or other period) of its life. This profile would start with recording the outlays involved in designing and planning the project, and would then turn to the disbursements for construction and equipment purchases. Then would come the whole history of the years of operation of the project. For each year, we would record the net cash flow linked to that year. This would cover the sale and other possible sources of cash inflows. From these inflows we would deduct all cash outflows -- labor and materials costs, taxes, costs of maintenance, repair and replacement of capital equipment, insurance costs, etc. In the end we would record for each year, starting with the design and planning stage and ending with the closing down of the project -- the net cash flow (i.e., inflows minus outflows) corresponding to that year. The end result of all this is a project profile. Such a profile would typically start with one or more years of net outflows, followed by a operating period in which the net flows of each year were (typically but not necessarily) positive. This project profile summarizes the key facts needed as inputs into the analysis.

But wait!! A little reflection should alert one to an additional issue that must be addressed -- the phenomenon of price inflation (or, conceivably, deflation). We all have experienced important price level movements and are fully aware that the dollar, the peso, or the rupee of one year hardly ever has the same real purchasing power as that of other years. So if we want to avoid fooling ourselves about the economic worth of the project, we have to convert our project profile into real terms. Standard practice is to choose a numeraire -- usually either the country's consumer price index (CPI) or its GDP deflator, and to divide each year's net cash

flow by this numeraire, thus expressing it in real terms. If we are using the CPI, the resulting profile is expressed in consumer baskets -- how many net CPI baskets were expended during each of the investment years, and how many net CPI baskets were taken in during each of the productive years. If we use the GDP deflator as the numeraire, we equivalently express each year's outflow or inflow in terms of "producer baskets". That is to say, our resulting, inflation-adjusted profile would end up measured in dollars (or pesos, or rupees) of constant purchasing power, with that purchasing power being sufficient to buy a given bundle of either consumer or producer goods in each year of the project's life.

To analyze the profile further we need a discount rate. In principle this discount rate should reflect the genuine real opportunity cost of funds for the entity that is undertaking the projects. For a typical U.S. firm that real opportunity cost might be something like 8-10 percent per year. For businesses in most developing countries it would typically be significantly higher. It is important to recognize that in order to evaluate a profile expressed in real terms, one needs to utilize a discount rate that is also expressed in real terms. For a business firm one might consider the real opportunity cost of funds to be equal to a weighted average of its real cost of equity capital and its real cost of debt capital, the weights being the shares in which new investments by that firm typically financed by equity and debt.

The key thought that lies behind the idea of economic opportunity cost is that it reflects the true cost of raising the money used in this project and/or the true yield that the project operator would normally have obtained in the likely alternative use of that money. If these two numbers differ, then the relevant opportunity cost is the higher of the two.

We will here work through an example of a simple project of, say, a cattle-feeding operation. In period zero young animals are purchased for, say 1000. During the periods one and two, the maintenance and feeding costs are 300. In the period 3, the animals are sold for 2520. All these numbers should be thought of as being already expressed in real terms and as representing the net cash flows corresponding to each year. The profile of this project is thus -1000, -300, -300, +2520.

In the top panel of Table 1, a calculation is made of the net present value of this project, using an opportunity cost rate of 10% per year. The project starts out in period zero using

TABLE 1
Net Present Value and Internal Rate of Return

| | <u>Net Present Value</u> | | | |
|---|--------------------------------|--------------|--------------|--------------|
| Period: | 0 | 1 | 2 | 3 |
| Profile: | -1000 | -300 | -300 | -2520 |
| Net Present Value @ 10% | | | | |
| | -1000 | | | |
| | x 1.10 = | <u>-1100</u> | | |
| | | -1400 | | |
| | | x 1.10 = | <u>-1540</u> | |
| | | | -1840 | |
| | | | x 1.10 = | <u>-2024</u> |
| Net Present Value at 10% Calculated to Period 3 = | | | | +496 |
| | <u>Internal Rate of Return</u> | | | |
| Period | 0 | 1 | 2 | 3 |
| Profile | -1000 | -300 | -300 | +2520 |
| Net Present Value @ 20% | | | | |
| | -1000 | | | |
| | x 1.20 = | <u>-1200</u> | | |
| | | -1500 | | |
| | | x 1.20 = | <u>-1800</u> | |
| | | | -2100 | |
| | | | x 1.20 = | <u>-2520</u> |
| Net Present Value of Zero Confirms IRR = 20% | | | | 0 |

1000 of resources, thus “owing” the owner the sum of 1000. But since the appropriate opportunity cost rate of return is 10%, by period 1 this 1000 grows to 1100. Add to this the 300 of feeding costs in period 1 and the project now “owes” its owner 1400. This, in turn, grows (at 10%) to 1540 by period 2, to which must be added the 300 of feeding costs in that period, yielding a total of 1840 of what we call “capital-at-charge” at that time. This in turn again grows by 10%, to 2024 in period 3, when the sale takes place. This sale, for the sum of 2520, yields the owner a profit of 496 (= 2520 minus 2024), which is the present value (as of period 3) of the gain from the project, over and above (a) the 10% return that the owner could have gotten via a normal alternative investment or alternatively (b) the 10% real interest rate that the owner actually had to pay for the funds used in the project. If (a) and (b) are not the same, then our assumed 10% opportunity cost should be taken to represent the higher of the two.

Often one finds it difficult to pin down precise numbers for the opportunity cost concepts (a) and (b). In that case one can work with upper and lower bounds for the discount rate. This is an easy route to take when a project passes the profitability test at both the upper- and lower-bound discount rates, or when it fails to pass at either of these rates. But one must recognize that some projects will pass at the lower-bound rate and fail at the upper-bound rate, leaving unsettled the decision of whether or not to accept the project.

One statistic that can be useful in such situations, and in many others as well, is the internal rate of return of a project. The internal rate of return is like the yield-to-maturity of a bond. It tells you the precise rate of yield of the project, compounded over its whole life. Thus one can talk about a project A having a yield of 12%, project B having one of 9%, etc. The internal rate of return (IRR) is thus an extremely useful piece of information, but readers should be warned that the IRR should not be thought of as a determining criterion of project choice. That is, one should not, when given a choice between projects A and B, always choose the one with the highest IRR. (This should not detract from the general utility of the IRR as a summary statistic. After all, intelligence is a relevant factor in one's decision of whom to marry, but that does not say that from among her suitors, a girl should always choose the one with the highest IQ.)

A particular advantage of the internal rate of return is that it is an attribute of the project profile itself. It can be calculated directly from the profile data. And for this purpose one doesn't have to know what is the relevant opportunity cost of capital. That is why many boards of directors (including, for example, the executive boards of the World Bank and of most regional development banks, ask that the IRR be calculated for each project that is submitted for their approval. Once again, it is not a determining criterion,¹ but it is a very useful and informative statistic.

The lower panel of Table 1 illustrates the calculation of the internal rate of return of the cattle-feeding project. In point of fact one finds the IRR by trial and error, or more realistically finds a computer that is programmed for this purpose. In this case the example was created so as to have an "easy" IRR of 20%. Thus, the initial capital-at-charge of -1000 in period zero grows to -1200 by period 1, and is augmented by costs of -300, leaving a capital-at-charge of -1500 in period 1. This grows to -1800 in period 2 and is again augmented by -300 of feeding costs. The resulting capital-at-charge of -2100 in period 2 grows (at 20%) to -2520. This accumulated cost is precisely canceled by the sale price of +2520 in period 3, leaving a net present value of zero. This is how the IRR is defined -- it is that rate of return which, when applied to a given profile, yields a net present value of zero.

¹Suppose you are asked to choose between project A, with a profile of -100 + 130, and project B, with a profile of -1000 + 1200. The internal rate of return of A is 30% while that of B is only 20%. Yet if the relevant opportunity cost of capital is 10%, the net present value of B calculated as before to the closing period is 100 (= 1200-1100), while that of A is only 20(= 130-110). This is a very simple example of why the IRR is not a reliable sole criterion for project choice, but the principle it illustrates extends to a wide range of real-world projects. The principle states that the larger of two projects may be preferable to the smaller one, even if it (the larger) has a lower IRR -- because, in spite of its lower IRR, its net present value can end up being significantly higher than that of the smaller project. This principle assumes that the relevant choice is between only one of each of the alternatives, but that is the real situation in a great many cases -- a high dam (more expensive) versus a low dam (cheaper), a concrete highway (more expensive) versus a gravel road between the same points, a bridge to a nearby island (more expensive) versus a ferry project.

“Economic” Cost-Benefit Analysis

What makes the analysis of a pure business project so easy is the fact that the profile of benefits and costs is so easily defined. “Money coming in” is good, a plus; “money going out” is bad, or minus. Some public projects may be, effectively, business ventures that can be analyzed in this way. But this is certainly not the case for the vast majority of public projects. Here the benefits may take the form of a saving of travel time for road users, the improvement of farmers’ crop yields, a better-educated labor force, a healthier and longer-lived population, etc. Benefits of these kinds do not come as increased cash flows to the project, yet they are clearly benefits from the point of view of society, and should undoubtedly be taken into account in the evaluation of the projects that generate them. This reflects the broader vision that distinguishes an “economic” from a purely business or financial analysis.

Readers will easily appreciate how difficult it can be to place a monetary value on some of these non-cash benefits. Such valuation has been one of the major challenges that economic project (and program) analysis has had to face. It is a struggle that is ongoing, in which victories are hard to achieve, and typically only clear up small patches of a large and cloudy panorama. Highway benefits are easier to evaluate than those of irrigation projects; irrigation benefits are easier to quantify than education benefits; the latter, in turn, are easier to handle than health benefits.

In light of the above-mentioned difficulties, this introduction to the “economic” facet of cost-benefit analysis will focus on the most straightforward type of externalities -- namely, taxes and subsidies. These come in the form of cash, but they represent benefits or costs that accrue to the government rather than the project entity. Thus, for example, we can think of undertaking an economic cost-benefit analysis of a textile mill or a cattle ranch or a private electric power plant. The economic analysis would differ from the business or financial analysis of these firms by recognizing that the taxes they pay, while a cost to the firm, are a benefit to the government, and that any subsidies they receive, though properly counted as a benefit from the business point of view, have also to be counted as a cost to the government when the analysis is undertaken following the broader “economic” point of view.²

Typically, the economic profile of a project will show higher benefits than the corresponding financial profile. Most projects involve the use of materials and other inputs that

²Cost-benefit terminology has changed somewhat over recent decades. What we now call economic benefits and costs were formerly labeled “social” benefits and costs. As late as 1985, I entitled a major paper “Reflections in Social Cost-Benefit Analysis”, yet if I were to do it today, the title would refer to “economic” cost-benefit analysis. The reason for the change is that many people interpreted the term “social” to refer to the items usually dealt with in what are labeled social programs -- inoculations for children, unemployment benefit payments, government health care programs and subsidies, poverty relief, etc. These programs can be subjected to an economic cost-benefit analysis, alright, but there is nothing about economic C-B analysis that is in any way limited to such programs. It was in order to avoid inaccurate interpretations that our professional terminology has now for some time been using the term “economic” in places where “social” would earlier have been employed.

are internationally traded, and to which import tariffs and possibly other taxes apply. The financial analysis of such a project would count those taxes as a cost, thereby reducing the net flow of benefits that is calculated for the project. In the project's economic profile, such taxes are not counted as costs, leading to a higher flow of net benefits.³

The same story applies with respect to all taxes that are either directly paid by the project or embodied in the costs of items that the people purchase. This includes corporation income taxes, sales and excise taxes, franchise taxes, etc.

Economic Opportunity Costs

The concept of economic opportunity costs (sometimes called "shadow prices") is quite central to economic cost-benefit analysis. It is also a concept that is unfamiliar (and hence often quite puzzling) to many non-specialists. In this section we will try to convey the basic idea at an intuitive level, and with the least possible complications. We will consider here the economic opportunity costs of foreign exchange and of capital.

Consider a case in which a country has an average import tariff of, say, 20%. When our project goes into the market to buy foreign exchange, that foreign exchange will ultimately come from a combination of "displaced other imports" and "newly stimulated exports". Suppose that the project buys \$100 of foreign exchange, of which \$60 comes from displaced imports and \$40 comes from newly stimulated exports. Suppose, too, that the project is in a "peso" country, with an exchange rate equal to 10 pesos per dollar. The project pays 1000 pesos for the \$100 it buys, but this cost does not include the 120 pesos of tariff revenues that are lost as \$60 worth of imports are displaced. The total economic cost of the \$100 of foreign exchange is thus 1120 pesos, not the 1000 pesos paid by the project. Thus, the economic opportunity cost of a dollar of foreign exchange would be 11.2 pesos, not the 10 pesos that the project had to pay when it bought the dollars in the foreign exchange market. The extra cost of 120 pesos (= 1.2 pesos per dollar) is an externality, represented by the import tariff revenue that the government forgoes, as a consequence of the displacement of other imports, as our project enters the market to buy foreign exchange.

The concept of economic opportunity cost would be useless, for all practical purposes, if it did not refer to repetitive operations. But in point of fact the foreign exchange market is totally impersonal -- its participants do not know who is behind each \$100 increment in demand, nor do

³The easiest way to think of the adjustment is to note that the financial profile does not count the benefit to the government that these taxes represent. This benefit is an "externality" so far as the project itself is concerned, and therefore should be added as we pass from its financial to its economic profile. At a perhaps more subtle level, thinking all the time about the benefits and costs of a project from the point of view of society as a whole, the taxes are a transfer payment, being a cost to the project itself, and a benefit to the government. Thus if we start by counting benefits and costs other than taxes, we can proceed by simply not making any adjustment -- i.e., treating taxes as a transfer payment, a non-cost. When so treated they are a part of economic "profits", as they have not been deducted from the sales or other gross benefits of the project. But if we take the step of counting the taxes as a cost (in creating the project's financial profile), we then have to count them again, but as a benefit, when moving from the financial to the economic profile.

they know for what purpose that \$100 will be spent. All the market senses is the added pressure coming from \$100 of additional demand. Hence, so far as the market's reactions are concerned, of import displacement and export stimulation, every extra \$100 is like every other extra \$100. We must therefore consider that the division (60-40 in our case) of the effects of an incremental demand does not depend on who is the demander or the purpose for which the foreign exchange will be used -- it depends instead on the conditions of supply and demand in the foreign exchange market.⁴

Thus, the economic opportunity cost of foreign exchange can be applied to all foreign exchange transactions. In our case, any purchase of foreign exchange will represent an economic cost of 11.2 rather than 10 pesos per dollar and every act of generating new foreign exchange revenues (e.g., via incremental exports) will bring about an economic benefit of 11.2 pesos per dollar. This means that we can amalgamate all of a project's purchases and sales of foreign exchange into a single net dollar value, and make the necessary adjustment (moving from the project's financial to its economic profile) by augmenting the net dollar cost or the net dollar benefit by the "foreign exchange premium" of 12% (in our example).

If the operations involved were not repetitive -- with all of the increments to demand for foreign exchange having a similar effect, this whole simplifying exercise would not be possible. We would then have to treat each act of buying foreign exchange as a separate exercise, finding for each of them its own fraction of import displacement and export stimulation, and maybe even end up with different purchases of foreign exchange displacing different imports with different tariff levels, etc. This would be an utter nightmare for the analyst, but we are lucky -- instead of calculating the economic opportunity cost of foreign exchange (EOCFX) 1000 different times for 1000 different projects, we only have to calculate it once, for all of them. This is ideally, therefore, a job for each country's project authority itself, or perhaps for an international agency like the IMF or the World Bank to do for small countries with limited technical expertise. Since EOCFX is an important parameter, and since it only has to be calculated once for each time period, it pays for those responsible to do a careful, professional job when carrying out such a calculation.

Readers should recognize that the EOCFX applies both to the acquiring of foreign exchange that the project will then spend, and to the disposition of foreign exchange that the project may have generated. It deals thus with half-a-picture, not the whole picture of a foreign exchange operation. The other half of the picture concerns how the project spends the foreign exchange that it buys, and how the project generates the foreign exchange that it sells. This part of the story is very clearly not repetitive. One project may use its foreign exchange to buy airline tickets, another to import wheat (with a zero tariff) and yet another to import a BMW car (with a 50% tariff). The tariffs and taxes paid in these project-specific operations have to be counted (as indicated in the previous section) as external benefits and costs of the project. The foreign exchange premium (or EOCFX) does not come into play on this side, but the project-specific tax and tariff externalities do operate as an offset, or balance wheel, to the extra costs involved when

⁴The simple, traditional example here is that where the elasticities of imports demand and export supply are the same, the division is 50-50. If the elasticity of import demand is twice that of export supply the division will be 2/3 vs. 1/3. If the export elasticity is twice that of imports, the division will be 1/3 vs. 2/3.

the foreign exchange premium is applied. Thus, in our example, we assigned a foreign exchange premium of 120 pesos to the \$100 our project bought, but when that foreign exchange was used it might have brought an extra benefit of 200 pesos (if the imported items carried a tariff of 20%) or of 500 pesos (if the tariff on this item was 50%) or of zero (if the imported item entered duty free).

It is easy to see that there is nothing repetitive about the spending of foreign exchange that is bought by the project. But fortunately, project analysts should have readily at hand the project's planned purchases, so the "other half of the picture" is typically relatively simple to compute, where project imports are concerned.

The export side should be even easier, for most projects generate no export revenue at all. Those projects that do generate export receipts will typically have only one or two or at most a very few export products, most of which will carry no export tax or export subsidy. In such cases no second-side adjustment has to be made. One simply applies the foreign exchange premium as an extra benefit on the foreign exchange that the project generates.⁵

Before leaving the subject of the foreign exchange premium, we should stop to explain why we say that each purchase of foreign exchange is in some fundamental sense "sourced" either from displaced imports or newly-stimulated exports. This principle is derived from the idea of exports being the main source of a country's foreign exchange earnings, and imports being the main use to which these earnings are put. Obviously, if these are the only source and the only use, the principle becomes a virtual tautology. In point of fact, however, borrowing can be thought of as an additional source of foreign exchange for a country. But -- and here's the rub -- one should operate on the presumption that such debts will be repaid, with interest. Thus in a present value sense one is still driven back to the notion that sooner or later the foreign exchange that we extract from the market will be reflected in lower imports and more exports than would otherwise appear. And the same thought works in reverse -- if today's export earnings are not directly reflected in more imports and less other exports, that means they are being lent or invested abroad. But when these loans are repaid with interest or those investments are repatriated along with the dividends they have accumulated, they will at that time be covering an excess of the country's imports over its exports that would not otherwise exist.

Of course there remains the possibility that gaps between imports and exports will be covered by foreign aid, by charitable donations, or by remittances from emigrants. These are recognized parts of the total picture, but they are rightly considered to be determined quite independently from any given project's purchases or sales. Thus they are not considered to play any role in covering a project's demand for foreign exchange or in absorbing its supply. This, I hope, helps explain why we are so adamant about insisting that a project's effects with respect to foreign exchange take place in the world of imports and exports, and not somewhere else.

⁵The way this would work, following our example, would be that as \$100 of export receipts is sold on the foreign exchange market, \$60 of extra imports are stimulated, and \$40 of other exports are displaced. On the \$60 of newly-stimulated imports the average 20% tariff would apply, so there would be an external benefit of 120 pesos, to be added to the 1000 pesos received from selling the \$100 in the foreign exchange market.

Having come this far, we still have to take one further important step before this exposition of the foreign exchange premium is complete. This next step entails shifting our focus from a project's direct purchases and sales of foreign exchange to that project's purchases and sales of internationally traded goods and services. We start with the notion that the prices of internationally traded goods and services are fundamentally determined in the world marketplace. With few exceptions, individual importing and exporting countries have little or no influence on the world prices of the tradables that they buy and sell. This has an important implication for our analysis. Argentina is an exporter of beef, so if a project located there buys beef in the local market, that means that Argentine exports of beef will fall by an equivalent amount. Thus, Argentina's export earnings will fall, leading to reduced imports and newly stimulated other exports to fill the gap.

Similarly, the U.S. is an importer of copper. So even if a U.S. project buys additional copper from the U.S. producer located in Montana, that additional purchase of local copper will mean that a like amount will end up being imported (at the given world price) by other domestic users of copper. That is, it will have the same effects on the foreign exchange market as would have occurred if the project had gotten its copper via direct imports.

The bottom line of all of this is that we should apply the foreign exchange premium to the project's purchases or sales of all internationally traded goods and services -- not just to its purchases of imports and sales of exports. This extension conforms to the lessons of modern open-economy macroeconomics, which focuses sharply on the dichotomy between "tradables" and "nontradables". Tradables in turn are broken down into "importables" and "exportables". And, finally, imports represent the excess of demand for importables over their supply, and exports represent the excess of the supply of exportables over the demand. From these simple relationships we can easily derive that if a country's exports exceed its imports by a given amount, its supply of tradables will exceed its demand for tradables by that same amount. And similarly, when a country's imports exceed its exports, precisely the same gap will exist between its total demand for tradables and its total supply of them. This is the basic foundation for the rule, in cost-benefit analysis, that the foreign exchange premium should apply to a project's total demand for tradable goods and services and to its total supply of them

Implicit in the above but not yet mentioned, the foreign exchange premium should also apply to any direct purchase of foreign currency for investment abroad and to any dividend or interest or capital repatriation flows that the project may receive as net inflows of foreign currency, to be sold on the national foreign exchange market.

The Economic Opportunity Cost of Capital

For simplicity, we will start this exposition by assuming what we call a "closed economy", without international trade or capital movements. Later, of course, we will extend the analysis to cover the real-world case of the open economy. Most readers will have learned, in some early economics course how, when the accounts of a year are drawn up for a national economy, it has to be true that its investments have to equal its savings. The total of what is produced ends up as either consumption or investment, since national savings are defined as total production minus total consumption.

We build on the equality of saving and investment to draw the conclusion that just as a demand for foreign exchange has to come for either displacing other imports or generating

newly-stimulated exports, so in the capital market a project's demand for funds must come either from displaced other investments or from newly-stimulated savings. The general picture is just the same, whether we are talking about the foreign exchange market or the capital market. It's just that in the one case the sources of funds are displaced imports and new exports, while in the other case the funds obtained from the project come from displaced investments and from new savings.

In the early years, when formalized work on cost-benefit analysis was just getting underway, the most common procedure was to consider that the alternative to "this" project's (the one being studied) use of capital funds was a "standard" investment in the rest of the economy. The investment "here" was thought to be a good one if the capital used "here" would be more productive than the same capital would be if used "there". This way of thinking naturally led to a number of attempts to measure the economic productivity of capital in the overall economy, or in its private sector, or in its "business" sector.

As thinking in this area evolved, it was soon realized that real-world mechanisms for raising money for a project typically entailed drawing those funds at the expense of both consumption and investment. In a capital-market model, one key element was a demand curve representing the demand for funds for investment in the rest of the economy (not counting "this" project). That demand curve was then juxtaposed to a supply curve of funds, both demand and supply being expressed as functions of the rate of interest. When analyzed in this way, the extraction of project funds from the capital market had two sources -- some of the funds came, in effect from displacing other investments, while the remainder came from moving up along the savings curve -- i.e., stimulating additional savings in the economy. Obviously, additional savings means reduced consumption so that project funds were seen coming part from displaced investment and part from displaced consumption.

That same dichotomy applied in the case where project funds were viewed as having come from incremental tax revenues. Tax money also comes from somewhere!! Investment and consumption were again the relevant sources, but of course the proportions in which they would come at the expense of consumption and investment would be different when the money came from taxes vis-a-vis when it came from the capital market.

The thought of considering taxes to be the relevant source of funds for a public project has a lot of immediate appeal, since most governments raise most of their money through taxes. Yet this pathway led quickly to a swamp -- different taxes had quite different effects on consumption and investment (consider personal income versus corporate income taxes, or capital gains taxes versus excise taxes, or estate or inheritance taxes versus consumption taxes). Moreover, the changes in tax laws passed in one year bear little relation (in terms of their effects on consumption and investment) to those that passed in other years. All of this led to the conclusion that there is no "standard" way of extracting additional resources via taxation.

Things look much better when one considers the capital market as the source of funds. Here we see that any new demand for funds just adds an additional demand to the original picture juxtaposing investment demand against the supply of savings. That market, like the foreign exchange market mentioned earlier, is fundamentally impersonal, it "feels" additional pressure as new demands are added, and it reacts to that pressure -- not to the specific purpose for which these funds will be used. This gives us a good reason for treating the capital market as

the standard source of project funds, but that would be just wishful thinking if it bore little or no relation to reality.

Luckily, we have a sound basis for treating the capital market as the “standard” source of project money. For the capital market is in fact the “sponge” that absorbs unexpected extra tax revenues as they appear from month to month, and that is indeed the source to which governments turn when they face a shortfall of monthly revenues vis-a-vis their budgeted outlays and/or budgetary overruns. This reality simply confirms that the capital market is indeed the marginal source of government funds.

Our picture, then, is one of taking funds out of the capital market, displacing investment and stimulating savings in the proportions that are determined by the supply-and-demand picture in that capital market. Fortunately we know quite a bit about how the demand for investment funds and the supply of savings respond to pressures in the capital market. Briefly, we know that the demand for investment funds is much more elastic than the supply of savings. It is quite within the range of plausibility that incremental funds drawn from the capital market would be divided 90-10 between displaced investment and newly-stimulated savings. On the other hand, it would stretch plausibility a bit to assume that this division was 75-25 rather than 90-10. Thus, the old literature that assumed that 100% of the funds came at the expense of investment was not absolutely right, but was pretty close in quantitative terms.

A simple example may help readers see the essential structure we are working with. Assume that 1000 of funds are raised in the capital market, 750 of which come at the expense of displaced investment, and 250 of which derive from newly-stimulated savings. Assume, too, that the displaced investments would have yielded a 12% real rate of return, and that the people doing the extra saving require a 4% rate of return (i.e., 4% is their supply price of saving, at the margin). Under these circumstances we would estimate the economic opportunity cost of capital (EOCK) at $(.75)(12\%) + (.25)(4\%) = 10\%$.

This gives the basic idea behind the EOCK but it does not go deep enough. Readers should realize that, just because we pulled that 1000 out of the capital market the economy has really lost the product that the 750 of displaced investment would have produced. At 12%, that amounts to a stream of lost benefits equal to 90 per year, for an extended future period. This 90 per year can be seen as a “debt” that our project “owes” to the economy -- if that debt is not “paid”, our project is not justified. Likewise, the 4% return (required by savers to leave them just barely indifferent versus not having saved an extra 250) is another “debt” our project has to “pay” in order to be worthwhile. This amounts to an added 10 per year.

Taken together, the 90 plus 10 per year is a challenge to the project. If it produces net benefits of exactly 100 per year, it will have just barely covered the true costs involved when the original 1000 was raised in the capital market. Only after these costs have been covered can it be said that our project was economically better than the standard alternative that would have emerged, had the 1000 of funds simply been left in the capital market (i.e., had we simply not embarked on the project).

In these circumstances, the whole story just told is automatically accomplished simply by using a 10% discount rate in analyzing the economic profile of the project. If this process yields a yet present value that is positive, the project is worthwhile, if negative, the project is not worth doing based on strictly economic criteria. If there are two projects (A and B) that are

alternatives, so that only one of them can appropriately be carried out, then the choice should favor the one that yields the higher net present value, using a discount rate of 10%.

On Drawing Funds From Abroad

Our discussion so far has implicitly assumed a closed economy, in that we have only considered displaced domestic investment and newly-stimulated domestic savings as the ultimate sources of the money we extract from the capital market. In short, what we have done so far is represented below:

$\rho =$ marginal productivity of domestic investment = 12%

$r =$ supply price of domestic savings = marginal rate of time preference = 4%

$f_1 =$ fraction of project funds coming at the expense of displaced investment = $(750/1000)$
= .75

$f_2 =$ fraction of project funds coming from newly-stimulated domestic savings = $(250/1000)$
= .25

EOCK = economic opportunity cost of capital

$$= f_1\rho + f_2r$$

$$= .75(12\%) + .25(4\%) = 10\%$$

Now we will make the exercise more realistic by incorporating the linkage to the world capital market. The essence of this operation can be easily seen. We now might have:

$f_1 =$ fraction of project funds coming at the expense of domestic investment = $(700/1000)$
= .70

$f_2 =$ fraction of project funds coming from newly-stimulated domestic savings = $(100/1000)$
= .10

$f_3 =$ fraction of project funds represented by a net increase in “foreign savings” coming to the country in question = $(200/1000) = .20$

We keep: $\rho = 12\%$, as before.

$r = 4\%$, as before.

and we add:

MCFF = marginal cost of foreign funds = 8%.

Under these assumptions we have:

$$\text{EOCK} = f_1\rho + f_2r + f_3\text{MCFF}$$

$$= .7(12\%) + .1(4\%) + .2(8\%) = 10.4\%$$

The marginal cost of foreign funds is very difficult to estimate, and so too is f_3 , the fraction of incremental funds coming from abroad. What we know in this regard is:

- a) the capital market linkage of a country to the rest of the world is not “perfect”. The country can normally draw additional funds from abroad, but not at the same price. The supply curve of foreign funds available to a given country is not infinitely elastic -- instead, it has an upward slope.
- b) The upward slope of any supply curve means that the marginal cost (of getting an extra unit as one moves along that supply curve) is higher than the supply price.⁶

When considering foreign sourcing of funds we have to recognize that the behavior involved is not as straightforward as the domestic demand for investment and the supply of domestic savings. That is to say, the supply of “foreign savings” shifts up and down more than that of domestic savings, and its elasticity is also less easy to pin down to a narrow range. But, at least for many countries, it would be totally inappropriate to neglect the supply of foreign savings when trying to estimate the economic opportunity cost of capital.

In dealing with these uncertainties one looks for ways, at least of pinning down some limits for the marginal cost of foreign funds. The best approach that I know of is to try to estimate the average cost (in real terms) of foreign funds to the country in question. This number should incorporate both equity and debt financing, and should be expressed in real terms. Suppose this rate is 6%. One then estimates the corresponding “riskless” rate in the world capital market. Suppose this rate is 4%. Under these circumstances a plausible band for the MCFF would be from 8 to 12 percent. The 8% figure is found by taking the excess of the country’s average rate (6%) over the world riskless rate (4%), and adding this excess (2%) to the average rate (6%). The 12% figure is obtained simply by doubling the country’s average rate (6%).⁷

In advanced countries like the U.S., Canada and Western Europe, estimates of the economic opportunity cost of capital have ranged around 8 to 10 percent. This rate can be quite a lot higher in developing countries, particularly those with vibrant economies and ample investment opportunities. One guidepost can be derived from the policies and practice of the

⁶From standard economic textbooks:

$$\begin{aligned}d(pq) &= pdq + qdp \\d(pq)/dq &= p + q(dp/dq). \\d(pq)/dq &= p(1+1/\varepsilon)\end{aligned}$$

$$\text{marginal cost} = \text{average cost times } \left(1 + \frac{1}{\varepsilon}\right),$$

where ε = elasticity of supply.

Since the supply of foreign funds is upward sloping, ε is positive, and the marginal cost of funds must be greater than the average cost.

⁷Both of these figures follow from assuming a tangent line to the average cost curve (at its current equilibrium position) equal to $p = a + bq$. Marginal cost is then equal to $p + bq$. If a equals the riskless rate, then $bq = p - a$ (2% in our example, and this has to be added to the average cost (6%) in order to get marginal cost (8%)). If the tangent line goes through the origin, then a is zero, and we have to add another “ p ” to the average cost, yielding MCFF = 12%.

World Bank over its 60-odd years of history. That organization has faced a tricky problem in the sense that it does not seem “correct” for such an international organization to insist on a 14% criterion in one country while accepting, say, an 8% criterion in another. Probably motivated more by this dilemma than by purely technical considerations, the World Bank has for decades used a real rate of 10% per annum as the “standard criterion rate” for judging whether a project merits a World Bank loan. This by itself does not reveal much about the true opportunity cost of funds in different countries around the world, but the actual experience of World Bank projects can be generally instructive. We have the fact that the World Bank has faced no shortage of projects, over the years, that were able to pass the 10% test *ex ante*. Moreover, the World Bank has done extensive examinations *ex post*, in order to assess whether its previously-financed projects had performed up to expectations. These reassessments have regularly been quite favorable, with many projects having estimated real rates of return (*ex post*) of 15% or more.

All of this suggests that it is not a mistake for people to think of 10% as a sort of plausible benchmark real rate of return that one should normally expect public sector projects to achieve.

A particular problem sometimes emerges when a proposed project has the prospect of outside financing at a very favorable rate. The accepted wisdom on this topic is that financing at a cheap rate should not be taken as justification for accepting projects of correspondingly low expected benefits. The country really gains when it generates a 10% yield, when the project is fortunate enough to be financed by a 2% loan. Obviously, the country would not gain if the total benefits of such a project only amounted to a 2% real return.

More important in this area is the general “fungibility of funds”. Certainly, it is better to use cheap money for good projects than for bad ones. But it is highly likely that doing a bad project using cheap money will in fact end up squeezing out good projects that could otherwise be financed. Cheap money is almost certain to be “inframarginal” -- to be located in the “early” reaches of the supply curve of funds to the country,. Its most natural economic use would thus be for the most productive projects, not for ones of patently low rates of net return. In short, if a lender expresses willingness to provide money at 2%, that should always be welcome. But if that lender tries to convince you to invest that money in projects that yield only 2 or 3 percent, every effort should be made to display the availability and to argue the desirability of higher-yielding projects reflecting at least the true economic opportunity cost of capital for the country.⁸

⁸The whole exercise of calculating the EOCK can be regarded as a 2-way street. If going into the capital market generates costs which can only be covered by a project yield of, say, 10.4%, then the simple act of dumping new funds into the country’s capital market will have an economic yield of 10.4%. If nothing else, then, money borrowed at a rate of 2% could buy a big benefit to a country by the simple act of using those funds to feed the country’s capital market.